Optimal Pooling of Performance Information in Dynamic Contracting

Rui R. Zhao*

January 18, 2008

*Department of Economics, University at Albany - SUNY, Albany, NY 12222. E-mail: rzhao@albany.edu.

Abstract

In a long-run principal-agent relationship the principal may be able to draw more precise inference on the agent's performance by pooling observations across multiple periods. This paper studies the factors that determine the optimal pooling of performance information, including the verifiability of evaluation, the risk attitude of the agent and whether the agent faces limited liability constraints. Pooling benefits the principal by minimizing the *wedge* between the payoff of the principal and the reservation utility of the agent; the wedge may be a rent accrued to the agent or a deadweight social cost or the sum of the two. The specific wedge determines the effectiveness of pooling and the form of optimal compensation policy. Under alternative informational structures, optimal contract may display a variety of features such as deferred compensation, bonuses, deferred punishment, and sticky wages.

1 Introduction

In a long-run principal-agent relationship the principal may be able to draw more precise inference on the agent's performance by pooling observations across multiple periods. In this paper we study the factors that determine the optimal pooling of performance information, by identifying two functions of a review phase that spans multiple periods. The first is an *informational function*. The key is whether there is any gain from gathering information over a longer period of time, which in turn depends on *what* information need be gathered for designing incentive contracts. The second function is *consumption allocation*. By allocating wage payments across periods, the principal may take advantage of intertemporal substitution and economizes incentive provision. The interaction between these two functions depends on several factors central to dynamic contracting, including the verifiability of performance evaluation, the risk attitude of the agent and whether the agent faces limited liability (or wealth) constraints. In some cases there is perfect harmony between the two functions and the longer the review phase, the better off the principal. In other cases, consumption allocation may interfere with information gathering and the principal has to strike a balance between these two functions; then there can be an optimal review length.

The main findings of this paper are summarized as follows. If performance information is *publicly* verifiable (Section 2), whether the agent earns rent becomes central to whether there is any benefit from pooling information across periods. When the risk-neutral agent faces limited liability and hence must be incentivized by rewards, the principal can minimize the rent paid to the agent by pooling performance information across periods. The idea can be explained as follows. Suppose the principal divides the total duration into two separate review phases. To minimize incentive cost, in each review phase the principal should only reward the agent when she sees an output signal that is most likely if the agent chooses the principal's preferred action rather than something else. In particular, in each review phase i every (expected) dollar cost spent on rewarding the agent brings the principal an incentive benefit equal to $1-r_i$, where $r_i \leq 1$ is the minimum value of a likelihood ratio between defection and conformity. By pooling information across two review phases, the principal can raise the benefit-cost ratio from $1-r_i$ to $1-r_1r_2$. Therefore, there is increasing return in incentive provision as the review phase grows longer, and the principal can efficiently motivate the agent by evaluating and rewarding him until the end of the relationship. Deferred compensation is a necessary feature of the optimal wage policy.

By contrast, when the risk-neutral agent does not face limited liability, information pooling and deferred compensation have no effect and the principal can treat each period separately without loss. When the agent earns rent but is risk averse, it becomes increasingly costly, due to inefficient consumption smoothing, to motivate the agent using deferred rewards farther into the future. Then an optimal review length exists.

If performance information is *privately* held by the principal (Section 3), the roles of risk aversion and limited liability are completely reversed. To solicit the principal's private information about the agent's performance, the principal is essentially required to put aside a fixed amount of money and the mechanism freely allocates the money to the risk-averse agent across periods for both insurance and incentive purposes. Due to her own incentive problem, the principal pays a deadweight social cost in addition to what is received by the agent. Without limited liability, there is perfect harmony between consumption smoothing and information pooling across periods, and as a result the longer the review phase the better off the principal.

The intuition can be illustrated as follows. If the principal divides the entire duration into two review phases, then in each review phase every unit of social cost the principal pays brings an incentive benefit equal to $r_i - 1$, where $r_i \geq 1$ is the maximized value of the likelihood ratio of signals between deviation and conformity. The benefit-cost ratio when compensating the agent based on pooled signals is at least $r_1 r_2 - 1$, which is larger than each $r_i - 1$. Therefore, to minimize the social cost, optimal contracts should pool the agent's performance information over the entire duration. At the same time, since the social cost need only be paid at the end of the relationship, the principal can smooth the agent's consumption and pay the same wage until the last period, which is perfectly consistent with insuring the risk-averse agent against consumption shocks. Optimal compensation policy thus features deferred punishment and *sticky wages* over time; the stickiness is especially robust as it is derived under arbitrary asymmetries in productivities across periods.

Finally, if evaluation is private and the agent is protected by limited liability (Section 4), then deferred punishment won't work and the principal can not always gain by pooling performance information across periods. In the previous two cases, the principal benefits from pooling because it reduces either the agent's rent or the social cost. But in the current situation the principal needs to minimize the *sum* of the agent's rent and the social cost. For this purpose, the principal should reward any signal that indicates conformity and punish (by withholding rewards) any signal that indicates deviation, not just the ones that are most informative. Unfortunately, different pooled signals are better at detecting different deviations and an order of informativeness may not exist among pooled signals even if it exists for signals within individual periods. As a result, pooling may not always provide the informational gain needed for more efficient incentive provision and hence may not benefit the principal.

In summary, pooling works to minimize the *wedge* between the payoff of the principal and the reservation utility of the agent; the wedge may be a rent accrued to the agent or a deadweight social cost or the sum of the two. The specific wedge determines the effectiveness of pooling and the forms of optimal compensation policy.

Literature review: The basic idea that repetition generates better inference is important in reality and certainly is not new in the literature. In some studies on this issue the principal uses law-of-large-number type arguments to generate better inferences on the performance of the worker (e.g. Radner, 1985). The reliance on limiting arguments, however, provides limited insight on the structure of optimal contracts when the number of periods is not very large, an arguably relevant case in reality.

In a seminal paper, Abreu, Milgrom, and Pearce (1991) study repeated partnerships with imperfect monitoring and show that when signals are public but are observed by agents with a t-period lag, the average incentive cost per period due to signal imperfections are reduced to the order of $1/t$ of the original level when there is no reporting lag. Their ingenious idea, equally applicable to a small number of periods, is to "reuse" a single punishment at the end of a t-period review phase to enforce action choices across all t periods, which economizes the incentive cost compared with reviewing performance on the basis of individual periods.[1]

In an interesting paper, William Fuchs (2007) applies the idea of Abreu, Milgrom, and Pearce to repeated principal-agent contracting with private evaluation, a model first introduced and analyzed by Jonathan Levin (2003) [MacLeod, 2003, extends Levin's analysis to a risk-averse agent in a static model]. Fuchs finds that the principal can benefit from delaying the release of the agent's performance evaluation, and more importantly derives a justification for the optimal use of efficiency wage and termination in the infinite horizon setting.[2]

[1] Radner (1985) and Holmström and Milgrom (1987) are precursors studying dynamic principal-agent problems. Compte (1998) and Kandori and Matsushima (1998) extend the idea of Abreu, Milgrom and Pearce (1991) to repeated games with imperfect private monitoring.

[2] We discuss further the relations and differences between the present paper and these other treatments in Section 3, where private evaluation is considered.

The contribution of this paper is to demonstrate the benefit of optimal information pooling in broader settings and to delineate the factors that influence the effectiveness of pooling; by so doing it shows how pooling can generate rich features of optimal contracts. From the perspective of the current framework, the studies of Abreu, Milgrom, and Pearce (1991) and Fuchs (2007) both focus on the use of punishment as incentive devices and in both papers the benefit of pooling is driven by a specific wedge, the social cost, between the payoffs of the principal and agent; also, the results in these papers are derived in stationary environments. The present paper poses the problem in general nonstationary settings and under several alternative informational structures that are common in dynamic contracting.[3] This generalization clarifies and unifies the different ways in which informational gain can be achieved through pooling and it demonstrates the different features that optimal contracts may display as a result. Specifically, as explained in the above, optimal contracts are heavily influenced by a set of factors central to wage contracting; depending on the specifications of these factors there may be either deferred rewards or deferred punishment; sticky wages over time may occur for quite different reasons: due to consumption smoothing in the private evaluation case and as a result of deferred compensation in the public evaluation case; moreover, the reusability of rewards or punishments suggested by Abreu, Milgrom and Pearce (1991) hinges upon the way in which informational gain is realized through pooling and it may not have any benefit when there is no such informational gain.

The remaining sections analyze public evaluation, private evaluation, and private evaluation with limited liability, in that order. Proofs omitted in the text are given in the Appendix.

2 Public Information and Deferred Rewards

2.1 The Model

A principal hires an agent to operate a project for $t > 1$ periods. At the beginning of date $\tau = 1$, the principal offers a contract to the agent - its details will be spelled out below. If the agent accepts the offer then he works for the principal throughout the t periods. In particular, in each period τ the agent privately chooses an effort level a_τ from some finite set A_τ, which generates some benefit or output y_τ for the principal. Output y_τ is stochastic and is drawn from some finite set $Y_\tau \subset \Re_+$

[3] Nonstationarity however makes it difficult to study the problem when the time horizon is infinite. See Abreu, Milgrom, and Pearce (1991) and especially Fuchs (2007) for an analysis of the infinite-horizon situation.

according to a probability distribution $p_\tau(y_\tau|a_\tau)$, with $p_\tau(y_\tau|a_\tau) > 0$ for all $\tau = 1,...,t$, all $y_\tau \in Y_\tau$ and all $a_\tau \in A_\tau$.

The agent and principal both maximize sum of expected discounted utilities, and their period-τ utilities are given respectively by $w_\tau - g_\tau(a_\tau)$ and $y_\tau - w_\tau$, where w_τ is the wage payment and $g_\tau(a_\tau) \geq 0$ is the effort cost.

There is a lower bound on wage payment, which is normalized to zero: $w_\tau \geq 0, \forall \tau$. This may be a legal minimum wage, an ongoing market wage, or a historically determined wage level in the organization. Whatever the reason, the effect of such a wage floor is that the principal can only count on bonuses as incentive devices. The agent's reservation utility at the time of signing the contract is also normalized to zero.

The principal's objective is to implement her most preferred effort levels in the t periods with minimum cost. Specifically, in every period τ there is an effort level $e_\tau \in A_\tau$ that the principal would like to implement.[4] To make the problem nontrivial, assume that for all τ, $g(e_\tau) > 0$ and there is a least costly action $\underline{a}_\tau$ such that $g(\underline{a}_\tau) = 0$.

Assumption 1. For all τ, there is an output signal $z_\tau \in Y_\tau$ such that for all $a_\tau \neq e_\tau$:

$$\frac{p(z_\tau|e_\tau)}{p(z_\tau|a_\tau)} \geq \frac{p(y_\tau|e_\tau)}{p(y_\tau|a_\tau)}, \text{ for all } y_\tau \in Y_\tau, \text{ with strict inequality for some } y_\tau.$$

This assumption is satisfied generically if A_τ contains just two effort levels (*work* and *shirk* for example). It is slightly more general than the standard monotone likelihood ratio property, under which signal z_t is the highest output level.

An immediate consequence of Assumption 1 is $p(z_\tau|e_\tau) > p(z_\tau|a_\tau), \forall a_\tau \neq e_\tau$.

The principal can choose when to reveal the agent's performance information y_τ and in this section the information is verifiable by a court upon its release.

Let $y^\tau = (y_1,...,y_\tau)$ be the sequence of output realizations up to date τ; other sequences are similarly defined. A contract then specifies a plan of wage payment $w(y^\tau) \geq 0$ for all y^τ and all $\tau = 1,...,t$ and a plan for the principal to release performance information.

Given a wage plan $w(\cdot)$ and an information disclosure plan, at each date τ and given a wage path $w^\tau \equiv (w_1,...,w_\tau)$, the agent forms a belief $\beta_\tau(y^\tau\,|\,w^\tau)$ over the set of all possible τ-period output histories y^τ. For instance, if wages are independent of output realizations and the principal

[4] As an example, the agent may be a CEO, divisional manager, or research head, whose best effort is critical for the benefit of the principal.

does not reveal any information, then the agent's belief is the same as distribution $P_\tau(y^\tau|e^\tau) = p_1(y_1|e_1)\cdots p_\tau(y_\tau|e_\tau)$. To induce the agent to choose effort sequence e^t, a contract needs to satisfy the *sequential incentive constraints* (SIC): given any wage path w^τ, $\tau < t$, and given the belief $\beta_\tau(y^\tau \,|\, c^\tau)$ and the continuation wage plan, the agent should find it optimal to choose the continuation effort sequence $(e_{\tau+1},...,e_t)$.

The complex plans for information release are of no use to the principal. Intuitively, letting the agent know more about his past performance beyond what is conveyed by the wage plan only serves to add more history-dependent sequential incentive constraints to the principal's problem without any extra value. This consideration leads to the following relaxed program, which ignores the (SIC) constraints altogether.

Program (P):

$$\max_{w(y^t)} \; -\sum_{y^t} \delta^{t-1} w(y^t) P(y^t|e^t)$$

subject to:

$$w(y^t) \geq 0, \quad \forall y^t$$

$$e^t \in \arg\max_{a^t \in A^t} \sum_{y^t} \delta^{t-1} w(y^t) P(y^t|a^t) - G(a^t), \tag{1}$$

where $P(y^t|a^t) \equiv p(y_1|a_1)\cdots p(y_t|a_t)$ and $G(a^t) \equiv g_1(a_1) + \cdots + \delta^{t-1} g_t(a_t)$ is the agent's total effort cost discounted to date one. In this program all wages are paid at the end of period t and the principal minimizes the expected wage payment subject to the only incentive constraint that effort sequence e^t is optimal for the agent at date one.

If $\tilde{w}(\cdot)$ is any wage plan that in conjunction with some information disclosure plan satisfies the sequential incentive constraints, then by postponing all interim wages to date t with proper discounting we obtain a wage plan $\hat{w}(\cdot)$ that satisfies (1). Therefore, all feasible wage plans in the original problem remain feasible so the principal's optimal cost will not go up in Program (P).

Given an optimal solution $w^*(\cdot)$ to Program (P), if the principal does not divulge any information about output history until the end of date t, then the agent will not have any incentive to deviate from e_τ in any period τ. This is because when the agent deliberates at time $\tau > 1$ after having chosen the required efforts in the past, the only difference with time 1 is that all benefits and costs are discounted less heavily, which amounts to multiply all payoffs by a constant and therefore does not affect his original incentives. Then $w^*(\cdot)$ together with the no-early-disclosure policy solves the original problem. We summarize this observation in the following lemma.

Lemma 1. *The principal can restrict attention to plans that withhold performance information and pay the agent until the end of date t.*

2.2 Deferred Rewards

The interesting questions are whether the principal is *strictly* better off by deferring the release of performance information and what the optimal wage contracts look like. We first answer the latter in the following proposition.

Proposition 1. *There is a solution to Program (P) of the following form:*

$$w(y^t) = \begin{cases} W^* \equiv \frac{G(e^t)-G(b^t)}{P(z^t|e^t)-P(z^t|b^t)} & \text{if } y^t = z^t \\ 0 & \text{if } y^t \neq z^t \end{cases} \tag{2}$$

for some effort sequence $b^t = (b_1, ..., b_t) \neq e^t$.

The argument for this result, which illustrates how incentives are provided efficiently in the repeated setting, can be sketched as follows. Consider a collection of problems that further relax Program *(P)*. In every such relaxed problem, the principal minimizes the expected wage payment subject to the constraint $w(y^t) \geq 0$, $\forall y^t$, and a single incentive constraint ensuring that the agent prefers e^t to a specific effort sequence a^t.

We may, without loss of generality, focus on effort sequences a^t with $G(a^t) < G(e^t)$.[5] To motivate the agent to choose e^t over a^t, rewards should be given to him when there is evidence that he has chosen the former. If the principal pays the agent one dollar when observing a signal y^t, her expected cost equals $P(y^t|e^t)$ and the benefit of incentivizing the agent equals $P(y^t|e^t) - P(y^t|a^t)$. By Assumption 1, the cost-benefit ratio is minimized when the signal $y^t = z^t$. Therefore an optimal contract is to reward the agent only when z^t is observed at the end of period t with a bonus equal to

$$W(a^t) \equiv \frac{G(e^t) - G(a^t)}{P(z^t|e^t) - P(z^t|a^t)}.$$

[5]If $G(e^t) \leq G(a^t)$ then the solution is trivial: $w(y^t) = 0$, $\forall y^t$. Namely, the low-cost effort sequence e^t can be implemented without any incentive pay.

Define $W(a^t) = 0$ if $G(a^t) \geq G(e^t)$ and let $W^* = \max_{a^t} W(a^t)$ and $b^t \in \arg\max_{a^t} W(a^t)$. Then we obtain the solution to Program *(P)* given in Proposition 1.

We now show that it is strictly beneficial for the principal to delay the release of performance information to the last date. Specifically, consider the following alternative review strategies: the principal chooses a sequence of dates $1 \leq t_1 < ... < t_k < t$; at the end of date t_j the principal evaluates the agent's interim performances in dates $t_{j-1}+1, ..., t_j$ and compensates the agent based on the evaluations. This is a common practice in performance evaluation. The agent knows nothing about the output realizations within a review period before the final date of that period. Therefore, effectively the principal divides the problem into k separate problems; each corresponds to a particular review period. Each of these problems can then be analyzed in the same way as we studied Program (P). Let $C(s,\tau)$ be the minimum cost, discounted to period 1, that the principal needs to pay in order to implement effort sequence $(e_{s+1}, ..., e_\tau)$.

Proposition 2. *For all $t > 0$, the principal strictly prefers to delay the release of performance information to the last period t in the sense that*

$$C(0,t) < C(0,t_1) + \cdots + C(t_k,t)$$

for all $1 \leq t_1 < ... < t_k < t$ with $k \geq 1$.

If the contract given in Proposition 1 is the unique optimal contract then delaying the review to the final period is strictly better because the optimal contract can not be replicated by any short-term contracts. Proposition 2 directly verifies the superiority of pooling by showing exactly how the incentive cost can be reduced.

Proof. It suffices to prove the result for two separate review phases as the general case then follows by induction. Suppose that the principal divides the t dates into two subperiods: the first from date 1 to some date $\tau < t$ and the second from $\tau + 1$ to t.

By Proposition 1, if the principal reviews the agent only at the end of period t, then an optimal scheme is to reward the agent only for signal z^t with a bonus equal to

$$W^* = \frac{G(e^t) - G(b^t)}{P(z^t|e^t) - P(z^t|b^t)}$$

for some effort sequence $b^t = (b_1, ..., b_t) \neq e^t$.

For ease of exposition, decompose each action sequence a^t into (a^1, a^2), where $a^1 = (a_1, ..., a_\tau)$ and $a^2 = (a_{\tau+1}, ..., a_t)$, and define the following probabilities

$$P_1 = P(z_1, ..., z_\tau | e^1) \text{ and } P_2 = P(z_{\tau+1}, ..., z_t | e^2)$$

$$Q_1 = P(z_1, ..., z_\tau | b^1) \text{ and } Q_2 = P(z_{\tau+1}, ..., z_t | b^2).$$

By Assumption 1, $P_i > Q_i$ whenever $b^i \neq e^i$. Let $G(e^i), G(b^i)$ be the costs of the respective effort sequences discounted to period 1; for instance, $G(e^2) = \delta^\tau e_{\tau+1} + \cdots + \delta^{t-1} e_t$.

If $G(e^i) > G(b^i)$ for both $i = 1, 2$, then the cost for implementing e^t satisfies

$$\begin{aligned} C(0,t) = P_1 P_2 W^* &= \frac{P_1 P_2 (G(e^1) + G(e^2) - G(b^1) - G(b^2))}{P_1 P_2 - Q_1 Q_2} \\ &= \frac{P_1 P_2 (G(e^1) - G(b^1))}{P_1 P_2 - Q_1 Q_2} + \frac{P_1 P_2 (G(e^2) - G(b^2))}{P_1 P_2 - Q_1 Q_2} \\ &< \frac{P_1 (G(e^1) - G(b^1))}{P_1 - Q_1} + \frac{P_2 (G(e^2) - G(b^2))}{P_2 - Q_2} \\ &\leq C(0,\tau) + C(\tau,t). \end{aligned}$$

The first inequality follows from $P_i > Q_i$ and the second follows because

$$\frac{P_i (G(e^i) - G(b^i))}{P_i - Q_i}$$

is the expected cost (discounted to period 1) for implementing e^i against the *only* deviation b^i and it is no more than the cost required to implement e^i against *all* possible deviations.

If $G(b^1) \geq G(e^1)$ (e.g. $b^1 = e^1$), then

$$C(0,t) \leq \frac{P_2 (G(e^2) - G(b^2))}{P_2 - Q_2} \leq C(\tau,t) < C(0,\tau) + C(\tau,t)$$

since $C(0,\tau) > 0$ for all $\tau = 1, ..., t$. The result is the same if $G(b^2) \geq G(e^2)$. □

The general intuition for this result is simple: working under a single pooled incentive constraint (e^t against b^t) instead of two separate constraints (e^i against b^i for each i) gives the principal more leeway in choosing wage plans. This extra latitude is indeed beneficial : When implementing each e^i by rewarding z^τ or $z^{t-\tau}$, a dollar of (expected) reward money the principal spends brings an incentive benefit (in terms of adding slack to the incentive constraint) equal to

$$\frac{P_i - Q_i}{P_i} = 1 - \frac{Q_i}{P_i}$$

and it is lower than the benefit-cost ratio when rewarding the joint signal $(z^\tau, z^{t-\tau})$, which equals

$$\frac{P_1 P_2 - Q_1 Q_2}{P_1 P_2} = 1 - \frac{Q_1}{P_1} \cdot \frac{Q_2}{P_2}.$$

The essential idea is that the product of the likelihood ratios $\frac{P_1}{Q_1} \cdot \frac{P_2}{Q_2}$ is no less than the individual ratios (given $P_i \geq Q_i$). This result reflects the complementarity between the individual efforts e_τ in raising the probability of receiving the reward, $P(z^t|e^t)$: each additional choice of e_τ has a greater impact if the agent has already chosen such efforts in other periods. Abreu, Milgrom and Pearce (1991) attribute this result to the "reusability" of rewards (they actually consider punishments): rewards paid for the joint signal $(z^\tau, z^{t-\tau})$ provides incentives simultaneously for effort sequences e^1 and e^2 and hence reduces the expected cost compared with separately motivating each individual effort sequence. Note that for this reusability to work, there should be informational gain by pooling observations across periods. Section 4 describes when such gains may not be available.

This result offers an alternative and complementary explanation for the use of *deferred compensation*, a practice widely documented and studied (see Prendergast, 1999, Section 3.1 for a good review). The leading incentive-based explanation by Lazear (1981) uses efficiency wage theory with the threat of firing: older workers are willing to exert effort for fear of losing the rents associated with their jobs. These rents are also attractive to younger workers because exerting effort increases their chances of surviving in the firm long enough to receive them. As a result, the compensation profile should be tilted towards later stages in a worker's tenure, which provides incentives for both young and old workers (see also Akerlof and Katz, 1989). In the present model the worker can not be fired and the benefit of deferred compensation comes from the efficient pooling of performance information across periods achieved by delaying its release. In fact, we have the following:

Observation 1. If performance information can not be hidden from the agent then the principal can not gain from deferred compensation.

The reason is that when the principal wants to induce the efforts e_τ in every period the same amount of rent must be paid to the agent regardless of the timing of compensation.[6]

[6]The idea may be formalized as follows. Suppose the principal and agent both observe the realized output y_τ in every period and the principal makes wage payments in the last period t. For $s = 1, ..., t$, let $R(e_s, ..., e_t) > 0$ be the agent's payoff net of effort costs (i.e. rent) that the principal has to pay for implementing effort sequence $(e_s, ..., e_t)$. Let $U(y_1)$ be the agent's expected payoff conditional on first-period output being y_1. Since the principal wants to implement

Since delaying performance review is better for the principal, will the agent's rent decrease as t increases? Will the rent vanish as $t \to \infty$? We may get a glimpse of the answers to these questions by inspecting the optimal bonus, given in Eq. (2). The principal's cost for implementing e^t and the agent's rent are respectively given by

$$C(t) \equiv \frac{P(z^t|e^t)(G(e^t) - G(b^t))}{P(z^t|e^t) - P(z^t|b^t)} = \frac{G(e^t) - G(b^t)}{1 - \frac{P(z^t|b^t)}{P(z^t|e^t)}},$$

$$R(t) \equiv C(t) - G(e^t) \geq 0.$$

Based on these calculations, we have the following:

Observation 2.

a. If the likelihood ratio $\frac{P(z^t|b^t)}{P(z^t|e^t)} \to 0$ as $t \to \infty$ then $C(t) \to G(e^t) - G(b^t)$. Since the agent's rent $R(t) = C(t) - G(e^t)$ is nonnegative and goes to $-G(b^t)$ as $t \to \infty$, then the binding effort sequence b^t must be the minimum effort sequence, i.e. $G(b^t) = 0$. In this case, as $t \to \infty$, the agent's rent vanishes: $R(t) \to 0$.

b. If the binding effort sequence b^t differs from the minimum effort sequence as $t \to \infty$, then $\frac{P(z^t|b^t)}{P(z^t|e^t)}$ must *not* converge to zero and the agent will receive positive rent even as $t \to \infty$.

When will either of these scenarios arise depends on the parameters of the model, in particular on the discount factor. To gain a sharper insight on this question, next we specialize the model to a stationary binary choice setting.

$(e_2, ..., e_t)$ regardless of y_1, we must have $U(y_1) \geq R(e_2, ..., e_t)$. Then, to motivate the agent to choose e_1 in period 1, we must have

$$\sum_{y_1} (U(y_1) - R(e_2, ..., e_t))p(y_1|e_1) - g(e_1) \geq R(e_1)$$

where $R(e_\tau)$ is the rent paid to the agent for inducing e_τ in the one-period problem. Since

$$R(e^t) = \sum_{y_1} U(y_1)p(y_1|e_1) - g(e_1)$$

we have $R(e^t) \geq R(e_1) + R(e_2, ..., e_t)$. By induction, $R(e^t) \geq R(e_1) + ... + R(e_t)$. Thus the principal must pay the agent as much rent as she would when dealing with each period separately.

2.3 Rent extraction as t becomes large

Assume in this section that the agent has the same feasible choices in each period: $\{0,1\}$. Low effort costs the agent nothing but high effort costs $c>0$. The finite set of output signals, Y, is also stationary over time. There is a signal $z\in Y$ such that $p(z|1)/p(z|0)\geq p(y|1)/p(y|0)$ for all $y\in Y$, with strict inequality for some y. For convenience, let $p=p(z|1)$ and $q=p(z|0)$; then $0<q<p<1$.

The principal's objective is to implement the effort sequence $e^t\equiv(1,....,1)$ with minimum cost. By Proposition 1, an optimal scheme for the principal is to reward the agent a bonus at the end of period t if and only if the realized output path is $Z\equiv(z,...,z)$. The optimal bonus W^*, given in Eq. (2), depends on the binding effort sequence b^t.

To find a binding sequence b^t, consider a collection of relaxed problems as follows. For $k=1,....,t$, the agent is only allowed to choose between effort sequence e^t and the sequence that differs from e^t *only* in the initial $k\leq t$ periods. By the previous argument, the optimal bonus $W(k)$, valued at date 1, is determined by the following incentive constraint:

$$(p^t-p^{t-k}q^k)W(k)=\frac{1-\delta^t}{1-\delta}c-\delta^k\frac{1-\delta^{t-k}}{1-\delta}c$$

and hence

$$W(k)=\frac{(1-\delta^k)c}{(1-\delta)p^t(1-\theta^k)},\ \text{ where }\ \theta\equiv\frac{q}{p}. \tag{3}$$

Note that if the agent deviates k times but not all in the initial k periods then he is worse off than if he deviated in the initial k periods; this is because the chances of getting the bonus is the same for all k-deviations but deviating in the initial periods saves effort costs due to discounting. Therefore, $W(k)$ is the minimum bonus needed to motivate the agent not to take any k deviations. It follows that the optimal bonus is given by

$$W^*\equiv\max_{k\in\{1,...,t\}}W(k). \tag{4}$$

Once a binding effort sequence involving k^* initial deviations is found, the principal's expected wage payment and the agent's rent can be calculated respectively as

$$C(t)=W(k^*)\cdot p^t=\frac{(1-\delta^{k^*})c}{(1-\delta)(1-\theta^{k^*})}$$

$$R(t)=C(t)-\frac{1-\delta^t}{1-\delta}c.$$

The following proposition characterizes the value of k^* and the conditions under which the agent's rent disappears as $t \to \infty$.

Proposition 3. *Let $\theta = q/p \in (0,1)$.*

(a) If $\delta > \theta$ then for t sufficiently large, the binding sequence $b^t = (0,..,0)$; moreover, the principal's cost goes to $\frac{c}{1-\delta}$ and the agent's rent goes to zero as $t \to \infty$.

(b) If $\delta = \theta$ then the agent is indifferent between e^t and every effort sequence that involves any $k = 1,\ldots,t$ initial deviations and the principal's cost equals $\frac{c}{1-\delta}$ for all t; moreover, as $t \to \infty$, the agent's rent goes to zero.

(c) If $\delta < \theta$ then the binding sequence b^t involves exactly one deviation in the first period; moreover, for all $t > 1$, the principal's cost equals $\frac{c}{1-\theta}$ and hence the agent receives positive rent $\frac{c}{1-\theta} - \frac{c}{1-\delta}$ even as $t \to \infty$.

As discussed earlier, efforts in each period are complementary to each other for raising the probability of receiving the reward. Therefore, working in each additional period brings larger extra return to the agent. On the other hand, efforts in different periods have different utility costs due to discounting, which makes exerting additional effort more costly: from the ex ante point of view, effort in period t has the lowest present-value cost so it is the easiest to induce; the marginal cost of effort monotonically goes up as we go backwardly from period t toward period 1. These are two opposite effects on the agent's overall marginal return from effort. When discounting is not too heavy, the first effect dominates as the marginal cost of effort goes up relatively slowly. As a result, the agent's marginal return of effort increases with total effort level and his utility is convex in effort level. Hence, to induce effort in all periods it is essential for the principal to prevent the agent from shirking in *all* periods. By contrast, when discounting is sufficiently heavy (δ sufficiently small), the second effect dominates and the agent's marginal utility decreases with total effort level. Then the last unit of the effort is the hardest to induce and the principal need go no further than deterring a single deviation by the agent in period one. As we shall see in the next section, the aforementioned two effects always work in the same direction for the private evaluation case, independent of the discount rate. But regardless of which effect dominates, there is always a synergy for incentive provision by pooling information across periods.

The current model is related to multitask principal-agent problems, for which the effort substitution problem has been a prominent source of concern for optimal contracting (see Holmström and Milgrom, 1991 and Laffont and Martimort, 2002). MacDonald and Marx (2001) study a multi-task principal-agent problem in which the tasks are complementary for the principal but the efforts are substitutes for the agent. In the current model, although efforts are not physically transferable between different periods they can be thought of as substitutes for the agent because their present-value utility costs satisfy a convex structure due to discounting.[7] In MacDonald and Marx (2001), in order to overcome the effort substitution problem (or adverse specialization problem as they call it) optimal contracts induce the agent to view efforts in different tasks as complements. This is also true in the present model when the discount factor $\delta > \theta$; but it is not true for $\delta < \theta$: given the optimal contract, efforts are still substitutes for the agent. In fact, we shall see in the next section that efforts will *always* remain substitutes for the agent in the private evaluation case but this *does not* prevent the principal from reaping the benefit of pooling.

The result that the principal is always better off by delaying the release of performance information implies that there is no optimal contract if the relationship lasts indefinitely: $t = \infty$. The principal however can get arbitrarily close to her limiting payoff by rewarding the agent with a large prize but with vanishing probability. It has the flavor of a "lifetime achievement award" and is opposite to the model of Mirrlees (1974) in which the principal can get arbitrarily close to her best outcome by punishing the agent severely but with a very small probability.

2.4 Discussions

It should be stressed that information pooling has no benefit in the current setting if the agent does not face limited liability constraints. The optimal contract, as is well-known (Innes, 1990), is to "sell" the firm to the agent at a fixed price. The agent then becomes the residual claimant and hence has the right incentive to work. It can be done on a period-to-period basis or with some information lag without affecting the outcome, because agent's rent can be eliminated in either case.

With limited liability, there are still other factors that may limit the usefulness of pooling. If the agent is very averse to risk then pooling information across periods may not be useful for the principal as the skewed and risky bonus scheme can be too costly. Also, as t becomes large, the

[7] The caveat is that the agent should increase his total effort backwardly from period t toward period one; a reverse order would make efforts technologically complements for the agent. The former is the relevant order in the current setting because ex ante the agent is free to choose when to exert efforts and he rationally chooses the less-costly ones first.

required bonus may become very large, which can be beyond the principal's ability to pay.[8] A third factor is the bargaining power of the agent. Since longer review phase decreases the agent's rent, the agent may demand more frequent reviews if he has some bargaining power, which will put a limit on how long the principal can delay the release of information. The discussion so far has been under the condition that there indeed is some informational gain from pooling observations across periods (Assumption 1). Failure of such a condition would of course restrain the usefulness of pooling.[9]

In the next section we investigate yet another factor: the principal's evaluation of the agent's performance may be subjective, non-verifiable, and even private. We will see that it has completely different implications for limited liability and risk aversion.

3 Private Evaluation and Deferred Punishment

In the preceding section, it was assumed that performance can be measured objectively and can be verified by a court. As Prendergast (1999) points out, however, in many cases some aspects of a worker's performance are hard to measure objectively; as a result, firms frequently resort to subjectively determined measurements to evaluate a worker's job performance. In this and the following sections, we analyze the implications of private evaluation on the nature of optimal dynamic contracts.

The basic setting is very similar to that in Section 2, which we recapitulate as follows. In each period $\tau = 1, \ldots t$, the agent chooses some hidden action a_τ from some compact set $A_\tau \subset \Re_+$; the action generates some stochastic benefit y_τ for the principal, which is drawn from some finite set $Y_\tau \subset \Re_+$ according to a probability distribution $p_\tau(y_\tau|a_\tau) > 0, \forall y_\tau$. Different from in the previous section, the realized output y_τ is *privately* observed by the principal; this is to capture the idea of subjective evaluation.

The agent maximizes the sum of expected discounted utilities and has a per-period utility function $u(c_\tau) - g(a_\tau)$, where either $u(c) = c$ for $c \in \Re$ or $u : (0, \infty) \to \Re$ is concave, continuous, strictly

[8] Of course, in any realistic setting the agent's job tenure is limited by factors such as turnover or retirement, so it would be quite rare that the length of a contract is so long that it requires bonuses that are beyond the principal's means. Thus this may not be such a key factor in restricting the principal's ability.

[9] We have not considered other more intuitive reasons. Some are in favor of a more frequent review: for instance, performance review can provide feedbacks to the agent and hence helps on-the-job learning. Others are in favor of a longer review phase: for instance, the agent's present effort may have impact on future outputs.

increasing and $\lim_{c\to 0} u(c) = -\infty$.

Therefore, different from in the last section, here the agent can be risk averse and is not protected by limited liability. The latter captures the idea that the principal can punish the agent for poor performance. For instance, the agent may be required to pay a fine or be punished via non-pecuniary means (e.g. loss of reputation) for unsatisfactory performance (see MacLeod (2007) for a survey of the use of reputation in contract enforcement); or the job may have a "lock-in" effect: if the agent is terminated after a period of time he may have difficulty finding a comparable job and hence will suffer losses. We will reinstate limited liability in Section 4.

As before, the principal would like to induce the agent to choose a desired sequence of actions $e^t = (e_1, ..., e_t)$ with minimum cost on wage compensation. The general mechanism the principal uses can be described as follows. In each period τ, the principal reports the realized output to a planner; based on the reported output history the planner asks the principal to pay some amount $w(y^\tau)$ and then pays the agent some $c(y^\tau) \leq w(y^\tau)$. There is no storage or borrowing by the agent, so $c(y^\tau)$ is his consumption. The important feature of this mechanism is that *the planner need not reveal the principal's reports to the agent.* Therefore, upon receiving $c(y^\tau)$ the agent can only infer but are not sure what the reports are, unless the compensation plan is such that $c(y^\tau) \neq c(\hat{y}^\tau)$ for all $y^\tau \neq \hat{y}^\tau$.

A mechanism needs to satisfy incentive compatibility and feasibility constraints as follows: (i) the principal should have incentive to truthfully report the output realizations; (ii) for all τ, given a consumption path $c^\tau \equiv (c_1, ..., c_\tau)$ and based on the specified consumption plan $c(\cdot)$, the agent forms a belief $\beta_\tau(y^\tau \mid c^\tau)$ over the set of τ-period output histories y^τ; given this belief and the consumption plan, the agent should find it optimal to choose the continuation sequence $(e_{\tau+1}, ..., e_t)$; (iii) budget constraint: $c(y^\tau) \leq w(y^\tau)$, $\forall y^\tau$ and $\forall \tau$; (iv) participation constraint of the agent (to be given later).

Compared with the case of public evaluation, an additional truth-telling incentive constraint for the principal is added. Since the planner has no additional instrument to elicit the principal's information, for the principal to tell the truth all the time her payoff has to be independent of her report at any history. This observation is summarized as follows.

Lemma 2. *A wage plan $w(\cdot)$ is incentive compatible for the principal if and only if there exists some constant $\widehat{w}$ such that for all output path y^t,*

$$w_1(y^1) + \delta w_2(y^2) ... + \delta^{t-1} w_t(y^t) = \widehat{w}. \tag{5}$$

Therefore, the principal can just set aside some amount of money $\widehat{w}$ at the beginning and distributes it to the agent over time so as to motivate the agent to choose e^t. The sequential incentive constraints for the agent may seem quite daunting, especially when compared with the previous case, because with a risk-averse agent consumption allocation has both informational and risk-sharing effects. Fortunately, much of the perceived complication turns out to be superfluous.

Consider the following relaxed problem for the principal.

Problem Q: $\max_{w, c(\cdot)} -w$

subject to:

$$\sum_{\tau=1}^{t} \sum_{y^\tau} \delta^{\tau-1} u(c_\tau(y^\tau)) P(y^\tau | e^\tau) - G(e^\tau) \geq \overline{U} \tag{IR}$$

$$e^t \in \arg\max_{a^t} \sum_{\tau=1}^{t} \sum_{y^\tau} \delta^{\tau-1} u(c_\tau(y^\tau)) P(y^\tau | a^\tau) - G(a^t) \tag{ICA}$$

$$w \geq c_1(y^1) + \ldots + \delta^{t-1} c_t(y^t), \ \ \forall y^t. \tag{BC}$$

where again $P(y^\tau | a^\tau) \equiv \prod_{s=1}^{\tau} p_\tau(y_s | a_s)$ and $G(a^\tau) \equiv g_1(a_1) + \ldots + \delta^{\tau-1} g_\tau(a_\tau)$. (IR) is the agent's participation constraint and $\overline{U}$ is the agent's reservation utility at the time of signing the contract in period one. Once the contract is signed the agent works for the principal throughout the t periods. In Program *(Q)*, (ICA) relaxes the sequential incentive constraint of the agent to a single ex ante constraint in period $\tau = 1$; (BC) relaxes the original period-by-period budget constraints to a single budget constraint in terms of the sum of all future payments. The next proposition significantly simplifies Problem *(Q)*.

Proposition 4. *The principal can restrict attention to contracts of the following form that do not reveal the agent's performance information until the end of period t: the agent is paid a constant amount $\bar{c}$ in periods $1, \ldots, t-1$ and the same amount in period t except after some specific output histories y^t at which some lower payments are made.*

Proof. Suppose $(w, c(\cdot))$ is a solution to Program *(Q)*. The key is to observe that for all action sequence a^t,

$$\delta^{\tau-1} u(c(y^\tau)) = \sum_{y^{t-\tau}} \delta^{\tau-1} u(c(y^\tau)) p(y_{\tau+1}) \cdots p(y_t). \tag{6}$$

Hence the agent's expected utility from consumption is given by

$$\sum_{\tau=1}^{t}\sum_{y^\tau}\delta^{\tau-1}u(c_\tau(y^\tau))P(y^\tau)=\sum_{\tau=1}^{t}\sum_{y^t}\delta^{\tau-1}u(c_\tau(y^\tau))P(y^t)=\sum_{y^t}\sum_{\tau=1}^{t}\{\delta^{\tau-1}u(c_\tau(y^\tau))P(y^t)\},$$

where the last equality is by changing the order of summation. Define

$$U(y^t)\equiv\sum_{\tau=1}^{t}\delta^{\tau-1}u(c_\tau(y^\tau))$$

as the sum of discounted utilities along every output path y^t. Then the participation constraint (IR) and incentive constraint (ICA) respectively become

$$\sum_{y^t}U(y^t)P(y^t|e^t)-G(e^t)\geq\bar{U}, \tag{7}$$

$$\sum_{y^t}U(y^t)\left[P(y^t|e^t)-P(y^t|a^t)\right]\geq G(e^t)-G(a^t),\ \forall a^t. \tag{8}$$

Thus, consumption plan $c(\cdot)$ matters only to the extent that it affects the plan $U(y^t)$.

We now construct another solution to *(Q)* that fully delays information release to period t. Let $\bar{c}=(1-\delta)w/(1-\delta^t)$ be the per-period consumption when the total amount w is distributed to the agent evenly across periods, i.e. $\bar{c}(1+\delta\cdots+\delta^{t-1})=w$. The new plan is to first set

$$\tilde{c}(y^\tau)=\bar{c},\ \text{ for all } y^\tau \text{ and all } \tau.$$

This perfect consumption smoothing across time maximizes the agent's utility along every output path y^t , i.e. $\widetilde{U}(y^t)=(1+\cdots+\delta^{t-1})u(\bar{c})\geq U(y^t)$. If for some y^t, $\widetilde{U}(y^t)>U(y^t)$, then reduce the period-t consumption $\tilde{c}(y^t)$ by some amount so that $\widetilde{U}(y^t)=U(y^t)$. In the end, $\widetilde{U}(y^t)=U(y^t)$ for all y^t. The new plan is feasible because both constraints (7) and (8) and the budget constraint remain satisfied and it does not increase the principal's spending, so it must solve Program *(Q)*.

Moreover, under the new plan, the agent's consumption is constant till the last period so the agent's beliefs about output histories are solely determined by his actions: nothing can be learned from consumption. Therefore the agent's sequential incentive constraints are also satisfied; hence the new plan also solves the original problem. □

The principal, due to her own incentive problem, is essentially required to put aside a fixed amount of money for compensating the agent. The mechanism freely allocates the money to the agent across periods for both insurance and incentive purposes. What matters for the agent's ex ante incentive is his total payoff $U(y^t)$ along every output path y^t; the maximum of such $U(y^t)$ is achieved by perfect consumption smoothing across periods; any lower payoff may be referred to as a punishment. Absent limited liability, a punishment can be handed out to the agent in any particular period. But to provide incentive in *all* periods, the punishment has to be deferred to the last one. Thus, *insurance and incentive provision are in perfect harmony in all but the last period.*[10]

This result shows that optimal wages are *sticky* at some fixed level even if there are arbitrarily asymmetric productivity shocks across periods. This is a strong form of low pay-performance sensitivity and it is consistent with the empirical findings documented in Jensen and Murphy (1990). It also suggests that empirical studies may find more support for incentive theory by looking at pay-performance relations at the actual times of performance review, which may span long periods of time.

To obtain a sharper characterization of the optimal contracts, make the following assumption about the information structure.

Assumption 2. For every $\tau = 1,\ldots,t$, there is an output signal $z_\tau \in Y_\tau$ such that for all $a_\tau \neq e_\tau$:

$$\frac{p(z_\tau|e_\tau)}{p(z_\tau|a_\tau)} \leq \frac{p(y_\tau|e_\tau)}{p(y_\tau|a_\tau)} \text{ for all } y_\tau, \text{ with strict inequality for some } y_\tau.$$

In contrast to Assumption 1, which assumes there is a *good* signal for detecting *conformity* to action e_τ, Assumption 2 requires that there exist a *bad* signal for detecting *deviation* from e_τ.[11]

Proposition 5. *Suppose Assumption 2 holds. Then there is an optimal contract of the following form: the agent is paid a constant wage $\bar{c}$ regardless of history for periods $\tau < t$; in period t the agent is paid the same wage $\bar{c}$ unless the performance path $z^t = (z_1,\ldots,z_t)$ has realized, which results in a wage payment $c(z^t) < \bar{c}$. Moreover, letting $\bar{u} = u(\bar{c})$ and $u(z^t) = u(c(z^t))$, we have*

$$\bar{u} = \frac{1-\delta}{1-\delta^t}\left(G(e^t) + \overline{U} + P(z^t|e^t)K\right) \tag{9a}$$

[10]This analysis also helps to clarify why delayed performance review may not be useful when evaluation is verifiable and the agent is risk averse. In that case, the principal's payment need not be independent of output realizations and the agent's consumption profile will in general affect the principal's *expected* wage payment. Insurance and incentive provision may not go along as well as they do here: If reporting lag is too long then since all the variations in consumption are put off to period t the cost for insuring the agent can be very high for the principal.

[11]Both assumptions are satisfied if the agent only has two alternative actions in each period or if the monotone likelihood ratio condition is satisfied.

$$u(z^t) = \bar{u} - \delta^{1-t}K \tag{9b}$$

where

$$K \equiv \frac{G(e^T) - G(b^t)}{P(z^t|b^t) - P(z^t|e^t)} \text{ for some effort sequence } b^t. \tag{10}$$

While Proposition 4 describes wage stickiness, a form of wage compression over time, Proposition 5 describes wage compression across performance levels. This latter result has been previously obtained by MacLeod (2003) in a static model (see the discussion below in relation to Fuchs, 2007). The underlying logic is essentially the same in the multiperiod setting once the connection to the static model is recognized through Proposition 4. Essentially, since signal path z^t is most informative about deviations from effort sequence e^t it is most cost effective to punish the agent only when z^t is realized.

We now show that it is strictly beneficial for the principal to delay the release of performance information to the last date. Specifically, consider the following alternative review strategies: the principal chooses a sequence of dates $1 \le t_1 < ... < t_k < t$ and at the end of date t_j the principal evaluates the agent's interim performances at dates $t_{j-1}+1, ..., t_j$ and compensates the agent based on the evaluations. The agent knows nothing about the output realizations within a review period before the final date of that period. Denote such a reporting strategy by $S(t_1, ..., t_k)$.

Proposition 6. *For all $t > 0$, the principal strictly prefers delaying the release of all performance information to the last period t to any reporting strategy $S(t_1, ..., t_k)$.*

If the contract given in Proposition 5 is the unique optimal contract then delaying the review to the final period is strictly better because the optimal contract can not be replicated by any short-term contracts. Proposition 6 directly verifies the superiority of lagged review by showing exactly how the incentive cost can be minimized. The idea can be sketched as follows.

Although the agent does not earn any rent, the principal nevertheless pays expected social cost $P(e^t)K$ over and above the agent's reservation utility $\overline{U}$ because of the informational constraint on her own part. This social cost has a close resemblance to the agent's rent in the case of verifiable evaluation. The important observation is that *the expected social cost is positively related to the principal's overall cost* so we may as well regard the social cost as the objective that the principal seeks to minimize. Suppose the principal divides the t periods into two review phases: from period

1 to τ and then from $\tau+1$ to t. In each review phase i, when implementing e^i by punishing the agent for signals z^τ or $z^{t-\tau}$, each unit of (expected) social cost the principal spends brings an incentive benefit (in terms of adding slack to the incentive constraint) equal to

$$\frac{Q_i - P_i}{P_i} = \frac{Q_i}{P_i} - 1$$

and it is lower than the benefit-cost ratio when rewarding the joint signal $(z^\tau, z^{t-\tau})$, which equals

$$\frac{Q_1 Q_2 - P_1 P_2}{P_1 P_2} = \frac{Q_1}{P_1} \cdot \frac{Q_2}{P_2} - 1.$$

Similar to the public evaluation case, the idea is that the product of the inverse likelihood ratios $\frac{Q_1}{P_1} \cdot \frac{Q_2}{P_2}$ is no less than the individual ratios (given $Q_i \geq P_i$). But contrary to the previous case, this result reflects a complementarity between individual *deviations* from the desired actions e_τ in raising the probability of *punishment* $P(z^t|a^t)$: each additional deviation carries a greater penalty or alternatively each additional effort brings less utility for the agent. Moreover, if the environment is stationary, then ex ante, efforts in different periods carry different utility costs due to discounting, and the marginal cost of effort monotonically increases as the agent complies in more periods, going from period t toward period 1. These two effects reinforce each other and hence the agent's overall marginal return from effort diminishes as he exerts effort in more periods. This explains that in a stationary setting a single deviation in period 1 is the hardest to prevent, as found in Abreu, Milgrom and Pearce (1991) and Fuchs (2007). In contrast, recall that in the public evaluation case if discounting is not too heavy it is the hardest to prevent the agent from deviating in *all* periods (Proposition 3a) due to the complementarity between individual efforts e_τ in raising the probability of *reward*. These differences notwithstanding, pooling is beneficial in both cases.

The analysis in this section is closely related to Fuchs (2007). Fuchs analyzes a stationary model in which the agent's effort choice and the output signal are both binary, and obtains results similar to Propositions 5 and 6 here.[12] The present model generalizes his model in two directions. First, we assume that the agent is risk averse and can consume in every period whereas Fuchs assumes that the agent is either risk neutral or risk averse but consumes at the end of the contractual relationship. By assuming risk aversion and flexible timing for consumption, we are able to provide a justification for sticky wages over time (Proposition 4); in contrast, when the agent is risk neutral the wage profile across time is largely indeterminate. Second, we place no symmetry assumption across periods and

[12] Fuchs also characterizes solutions to the infinite-horizon version of the model.

place weaker assumptions on the signal structure. The present formulation makes the results more general and also reveals the fundamental conditions (in terms of likelihood ratios) for pooling to work. This generality also brings out the connections to and differences from the case of public evaluation.

4 Limited Liability and Private Evaluation

One key element of the model in Section 3 is the principal's ability to punish the agent. In this section, we show that losing this ability in conjunction with private evaluation can put some restraint on the benefit the principal gets from pooling performance information across periods. To make the point as clearly as possible, we consider a stationary setting as follows.

The principal hires the agent for t periods. In each period, the agent can privately choose one of two effort levels: $a_\tau = 0$ or 1. Low effort costs the agent nothing but high effort costs $c > 0$. There are two output levels: h (high) and ℓ (low). Realized outputs are privately observed by the principal. Let p (resp. q) be the probability that output level h will occur given effort level $a_\tau = 1$ (resp. $a_\tau = 0$). Assume $0 < q < p < 1$. The agent is risk neutral and must be paid nonnegative wages in every period; his reservation utility equals zero. For simplicity, there is no discounting. The principal's objective is to implement effort level 1 in all periods with minimum cost on wage compensation.

With both private evaluation and limited liability, the principal not only pays rent to the agent but also pays a social cost, thus combing elements in the previous two sections. The analysis is more complex and importantly, as we show below, pooling may not always benefit the principal.

The difference in the use of information can be explained as follows. In the previous two sections, the principal effectively minimizes either the rent paid to the agent or the social cost. And the principal's cost is the *expected* value of rewards or punishments. As a result, rewards or punishments can be efficiently used on output signals that are *most* informative in terms of the value of a likelihood ratio. Importantly, the extreme vales of the likelihood ratio naturally extend to multiple periods as a result of multiplication, i.e. if two signals z_1 and z_2 maximize likelihood ratios in two separate periods then the joint signal (z_1, z_2) maximize the likelihood ratio across the two periods. This structure contributes to the simple form of optimal contracts in the previous two cases.

In contrast, in the current problem the principal needs to minimize the *sum* of the rent and social cost, which equals the *maximum size* of the rewards, not their expected value. For this purpose, the principal should reward any signal that indicates conformity and punish (by withholding rewards) any signal that indicates deviation, not just the ones that are most informative. Unfortunately, a pooled signal may indicate either a deviation or conformity depending on which effort sequence is referenced to; so a nice order of informativeness among pooled signals may not exist for checking against all possible deviations even if such an order exists for signals in individual periods. As a result, pooling may not provide the informational gain needed for efficient incentive provision and hence may not benefit the principal. This point is illustrated by the following result.

Proposition 7. *Suppose $t = 2$. If $p > \frac{1}{2}$ and $p + q < 1$ then contracting for each period separately is optimal so pooling has no advantage.*

The result shows that pooling is useless when a single period already provides good information on the agent's effort: high output is more likely with high effort in an absolute sense ($p > \frac{1}{2}$); high output is sufficiently less likely with low effort ($q < 1 - p < \frac{1}{2} < p$). The basic reason is that rewarding a single success is good for deterring two deviations but bad for deterring just one deviation. As a compromise, in the optimum a single success should only be rewarded with half of the total prize. The result is exactly the optimal contract that treats each period separately.

Despite this negative result, the following proposition identifies two scenarios in which pooling is useful and optimal contracts resemble what have been found previously.

Proposition 8. *If (a) $p^t - q^t > p^{t-1} - q^{t-1}$ or (b) $p < \frac{1}{t}$ then the principal's average cost per date decreases as contract length increases, up to t. Moreover, if (a) holds then an optimal contract for a t-period relationship is to reward the agent a bonus $B = tc/(p^t - q^t)$ only if the realized output path is $(h, \dots, h)$; if (b) holds then an optimal t-period contract is to reward the agent a bonus $K = \frac{c}{(1-p)^{t-1}(p-q)}$ for all output path except $(\ell, \dots, \ell)$. In both cases, the principal's cost equals the full bonus.*

Here, hypothesis (a) means that the probability of full success $(h, \dots, h)$ increases in an absolute sense when the agent works in more periods, compared with always shirking. In such a case, the punishing region includes every pooled signal except $(h, \dots, h)$; that is, anything short of full success indicates a lack of effort by the agent and should be punished by forfeiting the reward. The optimal contract resembles that in Section 2. On the other hand, hypothesis (b) says that the probability of success is rather low even with high effort. In this case, the punishing region only contains the signal $(\ell,\ldots,\ell)$; that is, anything better than a total failure is indicative of good effort and should be rewarded with the full bonus. The optimal contract then resembles that in Section 3. In both cases, pooling up to length t benefits the principal because the stated conditions are necessarily satisfied and hence the above logic works for all length $k \leq t$. For instance, if hypothesis (a) holds, then for a length-k relationship the principal's per-period cost equals $c/(p^k - q^k)$, which decreases as contract length k increases, up to t. If hypothesis (b) holds, then for a length-k relationship the principal's per-period cost equals

$$\frac{c}{k(1-p)^{k-1}(p-q)},$$

which, given $p < 1/t$, also decreases in k as long as $k \leq t$.

When these conditions fail at some t, pooling across more periods may not be beneficial so there can be an optimal contract length, at least locally. In fact, the conditions in Proposition 7 exactly nullify both hypotheses in Proposition 8 for $t = 2$, which explains why pooling was not useful there. The general analysis for $t > 2$ is complex as the rewarding region may include any number of successes and the binding deviating sequence(s) may involve any number of deviating periods.

5 Conclusion

Pooling performance information across periods may help the principal to design more efficient incentive contracts. This paper has studied a set of factors that determine the effectiveness of pooling. These factors, including the verifiability of performance evaluation and limited liability and risk attitudes of the agent, have different requirements on the informational structure for pooling to be beneficial for the principal. In each case, the benefit of pooling is driven by a specific wedge between the agent's reservation utility and the payoff of the principal; different wedges result in different wage profiles across time. These findings, by offering some direct links between optimal wage policies and the characteristics of jobs and employees, can shed light on the various compensation practices used across jobs and firms, such as deferred compensation, bonus pay, deferred punishment, and sticky wages.

Appendix

Proof of Proposition 1.

Let $w(y^t) \geq 0$ be a solution to Program $P(a^t)$. Suppose $w(y^t) > 0$ for some $y^t \neq z^t$. Then we can modify the wage plan as follows: first reduce $w(y^t)$ to zero; then to keep the agent's incentive constraint intact, we increase $w(z^t)$ by an amount v so that $vP(z^t|e^t) = w(y^t)P(y^t|e^t)$. The agent thus gets the same payoff and hence the principal's payoff also stays the same when the agent chooses effort sequence e^t. If the agent chooses the alternative a^t then his payoff would be changed by:

$$\begin{aligned}\Delta &= vP(z^t|a^t) - P(y^t|a^t)w(y^t) \\ &= w(y^t)P(y^t|e^t)\left(\frac{P(z^t|a^t)}{P(z^t|e^t)} - \frac{P(y^t|a^t)}{P(y^t|e^t)}\right) \\ &\leq 0\end{aligned}$$

where the last inequality follows from Assumption 1. In other words, the agent can not gain by deviating from e^t under the new wage plan and hence the incentive constraint remain satisfied. Therefore, we can focus on wage plans that pay positive wage only for output path z^t. The optimal bonus is then determined by the following incentive constraint:

$$w(a^t)P(z^t|e^t) - G(e^t) = w(a^t)P(z^t|a^t) - G(a^t).$$

which leads to the solution given in the proposition after comparing among all a^t. □

Proof of Proposition 3.

By Eqs. (3) and (4), to find optimal bonus W^* we only need to look for k that maximizes

$$\beta(k) \equiv \frac{1-\delta^k}{1-\theta^k}.$$

And the principal's cost can be written as $C(t) = \beta(k)\frac{c}{1-\delta}$.

If $\delta > \theta$ then $\beta(k) < 1$ for all $k \geq 1$ and $\beta(k) \uparrow 1$ as $k \to \infty$. Therefore as $t \to \infty$, the binding sequence must involve $k(t)$ deviations with $k(t) \to \infty$.[13] It follows that as $t \to \infty$ the principal's cost $C(t) \to c/(1-\delta)$ and the agent's rent $R(t) \to 0$.

[13] In fact, it can be shown that $\beta(k)$ is strictly increasing in k when k is large. Therefore, the binding effort sequence is $L = (0,\ldots,0)$ for large t. Specifically, a sufficient condition for $\beta(k+1) > \beta(k)$ is $\delta^k - \delta^{k+1} > \theta^k - \theta^{k+1}$, which given $\delta > \theta$ is true as long as k is large enough.

If $\delta = q/p$ then $W(k)$ is constant across all k and hence the agent is indifferent between e and any effort sequence with $k \le t$ initial deviations. The principal's cost always equals $p^t W(k) = c/(1-\delta)$ for all $t \ge 1$. Hence the agent's rent $R(t)$ also goes to zero as $t \to \infty$.

If $\delta < \theta$ then for $k > 1$,

$$\beta(k) = \frac{1-\delta^k}{1-\theta^k} = \frac{(1-\delta)(1+...+\delta^{k-1})}{(1-\theta)(1+...+\theta^{k-1})} < \frac{1-\delta}{1-\theta} = \beta(1).$$

Therefore, the agent's binding incentive constraint involves the effort sequence that has exactly one deviation from H in period 1. It follows that the principal's cost $C(t) = \frac{c}{1-\theta}$ for all t and the agent receives positive rent that as $t \to \infty$ converges to

$$\frac{c}{1-\theta} - \frac{c}{1-\delta}. \qquad \square$$

Proof of Lemma 2.

This is most easily seen in the last period. In period t and given any output path y^{t-1}, in order for the principal to honestly report the current output y_t, her payment has to be independent of her report, i.e. there exists $\widehat{w}(y^{t-1})$ such that $\widehat{w}(y^{t-1}) = w_t(y^{t-1}, y_t)$, $\forall y^{t-1}, \forall y_t$. Similarly in period $t-1$ and for each y^{t-2} there exists some $\widehat{w}(y^{t-2})$ so that

$$\widehat{w}(y^{t-2}) = w_{t-1}(y^{t-2}, y_{t-1}) + \delta \widehat{w}_{t-1}(y^{t-2}, y_{t-1}) \text{ for all } y_{t-1}.$$

Applying similar arguments recursively proves the necessity of condition (5). For sufficiency, suppose condition (5) holds. Then at every history y^τ, the principal's sum of current and discounted future payments is the same regardless of his current report y_τ. $\square$

Proof of Proposition 5.

The proof is similar to that of Proposition 1. By Proposition 4, we can restrict our attention to consumption plans in which $c(y^\tau) = \bar{c}$ for $\tau < t$ and $c(y^t) \le \bar{c}$, $\forall y^t$, where the annualized consumption $\bar{c} = \frac{1-\delta}{1-\delta^t} w$. The equivalent utility assignments are $\bar{u} \equiv u(\bar{c})$ and $u(y^t) \equiv u(c(y^t))$.

We again consider a collection of problems that relax Program *(Q)*. In each relaxed program, the agent has just two choices: the effort sequence to be implemented, e^t, and an alternative sequence a^t. Thus the principal chooses $\bar{u}$ and $u(y^t) \le \bar{u}, \forall y^t$, to minimize $\bar{u}$, subject to

$$\sum_{y^t} \delta^{t-1} u(y^t) P(y^t|e^t) + \frac{1-\delta^{t-1}}{1-\delta} \bar{u} - G(e^t) \ge \bar{U} \tag{11}$$

$$\sum_{y^t} \delta^{t-1} u(y^t)(P(y^t|e^t) - P(y^t|a^t)) \geq G(e^t) - G(a^t) \tag{12}$$

We can, without loss of generality, focus on effort sequences a^t with $G(a^t) < G(e^t)$.[14] Suppose there is a solution in which $u(y^t) < \bar{u}$ for some $y^t \neq z^t$. We show that there is another solution in which the agent is punished only when z^t occurs. First set $\tilde{u}(y^t) = \bar{u}$. Then reduce $u(z^t)$ by some ε so that

$$\varepsilon \cdot P(z^t|e^t) = (\bar{u} - u(y^t))P(y^t|e^t).$$

Thus the agent's payoff stays unchanged when he chooses e^t. But if he chooses the alternative sequence a^t, his payoff changes by

$$(\bar{u} - u(y^t))P(y^t|a^t) - \varepsilon P(z^t|a^t) = (\bar{u} - u(y^t))P(y^t|e^t)\left(\frac{P(y^t|a^t)}{P(y^t|e^t)} - \frac{P(z^t|a^t)}{P(z^t|e^t)}\right) \leq 0.$$

Therefore the new plan satisfies all the constraints of the problem and does not reduce the principal's payoff. We thus found a solution in which $u(y^t) = \bar{u}(a^t)$ for $y^t \neq z^t$ and $u(z^t) = \bar{u}(a^t) - k(a^t)$ where $\bar{u}(a^t)$ and $k(a^t)$ are determined by the participation and incentive constraints as follows:

$$\frac{1-\delta^t}{1-\delta}\bar{u}(a^t) - \delta^{t-1}P(e^t)k(a^t) = G(e^t) + \bar{U} \tag{13}$$

$$(P(z^t|a^t) - P(z^t|e^t))\delta^{t-1}k(a^t) = G(e^t) - G(a^t). \tag{14}$$

By repeating this process, we obtain $u(a^t)$ and $k(a^t)$ for all a^t with $G(e^t) > G(a^t)$. Importantly, by (13), the principal's average cost per period, $\bar{u}(a^t)$, and the punishment given to the agent, $k(a^t)$, are positively related. Therefore, the two maxima below

$$\bar{u}^* = \max_{a^t} \bar{u}(a^t), \;\; k^* = \max_{a^t} k(a^t), \;\; \text{s.t. } G(a^t) < G(e^t)$$

are simultaneously attained by some effort sequence b^t. The profile

$$u(y^t) = \begin{cases} \bar{u}^* & \text{if } y^t \neq z^t \\ \bar{u}^* - k^* & \text{if } y^t = z^t \end{cases}$$

then is a desired solution to Program *(Q)* as given in the proposition. □

[14] If $G(e^t) \leq G(a^t)$ then the solution is trivial: $u(y^t) = \bar{u}$, $\forall y^t$. Namely, the low cost effort sequence e^t can be implemented without any incentive pay.

Proof of Proposition 6.

It suffices to prove the result for two separate review phases as the general case then follows by induction. If the principal releases all performance information by the end of period t, then an optimal contract is given in Eq. (9) for some effort sequence b^t.

Suppose that the principal divides the t dates into two phases: the first from date 1 to date $1 \leq \tau < t$ and the second from $\tau + 1$ to t. Decompose each action sequence a^t into (a^1, a^2), where $a^1 = (a_1, \ldots, a_\tau)$ and $a^2 = (a_{\tau+1}, \ldots, a_t)$, and define the following probabilities

$$P_1 = P(z_1, \ldots, z_\tau | e^1) \text{ and } P_2 = P(z_{\tau+1}, \ldots, z_t | e^2)$$

$$Q_1 = P(z_1, \ldots, z_\tau | b^1) \text{ and } Q_2 = P(z_{\tau+1}, \ldots, z_t | b^2).$$

Let $\overline{U}_1$ and $\overline{U}_2$ be the expected utilities, discounted to date 1, that the agent will receive in the two phases 1 and 2 respectively. To satisfy the participation constraint at date 1, the principal sets $\overline{U}_1 + \overline{U}_2 = \overline{U}$. The relative magnitudes of $\overline{U}_1$ and $\overline{U}_2$ are unimportant for the result. Thus, the principal effectively solves two separate problems.

By Proposition 5, the principal's cost per date for implementing e^i in phase i equals $u^{-1}(\bar{u}_i)$, where

$$\bar{u}_i = \frac{1-\delta}{1-\delta^{t_i}} \left(G(e^i) + \bar{U}_i + P_i K_i \right), \text{ for } t_1 = \tau,\, t_2 = t - \tau,$$

and K_i is the size of the punishment. Again, all payoff terms are discounted to date 1.

Suppose $G(e^i) > G(b^i)$ for both $i = 1, 2$. By fully postponing performance review to date t, the principal pays an expected social cost equal to

$$\begin{aligned}
P(z^t|e)K &= \frac{P(z^t|e^t)(G(e^t) - G(b^t))}{P(z^t|b^t) - P(z^t|e^t)} \\
&= \frac{P_1 P_2 (G(e^1) + G(e^2) - G(b^1) - G(b^2))}{Q_1 Q_2 - P_1 P_2} \\
&= \frac{P_1 P_2 (G(e^1) - G(b^1))}{Q_1 Q_2 - P_1 P_2} + \frac{P_1 P_2 (G(e^2) - G(b^2))}{Q_1 Q_2 - P_1 P_2} \\
&< \frac{P_1 (G(e^1) - G(b^1))}{Q_1 - P_1} + \frac{P_2 (G(e^2) - G(b^2))}{Q_2 - P_2} \\
&\leq P_1 K_1 + P_2 K_2.
\end{aligned}$$

Here the penultimate inequality follows since $Q_i > P_i$ and the last one follows because

$$\frac{P_i (G(e^i) - G(b^i))}{Q_i - P_i}$$

is the expected social cost for implementing e^i against the alternative sequence b^i, which is no more than the cost required to implement e^i against *all* possible deviations a^i.

If $G(b^1) \geq G(e^1)$ (e.g. $b^1 = e^1$), then

$$P(z^t)K \leq \frac{P_2(G(e^2) - G(b^2))}{P_2 - Q_2} \leq P_1K_1 < P_1K_1 + P_2K_2$$

since $P_2K_2 > 0$. The result is the same if $G(b^2) \geq G(e^2)$.

It then follows that

$$\frac{1-\delta^t}{1-\delta}u(\bar{c}) < \frac{1-\delta^\tau}{1-\delta}u(\bar{c}_1) + \frac{1-\delta^{t-\tau}}{1-\delta}u(\bar{c}_2) \leq \frac{1-\delta^t}{1-\delta}u\left(\frac{1-\delta^\tau}{1-\delta^t}\bar{c}_1 + \frac{1-\delta^{t-\tau}}{1-\delta^t}\bar{c}_2\right)$$

and hence

$$\frac{1-\delta^t}{1-\delta}\bar{c} < \frac{1-\delta^\tau}{1-\delta}\bar{c}_1 + \frac{1-\delta^{t-\tau}}{1-\delta}\bar{c}_2.$$

Namely, the total cost is lower with full delay than that with intermediate review. □

Proof of Proposition 7*.*

The principal chooses w and $c(y) \in [0,w]$, $\forall y = (y_1, y_2)$, to minimize w subject to the incentive constraints,

$$\sum_y c(y)p(y|e) - 2c \geq \sum_y c(y)p(y|a) - c(a_1 + a_2), \text{ for } a = (1,0), (0,1), \text{ or } (0,0).$$

First, total success should always be rewarded with the full reward, i.e. $c(h,h) = w$, and total failure should never be rewarded, i.e. $c(\ell,\ell) = 0$. Otherwise, the incentive constraints can all be relaxed because $p(h,h|e) - p(h,h|a) > 0$ and $p(\ell,\ell|e) - p(\ell,\ell|a) < 0$ for all $a \neq e$.

Letting $x = c(h,\ell)$ and $z = c(\ell,h)$, the incentive constraints become

$$p^2w + p(1-p)(x+z) - 2c \geq q^2w + q(1-q)(x+z)$$

$$p^2w + p(1-p)(x+z) - 2c \geq pqw + p(1-q)x + (1-p)qz - c$$

$$p^2w + p(1-p)(x+z) - 2c \geq pqw + q(1-p)x + (1-q)pz - c.$$

Let $x + z = kw$, where $k \in [0, 2]$. Then the first constraint can be rewritten as

$$w((p+q) + (1-p-q)k) \geq \frac{2c}{p-q}. \tag{15}$$

And the second and third constraints add up to

$$w(2p+(1-2p)k) \geq \frac{2c}{p-q}. \tag{16}$$

The principal's cost can not go up if we only consider these two constraints.

Define $f(k) \equiv p+q+(1-p-q)k$ and $g(k) \equiv 2p+(1-2p)k$. Note that $f(k)-g(k) = (p-q)(k-1)$, and given $p+q<1$ and $p>\frac{1}{2}$, $f(k)$ increases in k and $g(k)$ decreases in k.

Therefore, for $k \in [0,1]$, $f(k) \leq g(k)$ and hence

$$w \geq \frac{2c}{p-q} \cdot \frac{1}{f(k)} \geq \frac{2c}{p-q} \cdot \frac{1}{f(1)} = \frac{2c}{p-q}.$$

For $k \in [1,2]$, $f(k) \geq g(k)$ and hence

$$w \geq \frac{2c}{p-q} \cdot \frac{1}{g(k)} \geq \frac{2c}{p-q} \cdot \frac{1}{g(1)} = \frac{2c}{p-q}.$$

In conclusion, the minimum cost equals $\frac{2c}{p-q}$, which is achieved by using separate contracts for each period. □

Proof of Proposition 8 . **(a):** Suppose $p^t - q^t > p^{t-1} - q^{t-1}$. First consider the relaxed problem in which the agent has only two choices: $e = (1,\dots,1)$ and $a' = (0,\dots,0)$. The principal chooses plan $c(y) \leq w$ to minimize cost w subject to the constraint that the agent weakly prefers e to a':

$$\sum_y (P(y|e) - P(y|a'))c(y) \geq tc.$$

We show that $P(y|e) \leq P(y|a')$ for all $y \neq (h,\dots,h)$ so the principal should pay the agent a positive amount only if $y = (h,\dots,h)$. This is because for every y in which h occurs $k<t$ times,

$$\frac{P(y|e)}{P(y|a')} = \frac{p^k(1-p)^{t-k}}{q^k(1-q)^{t-k}} < \frac{1-q}{1-p} \cdot \frac{(1-p)^{t-k}}{(1-q)^{t-k}} \leq 1.$$

The first inequality is by the following:

$$\left(\frac{p}{q}\right)^k \leq \left(\frac{p}{q}\right)^{t-1} < \frac{1-q}{1-p}, \ \forall k = 1,\dots,t-1, \tag{17}$$

which follows from Hypothesis (a). Therefore the optimal bonus B is determined by $(p^t - q^t)B = tc$, as stated in the proposition.

It only remains to verify that the agent has no incentive to choose any other effort sequence involving high effort in $1 \leq n \leq t-1$ out of t periods, which amounts to show that, for $n = 1,\dots,t-1$,

$$B(p^t - p^n q^{t-n}) \geq tc - nc,$$

i.e.

$$B = \frac{tc}{p^t - q^t} \geq \frac{(t-n)c}{p^t - p^n q^{t-n}}. \tag{18}$$

To prove this relation, define $P(n) = p^n q^{t-n}$, and observe that for $n = 1, ...t-1$,

$$P(n+1) - P(n) > P(n) - P(n-1).$$

This relation reflects the complementarity between high efforts in each period in raising the probability of output path $(h, ..., h)$: high effort in each additional period has a bigger impact. It then follows that for $n = 1, ..., t-1$,

$$\frac{P(t) - P(0)}{t} < \frac{P(t) - P(n)}{t-n},$$

which implies (18).

In conclusion, $B = tc/(p^t - q^t)$ indeed is the principal's optimal cost for a t-period contract. Since by (17),

$$p^k - q^k > p^{k-1} - q^{k-1},\ \forall k = 1, ..., t-1,$$

it follows that the principal's total cost equals $kc/(p^k - q^k)$ for a k-period relationship and the *per-date* cost $c/(p^k - q^k)$ goes down as the contract length k grows, up to t.

(b): Suppose $p < \frac{1}{t}$. This time, consider the relaxed problem in which the agent has $t+1$ alternatives: the principal's desired sequence $e = (1, ...1)$ and t other sequences 0_k, each involving exactly one deviation from e at date $k = 1, ..., t$. To solve the problem, further relax the t individual incentive constraints to a single pooled constraint:

$$\sum_y \left(tP(y|e) - \sum_k P(y|0_k) \right) c(y) \geq tc.$$

Consider an output path y in which high output h occurs $n > 0$ times. For $k = 1, ..., t$, each probability $P(y|0_k) = qp^{n-1}(1-p)^{t-n}$ or $p^n(1-p)^{t-n-1}(1-q))$ depending on whether signal h or ℓ occurs in date k. Then

$$\begin{aligned} tP(y|e) - \sum_k P(y|0_k) &= tp^n(1-p)^{t-n} - (nqp^{n-1}(1-p)^{t-n} + (t-n)p^n(1-p)^{t-n-1}(1-q)) \\ &= p^{n-1}(1-p)^{t-n-1}\left(n(p(1-p) - q(1-p)) + (t-n)(p(1-p) - p(1-q))\right) \\ &= p^{n-1}(1-p)^{t-n-1}(p-q)(n-tp) \\ &> 0, \end{aligned}$$

for all $n = 1,...,t$, because $p < \frac{1}{t}$.

Therefore, the principal should pay the agent the maximum amount K for all output path except $(\ell,...,\ell)$ and K is determined by

$$-K((1-p)^t - (1-p)^{t-1}(1-q)) = c,$$

which gives rise to the desired expression in the proposition.

To show that the agent has no incentive to deviate from e, we verify that for $n = 1,...,t$ deviations,

$$-K((1-p)^t - (1-p)^{t-n}(1-q)^n) \geq tc - (t-n)c,$$

which is equivalent to

$$K = \frac{c}{(1-p)^{t-1}(1-q) - (1-p)^t} \geq \frac{nc}{(1-p)^{t-n}(1-q)^n - (1-p)^t}. \tag{19}$$

Define $P(n) = (1-p)^{t-n}(1-q)^n$ and observe that for $n = 1,...,t-1$,

$$P(n+1) - P(n) > P(n) - P(n-1).$$

This relation reflects the complementarity between low efforts in each period in raising the probability of the bad output path $(\ell,...,\ell)$: low effort in each additional period increases the probability by a greater amount . It then follows that for $n = 1,...,t$,

$$P(1) - P(0) < \frac{P(n) - P(0)}{n},$$

which is equivalent to (19).

In summary, the principal's minimum cost in t-period contracting equals K. By similar arguments, for $k = 1,...,t-1$, the minimum cost for a k-period relationship equals

$$\frac{c}{(1-p)^{k-1}(p-q)}.$$

To show that the per date cost decreases as k increases up to t, we need only show that for $k = 2,...,t$,

$$k(1-p)^{k-1} > (k-1)(1-p)^{k-2},$$

which is equivalent to $p < \frac{1}{k}$ and it is implied by Hypothesis (b). $\square$

References

Abreu, Dilip, Paul Milgrom, and David G. Pearce, "Information and Timing in Repeated Partnerships," *Econometrica*, November 1991, *59(6)*, pp. 1713–33.

Compte, Olivier, "Communication in Repeated Games with Imperfect Private Monitoring," *Econometrica*, 1998, *66(3)*, pp. 597–626.

Fuchs, William, "Contracting with Repeated Moral Hazard and Private Evaluations," *American Economic Review*, 2007, *97(4)*, pp. 1432–48.

Holmström, Bengt and Paul Milgrom, "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica*, 1987, *55 (2)*, pp. 303–28.

_ **and** _ , "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership and Job Design," *Journal of Law, Economics and Organization*, 1991, *Special Issue 7*, pp. 24–52.

Innes, Robert D., "Limited Liability and Incentive Contracting with Ex-ante Action Choices," *Journal of Economic Theory*, 1990, *52(1)*, pp. 45–67.

Jensen, Michael C. and Kevin J. Murphy, "Performance Pay and Top-Management Incentives," *Journal of Political Economy*, 1990, *98 (2)*, pp. 225–64.

Kandori, Michihiro and Hitoshi Matsushima, "Private Observation, Communication and Collusion," *Econometrica*, May 1998, *66(3)*, pp. 627–52.

Laffont, Jean-Jacques and David Martimort, *The Theory of Incentives: The Principal-Agent Model*, Princeton University Press, 2002.

Lazear, Edward P., "Agency, Earnings Profiles, Productivity, and Hours Restrictions," *American Economic Review*, 1981, *71(4)*, pp. 606–20.

Levin, Jonathan, "Relational Incentive Contracts," *American Economic Review*, 2003, *93(3)*, pp. 835–57.

MacDonald, Glenn and Leslie M. Marx, "Adverse Specialization," *Journal of Political Economy*, 2001, *109 (4)*, pp. 864–99.

MacLeod, W. Bentley, "Optimal Contracting with Subjective Evaluation," *American Economic Review*, 2003, *93*, pp. 216–40.

_ , "Reputations, Relationships, and Contract Enforcement," *The Journal of Economic Literature*, 2007, *45 (3)*, 595–628.

Mirrlees, James A., "Notes on Welfare Economics, Information and Uncertainty," in M. Balch, D. McFadden, and S. Wu, eds., *Essays in Equilibrium Behavior under Uncertainty*, North-Holland, 1974.

Prendergast, Canice, "The Provision of Incentives in Firms," *Journal of Economic Literature*, 1999, *36(1)*, pp. 7–63.

Radner, Roy, "Repeated Principal-Agent Games with Discounting," *Econometrica*, 1985, *53(5)*, pp. 1173–98.