Contracting with Repeated Moral Hazard and Private Evaluations

William Fuchs*

May 18, 2006

Abstract

A repeated moral hazard setting in which the Principal privately observes the Agent's output is studied. It is shown that there is no loss from restricting the analysis to contracts in which the Agent is supposed to exert effort every period, receives a constant efficiency wage and no feedback until he is fired. The optimal contract for a finite horizon is characterized, and shown to require burning of resources. These are only burnt after the worst possible realization sequence and the amount is independent of both the length of the horizon and the discount factor (δ). For the infinite horizon case a family of fixed interval review contracts is characterized and shown to achieve first best as $\delta \to 1$. The optimal contract when $\delta << 1$ is partially characterized. Incentives are optimally provided with a combination of efficiency wages and the threat of termination, which will exhibit memory over the whole history of realizations. Finally, Tournaments are shown to provide an alternative solution to the problem.

Mainly due to the availability of data, most of the contracting literature has focused on occupations such as chief executive officers, golfers, mutual fund managers, tree cutters, windshield installers, and so on, for whom it is possible to construct objective measures of output. As pointed out by Canice Prendergast (1999), the analysis of these occupations has resulted in interesting insights on how incentives work in theory and practice. However, as

*I'm particularly indebted to Andy Skrzypacz and Tom Sargent for their valuable advice and suggestions. I thank Manuel Amador, Heski Bar-Isaac, Simon Board, Vinicius Carrasco, Ignacio Esponda, Narayana Kocherlakota, Ed Lazear, Jon Levin, Gustavo Manso, Paul Milgrom, Tomasz Sadzik, Yuliy Sannikov, Ilya Segal, Ennio Stacchetti and Pierre-Olivier Weill for their useful comments. Support from the B.F Haley and E.S Shaw Fund at Stanford University and the Mariscal the Ayacucho Fellowship is gratefully acknowledged.

he says, most people don't work in these types of jobs. Most workers are instead evaluated using subjective measures of performance.

If a measure of performance is subjective then it is inherently private. For example, if an analyst prepares a research report for his boss, the boss will form a subjective assessment of how good or valuable he finds the report after reading it. Furthermore, his opinion is private and he will only reveal it truthfully if it is in his own interest. Some procurement relationships, in particular services, also naturally fit this framework.[1] Our goal in this paper is to characterize optimal contractual arrangements in these environments where a moral hazard problem on the Agent's side is combined with private evaluations on the Principal's side.

We show in this paper how many features observed in labor markets – such as wage compression, involuntary unemployment, efficiency wages, review periods and tournaments– show up as alternatives to overcome the Agent's moral hazard problem when the Principal monitors the output privately.

In a related paper, Bentley W. MacLeod (2003) analyzed a one-shot model of a Principal who privately observes the performance of an Agent (who, in turn, first privately chooses his effort level). We first show that the intuition from MacLeod that to provide incentives for both players we need to break budget balance, extends naturally to the dynamic setup in the form of the need to destroy surplus either directly or by inefficient termination of the project. We then establish the main new result, that the dynamic contract provides no feedback to the agent until the end (endogenous of fixed) of the contract.

We start by characterizing the optimal contract for a the finite horizon problem. We show it is optimal to wait until the end and consider all the outcomes together before performing an evaluation. The Agent then receives a common high wage for all possible outcome histories except for the case in which the outcome is low every period. In that case he receives only part of the compensation, the rest of it is burnt. Hence, compensation is both compressed within a period, as characterized in MacLeod (2003), and also across histories. The compression within a period is driven by the strength of the statistical test used to determine deviations. Instead, the pooling across periods is related to the concept of "reusability of punishments" introduced by Abreu et al (1991). The idea is that one punishment can simultaneously be used to provide incentives for the Agent to exert effort over many periods. Since punishments entail inefficient money burning, "reusing" punishments is efficient. In order to take advantage of the reusability of punishments it is important not

[1] For example, if a supplier ships goods across the Atlantic, the condition in which those goods are received may be only privately observed by the buyer.

to give interim feedback to the Agent about his performance.

In the second part of the paper we consider the case when the horizon is infinite. One motivation to study this case is that it allows us to endogenize money burning through termination. Termination destroys the remaining surplus that could have been generated by the relationship. Explicitly studying how value is destroyed in the repeated game also raises several important issues that were not present in the finite horizon model. The fact that there is no 'last' period and that all the burning can no longer take place at the end of the game introduces new and important aspects to the problem. The question of the optimal timing for the release of information no longer has a straightforward solution. Even if the Principal sends no explicit messages to the Agent, the simple fact that he has not been fired provides the Agent with some information about the outcome realizations. This adds a series of new concerns when determining the optimal firing decisions.

From a technical perspective, this is an infinitely repeated game with private monitoring. Unfortunately, such games lack a tractable recursive structure. There is no equivalent to the tools developed by Abreu et al. (1990) or Yuliy Sannikov (2004) applicable to these types of games. Hence, characterizing the equilibrium value set (or even a given equilibrium) poses a great challenge.[2]

In spite of these difficulties, we show that in order to characterize the equilibrium value set, we can restrict our attention to a specific class of contracts. This class entails the Agent exerting effort every period, receiving a fixed efficiency wage and no bonuses nor feedback on his past performance until he is fired.

These results resonate well with the evidence presented by George Baker et al. (1994) that only 5% of workers claim to receive performance pay in the form of commissions or piece rate contracts and that only 25% claim to have bonuses based on subjective measures of performance. Also consistent with our results is Prendergast (1999) who documents the reluctance of managers to rate employees, especially when it impacts compensation. Furthermore, MacLeod and Daniel Parent (1999) show that piece rate or commission contracts are mainly used for jobs with relatively few tasks and with more verifiable output measures such as mutual fund managers, tree cutters and windshield installers. They also find that the use of bonus pay is more frequent when unemployment is low, this is consistent with the use of termination as an incentive devise.

Then we use a variation approach to provide a further characterization of the optimal firing rule. This method consists in showing that if a given contract does not have certain prop-

[2]See Kandori (2002) for a discussion of the difficulties associated with the analysis of private monitoring games.

erties, we can modify it slightly to obtain a weak improvement. This allows us to sidestep some of the difficulties associated with the absence of a tractable recursive representation; in particular, the fact that arbitrarily complex deviations by the Agent must be considered.[3] This technique allows us to construct a partial ordering over histories and show that the termination decisions exhibit memory. Outcomes from past periods affect termination decisions for the whole future. Therefore, optimal contracts cannot be replicated by short-term contracts.

Jonathan Levin (2003), who first analyzed this infinite horizon problem, side-stepped these difficulties by restricting his attention to contracts with the "Full Review Property". This property restricts the analysis to equilibria in which the Principal truthfully reveals the outcome to the Agent every period. Hence, the Principal does not retain any private information from period to period. We show that these contracts are generally suboptimal, especially, when players are patient. For comparison with Levin (2003) we also consider a subclass of contracts in which the Agent's performance is only reviewed at fixed intervals of length T. We refer to these conracts as T-period review contracts.[4] In fact, these contracts are easy to characterize and hence not only allow us to show that a strict improvement over the Full Review contracts is possible, but also that if the players are sufficiently patient they can achieve close to first best outcomes. This result cannot be achieved within the family of contracts with the "Full Review Property". Finally, these contracts capture the common practice of many firms and organizations to review their workers at predetermined intervals of time. Our analysis describes some of the tradeoffs involved in the choice of the optimal review length.

The best way to relate this paper to the more general literature on moral hazard problems is to classify the models along the following two properties. Is the output verifiable by a third party? Is the output common knowledge between the Principal and the Agent? The first papers in contract theory such as Bengt Holmstrom (1979) and Steven Shavell (1979) focused on the case in which the output was common knowledge and, more importantly, could be verified by a court of law. They studied static problems and their main concern was the risk sharing vs. incentives trade-off. In many circumstances output measures fail to be verifiable by the courts. This problem was first analyzed Clive Bull (1987) and MacLeod and James M. Malcomson (1989). They showed that if the performance is common knowledge, continuation values in long term relationships could be used to provide incentives. David Pearce and

[3] The main technical difficulty is that in general it is not sufficient to check only one-step deviations by the agent.

[4] The case $T = 1$ corresponds to contracts with the "Full Review Property".

Ennio Stacchetti (1998) further showed how the verifiable and non-verifiable measures of performance could be combined in optimal contracts. The typical examples cited in the literature of situations in which output cannot be verified by the courts generally involve complex output measures such as the quality of teaching or customer care. It is natural to think that the assumption of common knowledge over performance might also fail in these types of situations. Lifting this assumption leads to the work of MacLeod (2003), Levin (2003) and this paper. The new insight we present is the role of dynamic long-term contracts and in particular of feedback (or the lack of it) on the optimal provision of incentives.

This paper is also related to the literature on efficiency wages. The use of efficiency wages in macro models such as Carl Shapiro and Joseph E. Stiglitz (1984) has been criticized for lacking proper microeconomic foundations. We provide such foundations and show that the contracts studied by Shapiro and Stiglitz (1984) are actually optimal in an environment where the Principal privately observes the Agent's output. There are a number of empirical papers that suggest there is evidence that efficiency wage contracts are used in practice.[5] In addition to providing a stronger microfoundation for the use of efficiency wages, our model suggests new ways to test for their use in practice.

The rest of the paper is organized as follows: Section I presents the model. Section II studies the finite horizon problem. Section III extends the analysis to the infinite horizon. Section IV explores extensions of the basic model to allow for multiple levels and stochastic effort and also analysis the use of tournaments. Conclusions and suggestions for future research can be found in Section V. Finally, omitted proofs can be found in Appendix B.

I The Model

Two risk neutral players, an Agent and a Principal are matched at time zero and have the opportunity to trade at dates $t \in \{0, 1, ..., T\}$. We will consider both the case where T is finite and $T = \infty$. At $t = 0$, the Principal and the Agent agree on a compensation schedule and a termination policy.[6] Part of this agreement is determined by a long term contract ω signed at time zero. Its terms depend only on verifiable information that is enforceable by a third party. The rest of the agreement corresponds to a perfect Bayesian equilibrium (a self-enforcing contract) of the post-contractual repeated game Γ_ω we now describe.

[5]See for example Alan B. Krueger and Lawrence H. Summers (1988) or Peter Cappelli and Keith Chauvin (1991).

[6]The way the surplus is divided is not important for the analysis so we won't model it explicitly. We think about it as the outcome of some efficient bargaining.

The figure below depicts the order of play in each period of the game while the Agent is still employed. Informally, the timing is as follows: first, the Agent chooses whether to exert effort or not, conditional on his information. Next, the Principal privately observes the output realization y_t and a random variable ϕ_t. He then chooses a message m_t to send to the Agent. This message is verifiable and implies a bonus b_t he must pay to the Agent. The Principal also pays a base wage w_t that is not contingent on output but can be contingent on a public randomization device x_t. Finally, simultaneous decisions are made in every period on whether to continue the relationship into the future or not.[7]

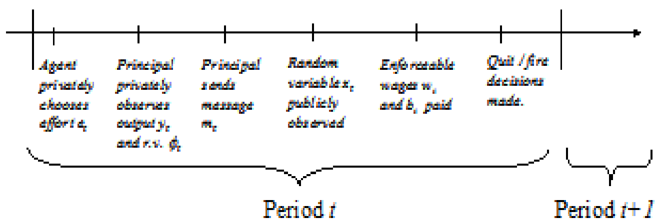

Given the particular information structure of the game we will define three different histories. First, we have the set of verifiable histories H_t^V where t denotes the length of the history. A typical element of this set is $h_t^V = \{m^t, k^t, \mathcal{A}^t, x^t\}$. Superscripts are used to denote sequences of realizations and subscripts are used to denote a particular realization. k_t is an indicator function that equals one if the Agent decides to stay in the job at time t. Similarly, $\mathcal{A}_t$ is an indicator function that equals one if the Principal decides to continue employing the Agent at time t. The exogenous public randomization device $x_t \sim U[0,1]$ is independent of all other elements of the game and identically distributed over time. Next, we have the set of Agent's private histories H_t^A. A typical element in this set is $h_t^A = \left\{h_t^V, e^t\right\}$. Finally, we have the set of Principal's private histories H_t^P. A typical element in this set is $h_t^P = \left\{h_t^V, y^t, \phi^t\right\}$. Where $\phi_t \sim U[0,1]$ is independent of all other elements of the game and identically distributed over time. For all of the histories defined above we take the convention that $h_{-1}^j = \emptyset$.

We can now formally define an enforceable contract $\omega \in \Omega$ as two sequences of functions: base wages and bonuses. Base wages w_t are random functions of all the verifiable history except the history of messages:

$$w_t : \left(H_{t-1}^V / M^{t-1}\right) \times [0,1] \to \mathbb{R} \quad for\ t = 0, ..., T\ .^8$$

[7]The assumption that these decisions are simultaneous could be easily modified without affecting the results.

[8]The $\times[0,1]$ stands for the public randomization device x_t.

Note that this formulation of the wages allows for severance pay. Bonuses are additionally functions of the messages sent by the Principal:

$$b_t : H_{t-1}^V \times M_t \times [0,1] \to \mathbb{R}^+ \quad for\ t = 0, ..., T\ .$$

The assumption that $b \geq 0$ is without loss of generality since we could always rescale $\{w_t\}$; also, without loss of generality we assume that for all t there is some m_t such that $b\left(h_{t-1}^V, m_t, x_t\right) = 0$.[9] This is useful because we want to allow the Principal to send messages at no cost, for example if those messages contain no information.

Each period if still employed, the Agent has a choice of exerting effort or not, $e_t : H_{t-1}^A \to \{0, 1\}$.[10] We assume no effort is costless and exerting effort is costly. Formally, we denote the cost of effort by $c(e)$ where, $c(0) = 0$ and $c(1) = c > 0$.

Output is also binary, $y_t \in Y_t = \{L, H\}$ and is privately observed by the Principal. Its realizations are stochastic and depend on the level of effort in the following way:

$P(y\|e)$	$y=L$	$y=H$
$e=0$	$1-q$	q
$e=1$	$1-p$	p

We assume that effort increases the likelihood of a high outcome:

$$1 > p > q > 0\ .$$

We use the common support assumption since it facilitates some of the discussion, in particular when considering off equilibrium behavior. The cases $p = 1$ and $q = 0$ are discussed later.

The expected output conditional on no effort is zero.[11] Lastly, the per period expected surplus from exerting effort will be denoted by s and assumed to be strictly positive:

$$\mathbb{E}[y \mid e = 1] - c = s > 0\ .$$

Upon observing the realization of y_t and ϕ_t, the Principal sends a public and verifiable announcement m_t, where $m_t : H_{t-1}^P \times Y_t \times [0,1] \to M_t$. The set M_t is the *algebra* generated

[9] Let $\hat{b}$ be the minimun bonus for all possible messages at time t. Let $\tilde{w}_t = w_t + \hat{b}\ \forall w_t$ and $\tilde{b}\left(h_{t-1}^V, m_t, x_t\right) = b\left(h_{t-1}^V, m_t, x_t\right) - \hat{b}\ \ \forall b_t$

[10] We consider mixed strategies and multiple levels of effort for the Agent in Section A.

[11] This is not purely a normalization since it implies that no effort no pay is identical to termination in terms of the surplus generated.

by Y^t.[12] That is, the Principal can choose to reveal only partially or not reveal at all the history of outcomes to the Agent.

The last event in each period is the simultaneous decision whether to continue the relationship or terminate it. Termination occurs if at least one of the parties decides to terminate the relationship. The Agent's decision whether to stay or not is denoted k_t where, $k_t : H^A_{t-1} \times e_t \times x_t \times M_t \to \{0,1\}$.[13] The Principal offers the Agent to remain employed with probability a_t where, $a_t : H^P_{t-1} \times Y_t \times M_t \times [0,1] \to [0,1]$.[14] If termination occurs both parties receive their outside options which are normalized to zero.[15] If a relationship has been terminated, the parties never trade again (the Agent exerts no more effort for the Principal) only severance pay transfers may still take place.

Let $\sigma^A_\omega = [\{e_t\}, \{k_t\}]$ denote the Agent's strategy conditional on having signed contract ω. Similarly let $\sigma^P_\omega = [\{m_t\}, \{a_t\}]$ denote the Principal's strategy conditional on having signed contract ω. We use Σ^A_ω *and* Σ^P_ω to denote the set of strategies for the Agent and the Principal respectively. We will use $\sigma_\omega = \left(\sigma^A_\omega, \sigma^P_\omega\right)$ to denote a strategy pair. The full strategies for the players have an additional element χ^i $i = A, P$ which denotes the choice of accepting a given contract $\omega \in \Omega$ with the self-enforcing agreement (to be defined in the next section) that $\sigma_\omega \in \Sigma^A_\omega \times \Sigma^P_\omega$ will be played in the continuation game. $\chi^i : \Omega \times \Sigma^A_\omega \times \Sigma^P_\omega \to \{0,1\}$. The complete strategy profile for each player is $\sigma^i = [\chi^i, \sigma^i_\omega]$. Denote by σ a strategy profile pair $\left(\sigma^A, \sigma^P\right)$ of the whole game.

We use $\delta \in (0,1)$ to denote the discount factor which is assumed to be the same for both players. The Agent maximizes the expected discounted flow of wages minus the effort cost.

$$V = \mathbb{E} \sum_{t=0}^{T} \delta^t (w_t + b_t - c(e_t))$$

The Principal maximizes the expected discounted value of output y_t net of the wages.[16]

$$F = \mathbb{E} \sum_{t=0}^{T} \delta^t (y_t - w_t - b_t)$$

[12] Alternatively: $m_t : H^P_{t-1} \times Y_t \times [0,1] \to \Delta M_t$. The set M_t is the *algebra* generated by Y^t and ΔM_t denotes the simplex over M_t.

[13] Formally $k_t : H^A_{t-1} \times \{0,1\} \times [0,1] \times M_t \to \{0,1\}$.

[14] Note, that we allow for mixed termination strategies for the principal but not for the agent. This is shown in Section 3 not to be a restricitve assumption.

[15] We could embed this model into a matching model such as that in Shapiro and Stiglitz (1984). What we need is for there to be quasi rents perceived when matched.

[16] This formulation implicitly rules out the possibility that part of the wage paid could be destroyed before it is received by the Agent. We relax this assumption in Section 4.

If the players follow a strategy profile σ we can define the values for the players from starting the game with a history h_{t-1}^{P} as follows:

For the Agent:

$$\begin{aligned}&V_t\left(h_{t-1}^{P},\sigma\right)\\=\ &\mathbb{E}\left[\sum_{j=t}^{T}\delta^{(j-t)}\left(w_j+b_j-c\left(e_j\right)\right)|h_{t-1}^{P},\sigma\right]\ ,\end{aligned}$$

and for the Principal:

$$\begin{aligned}&F_t\left(h_{t-1}^{P},\sigma\right)\\=\ &\mathbb{E}\left[\sum_{j=t}^{T}\delta^{(j-t)}\left(y_j-w_j-b_j\right)|h_{t-1}^{P},\sigma\right]\ .\end{aligned}$$

If the Players are following σ then the Principal knows both F_t and V_t. Note that if the Agent deviates, the Principal will be mistaken on his evaluations of F and V.[17] The Agent knows neither since they depend on the Principal's private information. He forms an expectation of V_t based on his information and the strategies being played.

$$v\left(h_{t-1}^{A},\sigma\right)=\mathbb{E}\left[V_t\left(h_{t-1}^{P},\sigma\right)|h_{t-1}^{A},\sigma\right]$$

A Equilibrium Definition

An equilibrium of the post contractual game Γ_{ω}consists of:

A pair of strategies $\left[\sigma_{\omega}^{A},\sigma_{\omega}^{P}\right]\in\Sigma_{\omega}^{A}\times\Sigma_{\omega}^{P}$ and a pair of posterior beliefs about the other player's type $\mu^{A}\left(h_t^{P}|e^{t}\right),\ \mu^{P}\left(h_t^{A}|y^{t}\right)$.[18]

Definition 1 *Best responses:*

A strategy $\sigma_{\omega}^{A}\in\Sigma_{\omega}^{A}$ *is a best response to* σ_{ω}^{P} *after history* h_{t-1}^{A} *if:*

$$\sigma_{\omega}^{A}\in\arg\max_{\tilde{\sigma}_{\omega}^{A}\in\Sigma_{\omega}^{A}}v\left(h_{t-1}^{A},\tilde{\sigma}_{\omega}^{A},\sigma_{\omega}^{P}\right)\ .$$

Similarly, a strategy $\sigma_{\omega}^{P}\in\Sigma_{\omega}^{P}$ *is a best response to* σ_{ω}^{A} *after history* h_{t-1}^{P} *if:*

$$\sigma_{\omega}^{P}\in\arg\max_{\tilde{\sigma}_{\omega}^{P}\in\Sigma_{\omega}^{P}}F\left(h_{t-1}^{P},\sigma_{\omega}^{A},\tilde{\sigma}_{\omega}^{P}\right)\ .$$

[17] As we will explain later this is one of reasons why in general it won't be sufficient to only consider one step deviations for the Agent.

[18] The opponents type is defined by his private information.

Definition 2 *A profile of strategies* $\sigma^*_\omega = \left[\sigma^{*A}_\omega, \sigma^{*P}_\omega\right]$ *and beliefs* μ^A, μ^P *form a Perfect Bayesian Nash equilibrium of* Γ_ω *if and only if:*

- σ^{*A}_ω *is a best response to* σ^{*P}_ω *after every history* h^A_{t-1} *and* σ^{*P}_ω *is a best response to* σ^{*A}_ω *after every history* h^P_{t-1} *given the beliefs* μ^A, μ^P.
- *The beliefs* μ^A, μ^P *are consistent with* σ^*_ω *and updated using Baye's rule, when possible.*

Additionally, an equilibrium for the complete game requires that individual rationality constraints for signing the employment contract ω at time zero are satisfied.

II The Finite Horizon Case

We start by analyzing the finite horizon model which is simpler to tackle and that provides us with some insight which will be useful when we move to the infinite horizon case. We first show a result that is both true of finitely and infinitely repeated relationships:

Proposition 1 *Providing incentives for high effort requires that in expectation some surplus be destroyed if* $e_t = 1$ *and* $y_t = L$

The destruction of value after a low outcome is necessary to have contracts in which the Agent's continuation value is contingent on the realization of outcome while leaving the Principal indifferent. This result extends Proposition 2 in MacLeod (2003) to games of arbitrary horizon.

When dealing with a risk-neutral agent in an environment in which output is verifiable a contract making the Agent the residual claimant could be written and it would deliver first best. Here those types of contracts cannot work because the Principal would always claim output was low. Furthermore, there is no contract that can deliver first best.

Corollary 1 *The first best is not achievable for* $\delta < 1$.

Proof. First note that the unconstrained optimum has the Agent exerting effort every period. By Proposition 1 if the Agent is to exert effort, there must exist some value burning which is inefficient. ■

For the finite horizon case, if we require budget balance there is no enforceable contract that can provide incentives for the Agent to exert effort. This is established in the following Proposition:

Proposition 2 *If $T = 1$, σ^*_ω is essentially unique and it has $e(\omega) = 0$ and $m(y)$ such that $b(m) = 0$.*[19] *Furthermore, for any finite T the game has as a unique equilibrium the repetition of the equilibrium for when $T = 1$ as a result of unravelling.*

If there is no remaining surplus in the relationship, incentives for effort cannot be provided because there is no way to burn resources in case of failure. Hence, we cannot provide contingent pay for the Agent leaving the Principal indifferent as required by Proposition 1.

To get around these negative results, MacLeod (2003) considered the possibility of having the wage being paid by the Principal differ from that received by the Agent, the difference between the two being burnt. He motivates his use of money burning as a shortcut to account for conflict in organizations (inefficient play in the continuation game if the game were to be repeated).[20] Actual money burning is also sometimes used. For example, some baseball teams have the ability to fine their players but the money collected from these fines is not pocketed by the club but rather given away to charity.[21,22] We will model explicitly the infinite horizon game in the next section, in this section we take the latter interpretation. An alternative way to break the balanced-budget between the Principal and the Agent is to introduce a second Agent. We study this setup in Section B.[23] For the remainder of this section we therefore define w^P and w^A as, respectively, the wage paid by the Principal and the amount the Agent receives.

In order to characterize the optimal contract with money burning we first show that without loss of generality we can focus on contracts in which wages are only paid in the last period and the Agent is to exert effort every period.

Lemma 1 *All payments and money burning can be done in the last period.*

Lemma 2 *For every contract with $V + F > 0$ there is a payoff equivalent contract in which the Agent exerts effort in every period.*

[19] "Essentially" since there could be more than one m such that $b(m) = 0$.

[20] MacLeod (2003 pp217): "These costs are generated when individuals either leave the relationship or carry out inefficient actions for several periods..."

[21] We thank Jon Levin for providing this example.

[22] Alternatively, if writting a good recommendation letter had no cost for the Principal and positive value for the Agent instead of burning money the Principal could promise the Agent to write a good recommendation letter if he performs well.

[23] This is reminiscent of the role of introducing a Principal two brake the budget balance between the team members in analysis of the moral hazard in teams problem by Bengt Holmstrom (1982).

These results are important for two reasons. First, when no bounds are placed on the amount of money burning that is allowed, it is sufficient to look at contracts with no termination. Destruction of value, if necessary, can be done via money burning. Second, since wage payments and money burning can be left for the last period, the Agent can be kept completely uninformed of the interim realizations of output. In this case, the Agent's dynamic problem of deciding in which periods to exert effort is identical to a static multitasking problem.

We make use of the static multitasking representation to characterize a contract that minimizes the amount of resources burnt (that maximizes surplus). The amount of resources burnt after a outcome sequence y^T will be denoted by $Z\left(y^T\right)$. We will use $P\left(y^T|\mathbf{e}\right)$ to denote the probability that outcome sequence y^T is realized conditional on the column vector of effort choices $\mathbf{e}$.[24] We use $\mathbf{L}$ to denote a vector of low outcome realizations. Finally we define $\mathbf{c}$ as the row vector of effort costs where the i^{th} element is given by $\delta^{i-1}c$.

Proposition 3 *An optimal contract that provides incentives for the Agent to exert effort for $T+1$ periods satisfies:*

$$\begin{aligned} w_T^A &= w_T^P - Z \quad if \quad y_t = L \ \forall t \\ w_T^A &= w_T^P \qquad\qquad otherwise. \end{aligned}$$

Where:

$$Z = \frac{c}{\delta^T (p-q)(1-p)^T}$$

and:

$$\frac{\sum_{t=0}^{T} \delta^t s_t}{\delta^T} \geq w_T^P \geq \frac{\sum_{t=0}^{T} \delta^t c}{\delta^T} + \frac{(1-p)c}{(p-q)} \ .$$

Furthermore, when $\delta = 1$, this is the only optimal contract.

This Proposition shows that not only is it optimal to have wage compression within a given period as shown in Proposition 6 of MacLeod (2003) but that it is also optimal to have compression across periods and only punish the Agent if the worst outcome is realized in every period. To understand better the intuition behind this result we consider 3 possible ways to provide incentives for the Agent to exert effort for 2 periods. This is shown in the figure below. The first option we consider is to simply treat each period independently and punish the Agent after each low outcome realization. This has an expected efficiency cost of

[24] We denote vectors in bold.

$\frac{c(1-p)}{p-q}(1+\delta)$. We can improve on this by telling the Agent that if he fails in the first period but succeeds in the second his first failure will be forgiven. This requires us to increase the punishment after two consecutive failures but lowers the inefficiency to $\frac{c(1-p)}{p-q}(1+p\delta)$. This is because one stick is providing incentives for the Agent to exert effort both in the first period and also in the second period if in the first period the output was low. We can use this one punishment even more effectively if after the first period we don't inform the Agent if the output was high or low. Now, since the Agent doesn't know if he is in the upper or lower branch of the tree we don't need to punish the Agent if he succeeded in the first period but failed in the second. Since he puts probability $(1-p)$ of being in the lower branch of the tree, the expected punishment he fears he could get if he fails is sufficient to motivate him to exert effort. This leads to an expect total efficiency loss of only $\frac{c(1-p)}{p-q}$. Note that this is the expected present value of money burning even for the one period game. The total efficiency loss is the same regardless of the horizon and the discount factor as formalized in Corollary 2. This implies that to lower the per period expected cost to motivate effort it is optimal to pool all periods together.[25] This is closely related to the results by Abreu et al (1991) on the reusability of punishments, in the sense that one threat of punishment is used to simultaneously deter the temptation to deviate in any period (and in any branch of the tree).

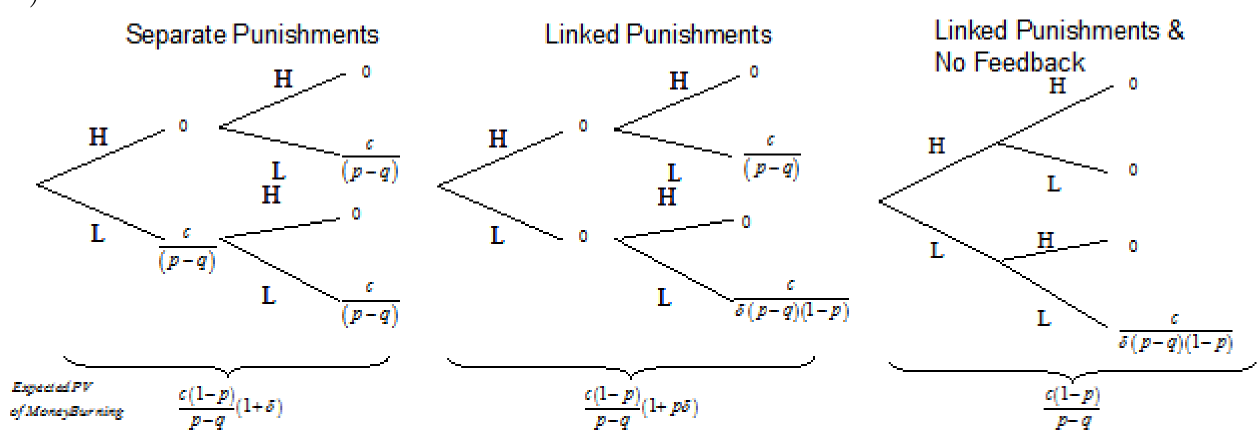

Corollary 2 *The expected present value cost to motivate effort is given by $c\frac{(1-p)}{(p-q)}$. It is independent of both T and δ.*

[25]We might be tempted to think that determining burning on the joint outcome of all periods is beneficial since it reduces the likelihood of burning. Although the likelihood of punishment occurring effectively decreases, the amount to be burnt when it occurs is simultaneously and proportionally increasing, in order to satisfy incentive compatibility eroding away all the gains from the lower frequency of occurrence.

The intuition behind this Corollary is as follows. In the proof of Proposition 3 we show that the only binding constraint is the one for effort in the first period. Since, the cost of the first period's effort, is not discounted by δ, varying δ has no effect. The invariance with respect to T also follows from there only being one constraint binding in the optimal contract. Since using Z^* once we provide incentives for the Agent not to shirk in the first period, the constraints of the remaining periods are automatically satisfied. The likelihood that burning takes place decreases in T but exactly offsetting this effect is the fact that, when required, there will be more burning for larger T. The intuition for this is that the less the Agent believes that the first period's outcome is pivotal in triggering the punishment, the harsher the punishment must be to prevent him from shirking.

Comparative static results on the expected present-value cost to motivate effort are straightforward and intuitive. First, the lower the cost of effort the lower the inefficiency. If it is not too costly for the Agent to exert effort, then it is easier to provide incentives. Next, the smaller the difference in the probabilities of success conditional on exerting or not effort $(p-q)$, the harder it is to provide incentives for the Agent. As $(p-q)$ gets smaller, the outcomes observed if he shirks are closer to the ones if he doesn't, hence there is little reward for his effort and therefore it is hard to provide incentives. On the other hand as p increases the probability that the Agent fails in the *first* period if he exerts effort decreases and hence there will be less burning in expectation. When $p = 1$ the optimal contract changes and has punishment taking place for all histories except $y^T = \mathbf{H}$ but since on the equilibrium path this will be the only history ever observed, this contract achieves first best.

Comparing the inefficiency of the optimal contract relative to the first best we can see that for $p < 1$ the only case in which this contract approximates first best is if both $\delta \to 1$ *and* $T \to \infty$. For all other cases, efficiency is bounded away from the first best. Intuitively this follows because the cost of inducing effort is independent of both T *and* δ and only as we take the value of the relationship to ∞ this cost becomes insignificant in relative terms.

Corollary 2 also implies that for there to be trade between the parties we need:

$$\sum_{t=0}^{T} \delta^t s \geq c\frac{(1-p)}{(p-q)} \ .$$

If the condition above is not met providing incentives for the Agent to exert effort would deliver a negative surplus. A potential resolution to this problem using tournaments is discussed in Section B.

We have shown that for the construction of the optimal contract, communication was not necessary. The following Proposition strengthens this result by showing that when $\delta = 1$

any feedback reduces the achievable efficiency.

Proposition 4 (No Feedback) *When $\delta = 1$, the optimal contract is not implementable if the Principal sends informative messages to the Agent.*

As we showed in the two period example above it is possible to use only one punishment to provide the Agent to exert effort after every possible history. This reusability of punishments is what makes the contract efficient. Releasing information to the Agent generates no benefits and might make the optimal contract infeasible. When $\delta < 1$ since all but the first constraint are slack, the Principal could send some very uninformative messages to the Agent and still be able to implement the optimal contract. When $\delta = 1$ on the other hand, there are no slack constraints and any information would reduce efficiency because the optimal contract would cease to be incentive compatible.

The implication that the Principal should not give any feedback to the Agent might be found to be somewhat extreme. We acknowledge this but nonetheless we consider this to be a useful benchmark to keep in mind. When considering motives that might suggest feedback is beneficial, such as the possibility of the Agent to learn from his mistakes, it is important to realize that feedback might come a cost and this model clearly highlights this factor.[26]

When we analyze the limit of the finite horizon economies as $T \to \infty$ the problem has no solution if we allow for arbitrary amounts of money burning. As $T \to \infty$, the optimal contract calls for burning infinite amounts infinitely far in the future. Therefore, we need to establish some bound on the amount of money that can be burnt. The most natural way we see to do this is to determine it endogenously. We achieve this objective by limiting the amount of value that can be burnt to the potential surplus in the relationship. Furthermore, as mentioned earlier, a motivation for the use of money burning has been the cost of conflicts between Agents and Principals. Leaving potential sabotages aside, the termination of a relationship or shirking by the Agent is the natural way in which to account for these costs. We take this approach in the next section.

III The Infinite Horizon Case

When looking at the optimal equilibria of the infinitely repeated game several new issues arise that are not present in the finite horizon model. In particular, since now the Agent

[26] Fuchs (2005) studies an environment in which there are different and unknown ex-ante match qualities in which communication has a role. In Appendix A we analize another potential motive for communication.

has to receive some information in finite time, we must take into account that the timing of punishments also induces the Agent to learn and potentially makes providing incentives harder in the future. Also, the Agent might have incentives to deviate in order to learn about the Principal's private information. Hence, analyzing the infinite horizon version of the game presents a great challenge. From a technical perspective, the difficulties arise because this is a game with private monitoring and, as pointed out by Michihiro Kandori and Hitoshi Matsushima (1998) and further discussed in Kandori (2002), private monitoring games lack a tractable recursive structure.[27] Hence, characterizing the equilibrium value set of infinitely repeated Private Monitoring Games is an objective that has not yet been achieved. A further complication related to the lack of a recursive structure is that when analyzing a candidate equilibrium it is not in general sufficient to consider only one-step deviations. This hinders any attempt at solving the problem in a direct way using Lagrangian methods. For these reasons, most of the papers in the area have focused on proving Folk Theorems and there has been little or no progress in characterizing optimal equilibria for $\delta << 1$.

To avoid the difficulties arising from the private monitoring, Levin (2003) focuses his analysis on equilibria with the 'Full Review Property'. This condition requires that the Principal reveal to the Agent what he observed after each period. For comparison with Levin (2003) we relax this condition in Section B and consider a class of equilibria in which the Principal only reveals his private information at fixed intervals of length T.[28] Furthermore, T is endogenously determined. We label the contracts with this property *T-period* review contracts. We show that using these contracts we are able to truncate the private histories and give the problem a simple recursive representation. An important property of these contracts is that the attainable per-period average payoffs approximate first best as $\delta \to 1$.

Although these contracts are asymptotically optimal as $\delta \to 1$, they are generally not optimal for $\delta << 1$. The Principal could provide incentives at a lower cost if instead of forgetting the past history after the Agent moves into a new review period he kept an account of it. Once we allow for this the problem becomes much more complex. We loose the recursive representation and we can no longer rule out multiple deviations by the Agent. We tackle this problem in section A and are able to make significant progress towards characterizing optimal contracts. In particular we show that efficiency wage contracts with no feedback are optimal. We also provide a partial characterization of the optimal termination rule. We return in this section to the original model in which we do not allow for exogenous money

[27] Kandori (2002) is the introductory article to a special edition of *The Journal of Economic Theory* on Private Monitoring Games.

[28] The case $T = 1$ corresponds to the contracts with the Full Review Property studied by Levin (2003).

burning.

A Characterization of the Optimal Contract

First we present our main result that shows that contracts that have full wage compression and in which all incentives are provided via efficiency wages are optimal in a general contract space. As a second step we then analyze the properties of optimal termination rule used by the Principal.

Sufficient Class of Contracts:

Theorem 1 *Given any contract* (ω, σ_ω) *that generates values* V *and* F*, we can construct a payoff equivalent contract* $(\hat{\omega}, \hat{\sigma}_\omega)$ *with the following properties:*
(i) The Agent receives a constant wage until he is fired and no bonuses.
(ii) The Agent exerts effort every period until he is fired.
(iii) The Principal gives no feedback (sends no messages) to the Agent.

This strong result provides a possible explanation for the high degree of wage compression observed. It shows that for any contract there is a payoff equivalent contract that eliminates the need for bonuses. Incentives for the agents to exert effort are instead provided via efficiency wages and the threat of termination. Empirical studies such as Krueger and Summers (1988) have found some evidence of the use of efficiency wages with interindustry data. Our results suggest that stronger evidence could potentially be found by looking at occupations rather than industries.

A common criticism of previous models of efficiency wages such as Shapiro and Stiglitz (1984) is that if the Agents were allowed to post performance bonds then this would be a more efficient way in which to provide incentives.[29] We show that even when allowing for performance bonds to be posted these will not be useful if output is monitored privately since the Principal would have a strong incentive to missreport in order to appropriate the bond.

Also worth remarking is the fact that the Principal need not give any informative feedback to the Agent about his performance. The intuition for this result is similar to that in the finite horizon. By keeping the Agent uninformed it is easier to have the firing depend on a longer performance history. In this way we can minimize the need for inefficient termination. In the Extensions Section we discuss a potential role for feedback or communication.

[29] See Malcomson (1999) for a precise discussion and further references.

Optimal Termination: From Theorem 1 we know that incentives for the Agent are

provided through the Principal's threat to terminate the relationship if the output sequence is indicative of cheating by the Agent. What remains to be determined is the efficient termination rule. Optimality requires the minimum possible amount of termination that will provide incentives for the Agent to exert effort.

The Principal chooses a sequence of history dependent functions $a(h^t)$ that determine the probability with which the relationship is continued after history $h^t = \{y^t\}$. The Agent on the other hand chooses a sequence of effort choices $\{e_t\}_{t=0}^{\infty}$ that maximize his expected value. We use $h^\tau \subseteq h^t$ to denote initial subhistories of length τ of a longer history h^t. We use $p(h^t)$ to denote the probability that a given history is realized.

The objective is to maximize expected surplus conditional on the Agent exerting effort every period. Formally, the problem can be stated as follows:

$$\begin{aligned} V &= \max_{\{a(h^t)\}} \sum_{t=0}^{\infty} s \left(\delta^t \left(\sum_{h^t} p\left(h^t\right) \prod_{h^\tau \subseteq h^t} a\left(h^\tau\right) \right) \right) \\ s.t.\ \ \{e_t = 1\}_{t=0}^{\infty} &\in \arg\max \sum_{t=0}^{\infty} \left(y^e - c\left(e_t\right)\right) \left(\delta^t \left(\sum_{h^t} p\left(h^t / e^{t-1}\right) \prod_{h^\tau \subseteq h^t} a\left(h^\tau\right) \right) \right) \end{aligned}$$

The following Proposition provides a partial characterization of the optimal contract using variational arguments.[30] When reading the proof which can be found in the Appendix it is important to keep in mind that altering any element of $\{a(h^t)\}_{t=0}^{\infty}$ impacts incentives in two ways. First, it affects the incentives of all previous periods through the effect the change has on the continuation values. Second, it impacts all future periods incentive constraints because the change affects the beliefs the Agent has about his past performance. Recall that the only information the Agent has about his past performance is the fact that he is still employed. Changing the probability of continued employment after a given history changes the beliefs the Agent has about his likely past performance. Hence, the perturbations we make of the original contract are constructed in such a way that both future and past periods incentives are kept constant but incentives for the period under examination are relaxed. We use the notation (h^t, h^j) to represent the history h^{t+j} where first h^t was observed and then h^j was observed. When h^j contains just one observation we will use (h^t, H) or (h^t, L).

[30] As mentioned before, the problem above is too complex to use Lagrangean methods.

Proposition 5 *A sequence of continuation rules* $\{a(h^t)\}_{t=0}^{\infty}$ *that does not have the following properties, can be weakly improved upon.*

(i) *For all* h^t, *if* $a(h^t, H) < 1$ *then* $a(h^t, L) = 0$.

(ii) *For all* $h^t \times h^j$, *if* $a(h^t, H, h^j) < 1$ *then* $a(h^t, L, h^j) = 0$.

(iii) *For all* t, *if* $a(L^t) > 0$ *then* $a(h^t) = 1$ *for all* h^t.

(iv) *For all* h^t, $V(h^t, H) \geq V(h^t, L)$.

(v) *For all* h^t, *if* $a(h^t, L) = a(h^t, H) = 0$ *then* $a(h^t) = 0$.

Parts (i) & (ii) of the Proposition establish a partial ordering of the outcome histories. Essentially, for two histories h^t and $\tilde{h}^t$, *if* h^t element wise dominates $\tilde{h}^t$ (i.e. if for all $\tau \leq t$ $y_\tau \geq \tilde{y}_\tau$), then $a(h^t) \geq a\left(\tilde{h}^t\right)$. This result is stronger than just an ordering since if $a(h^t) < 1$ then $a\left(\tilde{h}^t\right) = 0$. This tells us that in practice there will actually not be too much mixing used by the Principal. For most histories either the Agent is fired for sure or retained for sure. Furthermore, note that the contract is dependent on the whole past sequence of output realizations. Hence, the optimal contract cannot be replicated by a sequence of short terms contracts.

Part (iii) of the Proposition which follows directly from (i) and (ii) tells us that if the Agent is not fired after t straight failures then for sure he is not fired for the first t periods regardless of his performance. This result is important because for many parameter settings we will then have that the Agent is fired with zero probability for the first few periods. This could allow us to have upward sloping wage schedules as a function of tenure for the first t periods which is a closer to what is observed in practice.

Part (v) states that an Agent who has no future would rather be fired immediately than kept for an additional period. This result would be natural if the Agent knew he had no future since he would clearly have no incentives to exert effort in that case. On the other hand, when he doesn't know his standing, we could have thought that keeping him one more period could be useful because he generates one more period worth of surplus. This is shown not to be the case. Intuitively the idea is that the cost of keeping the Agent around for an additional period is that this weakens the incentives provided in that period since the Agent knows that no matter what he does there is some chance he will get fired anyway.[31] We can

[31] On the other hand, we can make up for this lost period of surplus by reducing the probability we fire the Agent after other (better) histories.

relate this to the problem that firms facing potential bankruptcy have in motivating their employees to work hard.

B Review Contracts

A T-*period* review contract can be described as follows: after the predetermined review length T, the Principal evaluates the Agent's performance. If the performance is favorable, the Agent is offered employment with a clear record at $T+1$. If performance is not satisfactory, then the Agent is fired with probability β. If he is not fired, he is also offered employment with a clear record at $T+1$. Thus the game is partitioned into independent review periods.

Within the literature on private monitoring games, similar type of equilibria appear in Oliver Compte (1998) and Kandori and Matsushima (1998). These, as our own, are based on the original work by Abreu et al (1991), who study the timing of information revelation in an imperfect public monitoring environment. The models differ since in Abreu et al (1991) the information is released by a third party and hence, there are no incentive compatibility considerations regarding the veracity of the reports in their analysis. Nonetheless, we observe that the potential for efficiency gains from the 'reusability of punishments' (attained by extending the length of the review) is still present. We exploit this and show that within this family of contracts we can obtain arbitrarily efficient contracts as $\delta \to 1$.

The T-period review contracts can also be thought of as a repetition of a T period game with money burning, where there is a limited amount of money that can be burnt. Another simplification that the T-period structure gives us is that there are predetermined dates for the release of information. Recall that in the finite-horizon game there was no release of information until the end, this of course cannot be true in an infinite horizon game. The question therefore arises of when is the best time to release information to the Agent. The longer the horizon the greater the efficiency, but the more that needs to be burnt if required. Since this amount is endogenously bounded, it in turn sets an endogenous bound on the review length.

Beyond their analytic convenience and asymptotic efficiency properties, these contracts are appealing because of the common practice in many firms to carry out periodic and predetermined evaluations of their employees. Our analysis does not make full justice of this review system but sheds some light on the frequency with which we might expect reviews to take place as a function of the properties of the task. Hopefully, this could help us understand why although most commonly firms review their employees in an annual or semiannual basis, on one hand some fast food chains do this in a monthly basis while on the other young academic economists on average get reviewed only every three years.

We focus on a particular type of contract that pays a fixed wage every period $w = pH + (1-p)L$ while the Agent is employed. Every T periods the Principal evaluates the Agent. If the output in all the T periods under review has been low, he terminates the relationship with probability β_T. This type of contract is very convenient for tractability reasons, since, as we show in the following Lemma, we only need to check one incentive compatibility constraint for the Agent.

Lemma 3 *With a T-period review contract, if it is not profitable for the Agent to deviate by exerting low effort in the first period, then there are no profitable deviations for the Agent.*[32]

The IC constraint for the Agent in the T-period review contract boils down to:

$$\text{(ICA T)} \qquad (1-p)^{T-1}(p-q)\left(\delta^T \beta V\right) \geq c \ .$$

Where the Agent's value is:

$$V = \frac{S_T}{1-\delta^T\left(1-\beta(1-p)^T\right)}$$

and S_T denotes the present value of the surplus in T periods:

$$S_T = \sum_{t=0}^{T-1} \delta^t s \quad .$$

The first term of the constraint $(ICA\ T)$ reflects the probability that the Agent's decision to put effort in the first period is actually relevant. Which is the case only if he believes he will fail in the rest of the periods. Then $(p-q)$ reflects the decrease in the probability of failure from exerting effort and finally $\delta^T \beta V$ plays the same role as Z did in the finite horizon case, here though it is endogenous an is determined in turn by the value of the relationship.

Using the fact that $(ICA\ T)$ binds in the optimal contract, we can re-write V as:

$$V = \frac{S_T - \frac{c(1-p)}{(p-q)}}{1-\delta^T} \quad .$$

Defining the value in this way highlights the similarity to a repeated T-period money burning problem. It also shows why implementing the longest possible review length is

[32] The proof, which is ommited, follows the same arguments as Part (iv) of Proposition 3.

optimal.[33] The review length will be bounded by the need to have $\beta \leq 1$. We can solve for the optimal β given T :

$$\beta^* = \frac{\left(1-\delta^T\right)}{\delta^T} \frac{c}{(p-q)(1-p)^{T-1}\left(S_T - \frac{c(1-p)}{(p-q)}\right)}$$

A necessary condition for the implementability of a contract of length T is that:

$$S_T > \frac{c(1-p)}{(p-q)} \tag{1}$$

This simply establishes that in period T there is something positive to be burnt. Of course this alone is not sufficient since the exact amount to be burnt depends as well on all the other parameters. In particular note that as players are more patient longer and longer review lengths can be implemented. Furthermore, as long as the per period surplus is positive ($s > 0$) if players are very patient, then an implementable contract always exists as shown in the following Lemma.

Lemma 4 *In any relationship with positive surplus $s > 0$ as we let $\delta \to 1$ there is some review length T large enough such that high effort is implementable.*

Definition 3 *The inefficiency $\lambda_T \in [0, 1]$ of an implementable contract, is defined by the value V it achieves, relative to the value of First Best:*

$$\lambda_T = 1 - \frac{V}{V^{FB}} \ .$$

Conditional on the contract being implementable, the inefficiency λ_T is given by:

$$\lambda_T = \frac{c(1-p)}{(p-q)} \frac{1}{S_T} \tag{2}$$

We can see from (2) that the same forces that were present in determining the inefficiency in the finite horizon play a role here. Additionally, the fact that we limit the amount of resources that can be burnt to the surplus in the relationship makes the inefficiency a function of the surplus that can be generated during the length of the review. As the per

[33] The $(ICAT)$ constraint allows us to conjecture that if we allowed for limited amounts money burning, both instruments would play the same role in the provision of incentives. Furthermore, that they would be used together to extend the review length as much as possible.

period surplus increases or as we are able to increase the review length, we are able to reduce the inefficiency. In particular, as $\delta \to 1$ we can make $\lambda_T \to 0$ since as $\delta \to 1$ we can take $T \to \infty$ and therefore $S_T \to \infty$. The Theorem below formalizes this result.

Theorem 2 (Folk Theorem) *For any $\varepsilon > 0$ and $s > 0$, there exists a $\delta^* < 1$ such that for all $\delta > \delta^*$ we can construct a $T(\varepsilon, s)$ period review contract which is implementable and has less than ε inefficiency.*

We can relate this result back to the characterization of the optimal contract with money burning. As $\delta \to 1$ the future value of the relationship grows and hence, the constraint that having the amount of burning endogenously determined by the future surplus of the relationship becomes weaker and weaker. Hence, as it was the case for the finite horizon with $T \to \infty$, we want to drive the review time as far out in the future as possible and only terminate after the worst history.

Although the T-period review contracts are very convenient to analyze and asymptotically efficient, for $\delta < 1$ they are generally not optimal. One exception is the special case $p = 1$, in this case even the one period review contract achieves first best. This result can be related to Shapiro and Stiglitz (1984). They use one-period review contracts in which the Agent gets all the surplus to derive a theory for efficiency wages and involuntary unemployment. They don't prove that this type of contract is optimal. Instead, they claim that other types of contracts, such as having bond posting by the Agent, could be hard to implement because "the firm could have an incentive to claim the worker shirked to claim the bond". In this paper we rigorously show that the one-period review contract is actually optimal over all contracts for the special case they analyze. As we showed in the previous section efficiency wage type contracts are generally opimal for $p < 1$ when the Agent's output is privately observed by the Principal.

IV Extensions

A Stochastic Effort, Multiple Levels of Effort and Communication

So far we have assumed that the only decisions available for the Agent were wether to work or to shirk a natural question to explore is what would happen if we allowed for random effort choices and also if we allowed for more degrees in the effort choice. We start by discussing the extension to multiple effort levels. Then, we show that with two levels of effort it is without loss of generality to focus on deterministic effort choices by the Agent. Finally, we argue that when we allow for both multiple levels of effort and mixed strategies by the Agent new types of efficiency enhancing contracts should be considered.

If we were to allow for many effort levels but restrict our attention to contracts in which the Agent follows a pure strategy the structure of the contracts would be very similar to the ones presented in this paper. There is one interesting feature of this generalization worth mentioning and which has already been discussed in MacLeod (2003), the effort level with private monitoring might be above the first best level of effort. This is because although it will be harder to motivate the agent to exert high effort, low output will be observed less frequently and therefore, value will need to be destroyed less often. Hence, depending on the specific parametrization of the problem it can be the case that it is beneficial to have the agent work harder than in the full information case. This same type of results carry on to the repeated case.[34]

In the contracts we have analyzed so far, the Principal has been doing all the monitoring. The continuation values are conditional only on his observations. We might think there is potential to grant some monitoring power to the Agent. If the Agent follows a mixed strategy for effort, he would have private information on what the outcome distribution should look like. This information can then be compared to the Principal's output reports. We could conjecture that this might lead to an improvement. Although a complete analysis of these issues is left for future work we want to provide two results of interest illustrating whether this is or not the case.

We first prove a negative result showing that within the framework of our model and focusing on one-period review contracts it is optimal to have the Agent exert high effort deterministically.

Proposition 6 *Allowing for contracts in which the Agent randomizes and the continuation values are determined as a function of the simultaneous announcements by the Agent and the Principal does not provide any improvement over the one-period review contract in which the Agent exerts high effort deterministically.*

Proof. See Appendix A. ■

[34] *Some more structure needs to be impossed on the cost function and the probability of success to keep the number of deviations to consider at bay.*

The intuition behind this negative result is that when the Agent mixes, the Principal actually has more incentives to claim L instead of H. This follows from the fact that the Principal cannot be punished for claiming low output if effort was low. Therefore, when the Principal observes H he knows there is some probability that the Agent had actually done low effort and hence that his lie will go unpunished. This intuition would seem to generalize when considering arbitrary contracts.

If instead the model allowed for an inefficiently high effort level, then things could work differently. Consider for simplicity the case in which if the Agent exerts a very high effort, then output is high for sure. Now, let the Agent mix between the efficient effort level and this higher effort level. The Principal is now more scared to declare L when H was realized. He knows that there is a chance that the Agent exerted the very high effort level and, that if this is the case, he would be caught lying if he claimed L. In Appendix A we provide an example that illustrates this point and shows how the Agent can use mixing to a higher effort level to monitor the Principal.

B Tournaments

The role of tournaments as an incentive device for the agents has been largely studied in the literature starting with the seminal paper of Lazear and Rosen (1981). The literature on tournaments has generally assumed that output is contractable. By changing that assumption, this paper highlights a virtue of tournaments that, with the exception of Malcomson (1984), has not received much attention in the past. The tournament structure allows for contracts that provide a constant continuation value for the Principal for all output realizations and upward-sloping schedules for the Agents without the need to burn any surplus. This reminds us somewhat of the moral hazard in teams problem studied by Holmstrom (1982) in which the agents use the Principal as a budget braker. In our setup, it is the Principal who uses a new agent to this effect.

Suppose the Principal can commit to give a prize b to the best performing Agent every period (or randomize in case of a draw). This way the Principal has no more constraints since he has to give a prize to some player but is really indifferent about which Agent he gives it to. The Agents have incentives exert high effort in order to win the prize.

For example for $T = 1$ with two Agents, to provide incentives for them to exert high effort we need:

$$\left(\frac{p-q}{2}\right) b - c \geq 0$$

The term multiplying the bonus is by assumption positive hence we can always find a bonus big enough to provide incentives for the Agents to exert high effort. Note that this might require the base wage to be negative depending on how much of the surplus is captured by the Principal and how much by the Agents.

Tournaments themselves are in turn susceptible to other problems. Examples of these shortcomings include: the Agents colluding against the Principal or the Principal colluding with one of the Agents against the other. The Agents might also engage in unproductive activities to try to win over the Principal's favor or directly to try to undermine the other Agent's work. Therefore, we would expect to observe tournaments in environments in which the information problem is of more importance than these other concerns. A full analysis of these issues is left for future work.

V Conclusions

Relaxing the assumption that output is common knowledge, we were able to explain many features of the labor markets that were at odds with the literature. In particular, we have provided a rationale for the observed wage compression and the use of efficiency wages together with the threat of termination to provide incentives.

This information structure can also be used to better understand questions of organizational design and potentially the theory of the firm. In particular we described how setting up tournaments can be used to simultaneously provide incentives for the Agents to exert effort and for the principal to be truthful.

Additionally our paper provides a useful benchmark in which there is no need for the use of communication between agents and principals. An important avenue for future research is to explore further possible reasons for the existence of communication. As shown in Section IV A a possible solution might come from allowing the Agent to monitor the principal by randomizing his effort choice to inefficiently high effort levels. Alternatively, the existence of match-specific productivity parameters might provide a rationale for communication. We hope to address these issues in future research.

VI Appendix A

A No Mixing No Talking

Consider the following one period review contract with random effort. The Agent plays $e = 1$ with probability $\alpha \in (0, 1)$, the Principal then observes the output and they simultaneously announce the actual effort and the actual output.

The following table summarizes the parameters (other than α) of the self-enforcing contract conditional on the announcements.

Announcements (e, y)	Wage	Termination Probability
$(0, L)$	w_{0L}	β_{0L}
$(0, H)$	w_{0H}	β_{0H}
$(1, L)$	w_{1L}	β_{1L}
$(1, H)$	w_{1H}	β_{1H}

Continuation Values (including wages not including cost of effort)

Principal	$y = L$	$y = H$
$e = 1$	F_{1L}	F_{1H}
$e = 0$	F_{0L}	F_{0H}

Agent	$y = L$	$y = H$
$e = 1$	V_{1L}	V_{1H}
$e = 0$	V_{0L}	V_{0H}

In order to simplify our analysis and exposition, we define the following new variables:

$\Delta_1 = \alpha\left(F_{1H} - F_{1L}\right)$

$\Delta_0 = (1-\alpha)\left(F_{0L} - F_{0H}\right)$

$\Delta_L = V_{1L} - V_{0L}$

$\Delta_H = V_{1H} - V_{0H}$

$\Delta_{cross} = V_{1H} - V_{0L}$

For the Principal we have the following truth-telling constraints:

Given that he observed $y = H$:

(PTTH) $$\Delta_0 \leq \frac{p}{q}\Delta_1$$

Given that he observed $y = L$:

(PTTL) $$\Delta_0 \geq \frac{(1-p)}{(1-q)}\Delta_1$$

We can represent these constraints graphically:

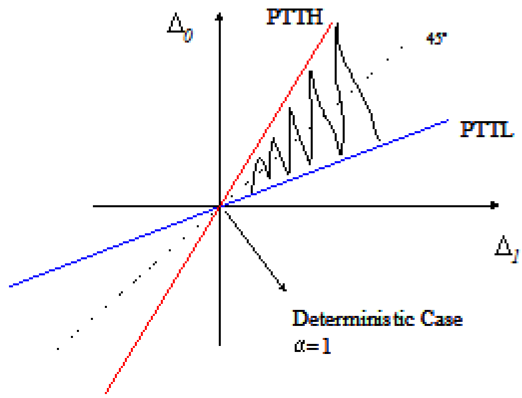

These constraints imply that $\Delta_0, \Delta_1 \geq 0$

For the Agent we have the following two truth-telling constraints:
Given he exerted high effort:

$(ATT1)$ $$\Delta_H \geq -\frac{(1-p)}{p}\Delta_L$$

Given he exerted low effort:

$(ATT0)$ $$\Delta_H \leq -\frac{q}{(1-q)}\Delta_L$$

We also need the Agent to be indifferent between high and low effort

(AIC) $$\Delta_H = \frac{c-(p-q)\,\Delta_{cross}}{q} - \frac{(1-p)}{q}\Delta_L$$

Graphically:

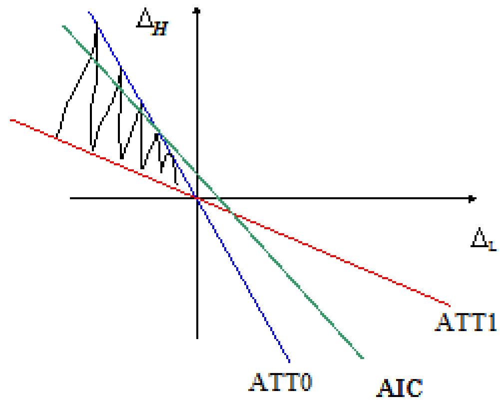

Note that (AIC) has always a slope between $(ATT0)$ and $(ATT1)$.

With these preliminaries, we now prove Proposition13 from the text.

Proposition 13 *Allowing for contracts in which the Agent randomizes and the continuation values are determined as a function of the simultaneous announcements by the Agent and the Principal does not provide any improvement over the one-period review contract in which the Agent exerts high effort deterministically.*

Proof. Suppose that the optimal contract has $\alpha < 1$ and $\Delta_1$and $\Delta_0 > 0$ and that $(PTTH)$ does not bind. Now consider the following: decrease w_{1L} and β_{1L} and increase w_{1H}. We can do this such a way that the resulting increase in V_{1H} is equal to $\frac{1-p}{p}$ times the decrease in V_{1L}.As a result, the Agent's IC constraints continue to hold and the truth-telling constraints are relaxed. The decrease in β_{1L} implies an improvement in efficiency since there will be less value burning. Therefore, $(PTTH)$ must be binding in the optimal contract.

Now suppose that$(PTTH)$ binds, but that $(PTTL)$ does not bind. If we could decrease.Δ_0 we would relax $(PTTH)$ and improve. This can be achieved by simply increasing α. Changing α has no impact on the Agent's constraints, it relaxes $(PTTH)$ and we can do it till $(PTTL)$ starts binding. If we iterate between this two steps (see graph below) clearly we can drive $\Delta_1$and Δ_0 all the way to zero. This means that either $\alpha = 1$ or that $(F_{0L} - F_{0H}) = (F_{1H} - F_{1L}) = 0$ but if this is the case we can set $\alpha = 1$ anyway and this is preferable since lower $\alpha's$ implies an efficiency loss. ■

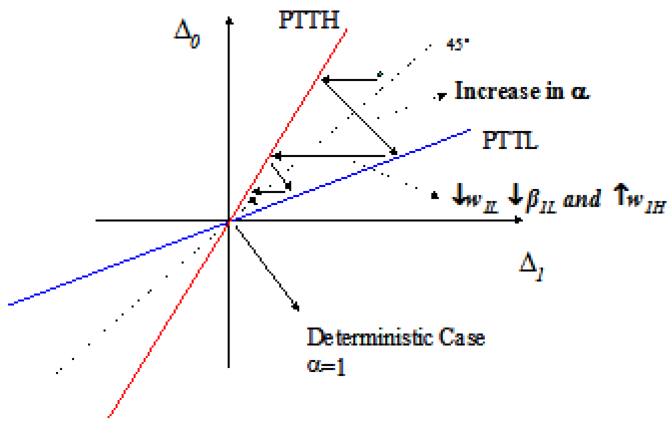

Mixing by the Agent fails as means to relax the Principal's truth-telling constraints when he observes a high realization of outcome. The intuition behind this negative result is that when the Agent mixes the Principal is actually getting some slack from claiming $y = L$ when $y = H$. This is because there is some chance that the Agent had actually done low effort and the Principal cannot be punished for claiming L when no effort effort was exerted. If instead we allowed the Agent to mix a higher effort level (even if it were inefficient from a productive perspective) then mixing might work. We illustrate this with an example in the next section.

B Monitoring with high Effort

Suppose that the Agent could additionally choose to exert a higher effort level. If this effort level leads to high outcome with higher probability, mixing by the Agent to this higher effort level relaxes the Principal's truth-telling constraints. The intuition for this is clearer when we consider the case in which very high effort leads to a high output for sure. Suppose the Agent mixes between these two effort levels, when the Principal observes high output he is now less tempted to claim low. He fears that the Agent might have exerted very high effort and that he would be caught lying if he did.

To illustrate more precisely how this works consider the following example: let $e = \{l, m, h\}$ and $c(e) = \{0, c, 1\}$ with $p(e) = \{0, p, 1\}$ respectively.[35] $H = 1$ and $L = 0$. Medium effort is the only productively efficient effort level $p - c > 0$.

Suppose we wish to have the Agent mix between high and medium effort and we allow

[35] Similar examples could be constructed for less extreme cases, we pick these parameters for expositional convenience.

continuation values including wages to depend on the simultaneous announcements. Denote these by $V_{\tilde{e},\, \tilde{y}}$ for the Agent and $F_{\tilde{e},\, \tilde{y}}$ for the Principal. We face the following constraints:

For the Agent:

Thruth-telling:

Given he exerted m effort:

$$pV_{mH} + (1-p)\, V_{mL} \geq pV_{hH} + (1-p)\, V_{hL}$$

Given he exerted h effort:

$$V_{hH} \geq V_{mH}$$

Incentive compatibility for mixing:

$$pV_{mH} + (1-p)\, V_{mL} - c = V_{hH} - 1 \geq V_{mL}$$

For the Principal:

Thruth-telling:

Given he observed L:

$$(TTL) \qquad F_{mL} \geq F_{mH}$$

Given he observed H:

$$(TTH) \qquad \alpha F_{hH} + (1-\alpha)\, pF_{mH} \geq \alpha F_{hL} + (1-\alpha)\, pF_{mL}$$

Consider a contract in which the Principal gets all the residual surplus.

Announcements	Wage	Termination Probability
(m, L)	$w_{mL} = 0$	$\beta_{mL} = 0$
(m, H)	$w_{mH} = \frac{c}{p}$	$\beta_{mH} = 0$
(h, L)	$w_{hL} = -\frac{p}{1-p}$	$\beta_{hL} = 1$
(h, H)	$w_{hH} = 1$	$\beta_{hH} = 0$

The Agent has no incentives to deviate since when he is called to exert high effort he is compensated for sure with the cost he had to bare. When he is called to do medium effort he gets zero if he obeys. If he chooses to exert low effort the best he can get is also a payoff of zero.

We must now make sure that the Principal does not have incentives to claim L when he observes H.[36] Using the table above we can re-write (TTH) as follows:

$$\begin{aligned} & \alpha\left(-1+\frac{\delta\left(\left(1-\alpha\right)\left(p-c\right)\right)}{1-\delta}\right)+\left(1-\alpha\right)p\left(-\frac{c}{p}+\frac{\delta\left(\left(1-\alpha\right)\left(p-c\right)\right)}{1-\delta}\right) \\ \geq\ & \alpha\frac{p}{1-p}+\left(1-\alpha\right)p\left(\frac{\delta\left(\left(1-\alpha\right)\left(p-c\right)\right)}{1-\delta}\right) \end{aligned}$$

We will use the following parameter values to illustrate the example: $p=\frac{1}{2},\delta=.9$ and $c=\frac{1}{8}$. Using these particular values and doing some algebra, we see that satisfying (TTH) reduces to finding an α that satisfies:

$$-\frac{27}{4}\alpha^2+3\alpha-\frac{1}{4}\geq 0$$

Any $\alpha\in\left[\frac{1}{9},\frac{1}{3}\right]$ is a solution to this, but since m is more productive than h, it is optimal to pick $\alpha=\frac{1}{9}$. In this case we get a total expected surplus of $3\frac{1}{3}$.The First Best surplus is $\frac{15}{4}=3.75$ and with the deterministic one-period review contract using our results from Section 5.1.1 we get an expected surplus of only $\frac{2}{3}FB=2.5$.This example shows that when considering a more general model, with many possible effort choices, ruling out mixed strategies by the Agent and communication might reduce efficiency.

VII Appendix B

The following definitions facilitate our analysis.

Definition 4 (Continuation Values) *$V_t^+\left(\left\{h_{t-1}^P,y_t\right\},\sigma\right)$ is the promised continuation value to the Agent in period t after e_t has been exerted and the Principal has observed y_t. Similarly, $F_t^+\left(\left\{h_{t-1}^P,y_t\right\},\sigma\right)$ is the principal's continuation value after observing and consuming y_t. Formally:*

$$V_t^+\left(\left\{h_{t-1}^P,y_t\right\},\sigma\right)=\mathbb{E}\left[w_t+b_t+\sum_{j=t+1}^{T}\delta^{(j-t)}(w_j+b_j-c\left(e_j\right))\middle|\left\{h_{t-1}^P,y_t\right\},\sigma\right],$$

[36] It is easily verified that (TTL) is satisfied.

$$F_t^+\left(\left\{h_{t-1}^P, y_t\right\}, \sigma\right) = \mathbb{E}\left[-y_t + \sum_{j=t}^{T} \delta^{(j-t)}(y_j - w_j - b_j) | \left\{h_{t-1}^P, y_t\right\}, \sigma\right] .$$

The Agent can only form an expectation of this value:

$$v_t^+ = \mathbb{E}\left[V_t^+\left(\left\{h_{t-1}^P, y_t\right\}, \sigma\right) | \left\{h_{t-1}^A, e_t\right\}, \sigma\right] .$$

Definition 5 *Given σ_ω and a history h_{t-1}^V such that $e\left(h_{t-1}^V, \sigma_\omega\right) = 1$, we say that σ_ω^A is incentive compatible with respect to effort for the Agent (ICA) after history h_{t-1}^V iff:*

$$\begin{pmatrix} \mathbb{E}\left[V_t^+\left(\left\{h_{t-1}^P, y_t\right\}, \sigma_\omega\right) | \left\{h_{t-1}^A, e_t = 1\right\}, \sigma_\omega^A, \sigma_\omega^P\right] - \\ \mathbb{E}\left[V_t^+\left(\left\{h_{t-1}^P, y_t\right\}, \sigma_\omega\right) | \left\{h_{t-1}^A, e_t = 0\right\}, \tilde{\sigma}_\omega^A, \sigma_\omega^P\right] \end{pmatrix} \geq c ,$$

for all $\tilde{\sigma}_\omega^A$ such that $\tilde{\sigma}_\omega^A \equiv \sigma_\omega^A$ for all $\tau < t$ and $\tilde{e}\left(h_{t-1}^V\right) = 0$.[37]

Definition 6 *Given σ^A and ω we say σ_ω^P is incentive compatible for the Principal (ICP) after history $\left\{h_{t-1}^P, y_t\right\}$ iff:*

$$F_t^+\left(\left\{h_{t-1}^P, y_t\right\}, \sigma\right) \geq F_t^+\left(\left\{h_{t-1}^P, y_t\right\}, \sigma^A, \tilde{\sigma}^P\right) ,$$

for all $\tilde{\sigma}_\omega^P$ such that $\tilde{\sigma}_\omega^P \equiv \sigma_\omega^P$ for all $\tau < t$.[38]

Proof of Proposition 1. Follows directly from *Lemmas* (5) and (6) stated below. ■

Lemma 5 *The ICA implies:*

(3)
$$\begin{aligned} & \mathbb{E}\left[V_t^+\left(\left\{h_{t-1}^P, y_t = H\right\}, \sigma\right) | \left\{h_{t-1}^A, e_t = 1\right\}, \sigma\right] \\ > \; & \mathbb{E}\left[V_t^+\left(\left\{h_{t-1}^P, y_t = L\right\}, \sigma\right) | \left\{h_{t-1}^A, e_t = 1\right\}, \sigma\right] , \end{aligned}$$

i.e. the expected payoff to the Agent is strictly increasing in the output when the Agent is supposed to exert effort.

[37] Note that $\tilde{\sigma}_\omega^A$ can include arbitrary future deviations both in the effort and in the termination, but agrees with equilibrium until $t-1$.

[38] In words, $\tilde{\sigma}_\omega^P$ is consistent with σ_ω^P up to $t-1$ but can include arbitrary future deviations both in the messages and in the termination.

Proof. Follows from the fact that effort increases the likelihood of $y_t = H$. ■

Lemma 6 *The IC for the Principal implies:*

$$F_t^+ \left(\left\{ h_{t-1}^P, y_t = H \right\}, \sigma \right) = F_t^+ \left(\left\{ h_{t-1}^P, y_t = L \right\}, \sigma \right) \quad \forall h_{t-1}^P \tag{4}$$

Proof. Suppose first that for some h_{t-1}^P :

$$F_t^+ \left(\left\{ h_{t-1}^P, y_t = H \right\}, \sigma \right) > F_t^+ \left(\left\{ h_{t-1}^P, y_t = L \right\}, \sigma \right) ,$$

and that $y_t = L$ is realized.

Now, consider the alternative strategy $\tilde{\sigma}_\omega^P$ for the Principal. $\tilde{\sigma}_\omega^P$ follows σ_ω^P except that it implies the same actions as σ_ω^P for histories with $y_t = L$ as it does for histories where $y_t = H$. This deviation is undetectable by the Agent and leads to an improvement for the Principal.[39]

A similar argument can be constructed if:

$$F_t^+ \left(\left\{ h_{t-1}^P, y_t = H \right\}, \sigma \right) < F_t^+ \left(\left\{ h_{t-1}^P, y_t = L \right\}, \sigma \right) .$$

Therefore, it must be the case that

$$F_t^+ \left(\left\{ h_{t-1}^P, y_t = H \right\}, \sigma \right) = F_t^+ \left(\left\{ h_{t-1}^P, y_t = L \right\}, \sigma \right) ,$$

for all h_{t-1}^P. ■

Lemmas 5 and 6 imply that for any strategy pair σ^* and history h_{t-1}^V such that σ^* induces $e\left(h_{t-1}^V\right) = 1$, the Agent's continuation value v_t^+ must depend on the outcome but the Principal's continuation value F_t^+ must be equal after either outcome. This implies that the total continuation surplus depends on the outcome realization.

Lemma 7 *Given σ^A and ω if for all $\left\{h_{t-1}^P, y_t\right\}$ σ_ω^P is incentive compatible for the Principal then σ_ω^P is a best response to σ^A given ω.*

[39] For the case when $p = 1$ and $q > 0$. We have to be careful when $e_t = 1$ because now the Agent can detect a deviation if the Principal claims $y_t = L$.

Proof. Follows directly from the definition of best response. There is no $\tilde{\sigma}_{\omega}^{P} \in \Sigma_{\omega}^{P}$ that can achieve a higher value for the Principal. ■

Technically, we are abusing the definition of best response slightly because we are not explicitly considering detectable deviations. On the other hand these deviations could be dealt with in a simple way.

Proof of Proposition 2. When $T = 1$, since this is the last period of the game, the continuation values for both, the Agent and the Principal, solely consists on the current compensation. By Lemma 6 the total compensation the Principal pays must be the same for either realization of output. Hence, providing incentives for the Agent to exert effort, which requires compensation to be dependent on the output, is not possible.

Next consider any arbitrary finite horizon. Suppose there exists a contract that contemplates the Agent exerting effort in some periods. Consider the last period in which the Agent is supposed to exert effort. By the arguments given to prove the case $T = 1$ it follows that incentives cannot be provided to exert effort for this last period hence contradicting the claim that there can be a contract for finite T in which effort is exerted. ■

Proof of Lemma 1. Suppose that after some history $h_{t-1}^{V} \times m^{t} \times [0, 1]$ a payment w_t^p was to be made by the Principal and $w_t^A \leq w_t^P$ to be received by the Agent. Instead of paying immediately the Principal can send a message that commits him to pay $\frac{w_t^p}{\delta^{T-t}}$ at time T. Also, any money burning that was supposed to take place at time t can also be delayed until time T. ■

Proof of Lemma 2.

Follows from Lemma (10) and noting that the destruction of value that can be achieved by termination can also be achieved by burning money. ■

Proof of Proposition 3. The proof to the first part of the Proposition is organized in four steps. First we write the problem, then we write a relaxed problem where only one-step deviations by the Agent are considered. The third step solves the relaxed problem and the last step proves that this is also a solution to the original problem.

i) To motivate effort for $T + 1$ periods which we denote $\mathbf{e}^{*}$, we face the following money burning minimization problem:

$$\min_{\{Z(y^T)\}y^T\in Y^T} \delta^T \sum_{y^t\in Y^T} Z\left(y^T\right) P\left(y^T|\mathbf{e}^*\right)$$

$$(5) \qquad s.t. \quad \mathbf{c}\left(\mathbf{e}^* - \mathbf{e}\right) + \delta^T \sum_{y^T\in Y^T} Z\left(y^T\right)\left(P\left(y^T|\mathbf{e}^*\right) - P\left(y^T|\mathbf{e}\right)\right) \leq 0 \;\; \forall \mathbf{e} \in \{0,1\}^{T+1}$$

$$Z\left(y^T\right) \geq 0 \quad \forall y^T \; .$$

Note that the Principal's incentive constraints are satisfied by construction since w^P is fixed.

ii) Instead of solving this problem directly, we first solve a relaxed problem in which we only consider one-step deviations by the Agent:

$$\min_{\{Z(y^T)\}y^T\in Y^T} \delta^T \sum_{y^t\in Y^T} Z\left(y^T\right) P\left(y^T|\mathbf{e}^*\right)$$

$$(6) \qquad s.t. \quad \delta^t c + \delta^T \sum_{y^T\in Y^T} Z\left(y^T\right) P\left(y^T|\mathbf{e}^*\right)\left(\frac{P\left(y_t|e_t=1\right) - P\left(y_t|e_t=0\right)}{P\left(y_t|e_t=1\right)}\right) \leq 0 \;\; \forall t$$

$$Z\left(y^T\right) \geq 0 \quad \forall y^T \; .$$

From (6) it is clear that in order to provide incentives in period t, we need:

$$\left(\frac{P\left(y_t|e_t=1\right) - P\left(y_t|e_t=0\right)}{P\left(y_t|e_t=1\right)}\right) < 0 \; .$$

In words, to provide incentives for effort in period t we need to have burning if the outcome of that period is low which, is more likely to occur if the Agent deviates. Therefore, the constraints can be written as:

$$\delta^t c \frac{(1-p)}{p-q} \leq \delta^T \sum_{y^T\in Y^T\cap(y_t=L)} Z\left(y^T\right) P\left(y^T|\mathbf{e}^*\right) \quad \forall t$$

Clearly $t = 0$ is the most binding constraint if $\delta < 1$ which puts a lower bound on the necessary expected money burning $c\frac{(1-p)}{(p-q)}$.

iii) Next we show that:

$$Z^* = \left\{ \begin{array}{ll} \frac{c}{\delta^T (p-q)(1-p)^T} & if \;\; y^T = \mathbf{L} \\ 0 & otherwise \end{array} \right\}$$

achieves the lower bound and satisfies all the IC constraints of the relaxed problem.

First note that:

$$P\left(y^T=\mathbf{L}|\mathbf{e}^*\right)=(1-p)^{T+1}.$$

Hence,

$$\delta^T\sum_{y^T\in Y^T}Z^*\left(y^T\right)P\left(y^T|\mathbf{e}^*\right)=c\frac{(1-p)}{(p-q)}\geq\delta^t c\frac{(1-p)}{(p-q)}\quad\forall t\,.$$

iv) We now show that Z^* is also a solution to the original problem.

Using Z^*, the constraints of the original problem (5) take the following form:

$$\mathbf{c}.\left(\mathbf{e}^*-\mathbf{e}\right)+\frac{c}{(p-q)(1-p)^T}\left((1-p)^{T+1}-P\left(y^T|\mathbf{e}\right)\right)\leq 0\ \ \forall\mathbf{e}\in\{0,1\}^{T+1}.$$

Letting n denote the number of periods the Agent deviates we can write the term that captures the expected cost of deviation as follows:

$$\frac{c}{(p-q)(1-p)^T}\left((1-p)^{T+1}-(1-p)^{(T+1)-n}(1-q)^n\right).$$

Now suppose the Agent was deviating in $n-1$ periods and now deviates for n periods. The increase in the probability that $y^T=\mathbf{L}$ is:

$$\left(\frac{1-q}{1-p}\right)^{n-1}(1-p)^T(p-q)\,.$$

Since $\left(\frac{1-q}{1-p}\right)>1$ the increase in the probability of punishment from an additional deviation is increasing in the number of previous deviations. Hence, if the Agent doesn't find it profitable to deviate once, he won't want to deviate at all. Given $\delta<1$ and Z^*, the most profitable single deviation is in the first period. This is not profitable by step (iii).

Finally, the conditions for w^P guarantee that individual rationality constraints are satisfied so the players sign the contract ω at time zero.

Uniqueness proof for the case when $\delta=1$:

Consider any other incentive compatible contract Z. Clearly, for Z to be optimal it must be that $Z\left(y^T=\mathbf{L}\right)<Z^*\left(y^T=\mathbf{L}\right)$ otherwise Z^* would imply strictly less expected money burnt. Denote all histories $y^T\neq\mathbf{L}$ for which $Z\left(y^T\right)>0$ and $y_0=L$ by $\dot{y}^T$. If there is no such history, then $Z\left(y^T=\mathbf{L}\right)\geq Z^*\left(y^T=\mathbf{L}\right)$ otherwise, the Agent would deviate in the first period. Hence, if Z is to be optimal it must be that there is some $\dot{y}^T$. Suppose there is one such $\dot{y}^T$. Let $j = min\ \{t : \dot{y}_t = H\}$. The incentive constraint for period j effort is given by:

$$\delta^j c + \delta^T \left(Z\left(y^T = \mathbf{L}\right)(1-p)^{T+1}\left(\frac{q-p}{1-p}\right) + Z\left(\dot{y}^T\right) P\left(\dot{y}^T\right)\left(\frac{p-q}{p}\right)\right) \leq 0$$

Now, let $\delta \to 1$

$$c + Z\left(y^T = \mathbf{L}\right)(1-p)^{T+1}\left(\frac{q-p}{1-p}\right) + Z\left(\dot{y}^T\right) P\left(\dot{y}^T\right)\left(\frac{p-q}{p}\right) \leq 0$$

This is the IC for period 0 effort. Note this implies that Z cannot be incentive compatible since this constraint held with equality for Z^* and $Z\left(y^T = \mathbf{L}\right) < Z^*\left(y^T = \mathbf{L}\right)$ and additionally, the last term in the left hand side of the inequality is positive. Therefore, Z^* is the only contract that is optimal for all δ. ■

Proof of Proposition 4 .

All of the Agent's IC constraints are holding with equality when he is completely uninformed. As soon as he can condition on any information, the probability he assigns to his effort being pivotal in any given period would change. Some IC constraints would therefore be violated when he updates this probability given the messages. ■

Restatement of Theorem 1

Given any contract (ω, σ_ω) *that generates values* V *and* F, *we can construct a payoff equivalent contract* $(\hat{\omega}, \hat{\sigma}_\omega)$ *with the following properties:*

(i) The Agent receives a constant wage until he is fired and no bonuses.

(ii) The Agent exerts effort every period until he is fired.

(iii) The Principal gives no feedback (sends no messages) to the Agent.

Proof. Properties (i) and (ii) follow from applying successively Lemmas (8), (9) and (10). Given properties (i) and (ii) Lemma (11) delivers property (iii). ■

Lemma 8 *For any contract* $[\omega, \sigma^*]$ *that generates values* V *and* F *there is a payoff-equivalent contract* $[\tilde{\omega}, \tilde{\sigma}^*]$ *such that for any history* h^V_{t-1} *with the property that* $e\left(h^V_{t-1}\right) = 0$ *all future actions are independent of the outcome* y_t.

Proof. Suppose the Principal conditioned his action on y_t. Now, instead let him condition his strategy $\tilde{\sigma}^P$ on ϕ_t in the following way: if $\phi_t < q$ let the Principal take the same actions that he did for σ^P when $y_t = H$ and if $\phi_t > q$ take the same actions he did for σ^P when $y_t = L$. The Principal had to be indifferent between the outcomes of y_t by Lemma 6 hence he will not have incentives to deviate from $\tilde{\sigma}$ if he didn't have incentives to deviate from σ. If the Agent follows the equilibrium strategy, his payoffs and information are the same as in the original equilibrium. Therefore, if he follows $e\left(h_{t-1}^V\right) = 0$ the incentives to deviate in later or earlier periods are unchanged. If he deviates in the current period to $e\left(h_{t-1}^V\right) = 1$ he is not affecting future payoffs nor obtaining any information hence such deviation costs him c without any benefit. ■

In words, to induce no effort we don't need to provide any incentives to the Agent. Therefore, we can always make his payoff unconditional of the outcome without inducing him to put the effort.

Definition 7 *A contract* $[\omega, \sigma^*]$ *is an input based compensation contract if:*

$$b_t + w_t = E\left[y_t | e_t\right] \;\; for \;\; all \;\; t > 0.$$

Lemma 9 (Input based pay) *For every contract* $[\omega, \sigma^*]$ *that generates values* V *and* F *there exists a payoff equivalent input based compensation contract* $[\tilde{\omega}, \tilde{\sigma}^*]$.

Proof. Modify $[\omega, \sigma^*]$ in the following way: let $\tilde{b}_t + \tilde{w}_t = E\left[y_t | e_t\right] \;\; for \;\; all \;\; t > 0$ and $\tilde{w}_0 = E\left[y_0 | e_0\right] - F$. Also, let $\tilde{a}_t = a_t k_t$ and $\tilde{k}_t = 1 \; \forall t$. This changes imply that the Principal gets F as in the initial contract but also, that for all t the Principal's continuation value is zero (equal to his outside option). Therefore, there are no incentive issues having the Principal, who is indifferent on whether to terminate or not, do all the termination. Additionally, there are no incentive problems regarding the principal's announcements. By construction, his announcements might change effort level and compensation, but they do not change his expected surplus. The Agent won't want to deviate with respect to $\{\tilde{k}\}$ since terminating when he is not supposed to, would make him worse off as he gets all the surplus in the continuation game. Lemma 8 implies we only need to check that there will be no deviations with respect to effort for those histories where effort is supposed to be exerted. Note that under the original contract, when the Agent was supposed to exert effort there had to be value burning in case of a bad outcome to ensure that the right incentives were provided for the Agent to exert effort. We showed already that the new contract will destroy the same amount of value after the same histories than the original contract, hence the Agent will have the same incentives to exert effort under the new contract as he had under the old. ■

Definition 8 *A contract $[\omega, \sigma^*]$ is a termination contract if it is an input based compensation contract and the Agent exerts effort every period until termination.*

Lemma 10 (Termination) *For any input based compensation contract $[\omega, \sigma^*]$ that generates values V and F there is a payoff equivalent termination contract $[\tilde{\omega}, \tilde{\sigma}^*]$.*

Proof. First we will use *Lemma* 8 to transform $[\omega, \sigma^*]$ into $[\hat{\omega}, \hat{\sigma}^*]$. Note that this implies that the destruction of value that is produced by having the Agent exert no effort after history h_{t-1}^V for a period is independent of y_t. When there is a period of no effort, value is being destroyed because there is a delay of one period until any future surplus is realized. Therefore if we let $\tilde{a}\left(h_{t-2}^P, y_t, x_t\right) = \hat{a}\left(h_{t-2}^P, y_t, x_t\right)\delta$ we destroy in expectation the same amount of value. Finally, we must adjust all strategies by one period since we are eliminating the slack period. ■

This shows that incentives for the Agent to exert effort can be provided via efficiency wages and the threat of termination.

Definition 9 *A termination contract $[\omega, \sigma^*]$ is a no communication contract iff the messages m_t are completely uninformative $\forall\ t$.*

Lemma 11 (No communication) *For any termination contract $[\omega, \sigma^*]$ there is a payoff equivalent no communication contract $[\tilde{\omega}, \tilde{\sigma}^*]$.*

Proof. First note that in a termination contract on the equilibrium path the Agent's actions are independent of the messages m_t. Having the Agent completely uninformed does not give the Principal any profitable deviation possibilities so he will still follow the original termination rule $\{a\}$. Finally, if for every message m_t that the Agent could have received, he chose to exert effort that means that if now he has no message on which to condition his action he would still choose to exert effort. ■

Restatement of Proposition 5 *A sequence of termination rules $\{a\left(h^t\right)\}_{t=0}^{\infty}$ that does not have each of the following properties, can be weakly improved upon.*

(i) For all h^t, if $a\left(h^t, H\right) < 1$ then $a\left(h^t, L\right) = 0$.
(ii) For all $h^t \times h^j$, if $a\left(h^t, H, h^j\right) < 1$ then $a\left(h^t, L, h^j\right) = 0$.
(iii) For all t, if $a\left(L^t\right) > 0$ then $a\left(h^t\right) = 1$ for all h^t.
(iv) For all h^t, $V\left(h^t, H\right) \geq V\left(h^t, L\right)$.

(v) For all h^t, if $a(h^t, L) = a(h^t, H) = 0$ then $a(h^t) = 0$.

Proof. (i) Suppose $a(h^t, H) < 1$ and $a(h^t, L) > 0$ for some h^t. Let $\tilde{a}(h^t, L) = a(h^t, L) - \varepsilon \geq 0$ for some small $\varepsilon > 0$. Also, instead of terminating with probability $1 - a(h^t, H)$, after history (h^t, H) .Let the Principal terminate with probability $1 - a(h^t, H) - \left(\frac{1-p}{p}\right)\varepsilon \geq 0$ and with probability $\left(\frac{1-p}{p}\right)\varepsilon$ switch for history (h^t, L) upon observing the high outcome at $t+1$. That is, after history (h^t, H) start the Agent at $t+2$ with history (h^t, L) with probability $\left(\frac{1-p}{p}\right)\varepsilon$. We must make sure these changes respect two constrains: first the expected values at time t must not be changed. This guarantees that incentives for all periods prior to $t+1$ are unaffected. Formally:

$$\begin{aligned} & p \times a\left(h^t, H\right) \times V\left(h^t, H\right) + (1-p) \times a\left(h^t, L\right) \times V\left(h^t, L\right) \\ = \; & p\left(a\left(h^t, H\right) \times V\left(h^t, H\right) + \left(\frac{1-p}{p}\right)\varepsilon V\left(h^t, L\right)\right) + (1-p) \times \left(a\left(h^t, L\right) - \varepsilon\right) \times V\left(h^t, L\right) \end{aligned}$$

Second, the probability distribution on the tree is unaffected so all incentive constraints for periods after $t+1$ are unaffected.

$$(1-p) \times \left(a\left(h^t, L\right) - \varepsilon\right) + p\left(\frac{1-p}{p}\right)\varepsilon = (1-p) \times a\left(h^t, L\right)$$

The key is that incentive constraints at period $t+1$ have been relaxed. Since the change has increased $V(h^t, H)$ and decreased $V(h^t, L)$ the Agent now has more incentives to exert effort at $t+1$.

(ii) Suppose $\exists\ h^t, h^j \quad s.t.\ a(h^t, H, h^j) < 1$ and $a(h^t, L, h^j) > 0$ and consider the following changes to $\{a(h^t)\}$.Reduce the termination probability after history (h^t, H, h^j) to $1 - a(h^t, H, h^j) - \varepsilon$ where $\varepsilon > 0$. With probability ε instead of terminating assume the history is (h^t, L, h^j) and treat the Agent accordingly. Let $\tilde{a}(h^t, L, h^j)$ be the new continuation probability after (h^t, L, h^j) .Before the changes, the probability an Agent was hired to begin work in period $t+j+2$ with history (h^t, L, h^j) was:

$$P\left(h^t, L, h^j\right)\left(\prod_{\varnothing \leq h^u \leq (h^t, L, h^j)} a\left(h^u\right)\right)$$

Now that probability is:

$$\begin{aligned}
& P\left(h^t, L, h^j\right)\left(\prod_{\varnothing \leq h^u<\left(h^t, L, h^j\right)} a\left(h^u\right) \times \tilde{a}\left(h^t, L, h^j\right)\right) \\
& +P\left(h^t, H, h^j\right)\left(\prod_{\varnothing \leq h^u<\left(h^t, H, h^j\right)} a\left(h^u\right)\right) \times \varepsilon \times \tilde{a}\left(h^t, L, h^j\right) \\
= & \tilde{a}\left(h^t, L, h^j\right) P\left(h^t, L, h^j\right)\left(\prod_{\varnothing \leq h^u<\left(h^t, L, h^j\right)} a\left(h^u\right)+\varepsilon \frac{p}{1-p}\left(\prod_{\varnothing \leq h^u<\left(h^t, H, h^j\right)} a\left(h^u\right)\right)\right)
\end{aligned}$$

This probability must be kept constant to guarantee incentive constraints in the future hold hence:

$$\tilde{a}\left(h^t, L, h^j\right)=a\left(h^t, L, h^j\right) \frac{\left(\prod_{\varnothing \leq h^u<\left(h^t, L, h^j\right)} a\left(h^u\right)\right)}{\left(\prod_{\varnothing \leq h^u<\left(h^t, L, h^j\right)} a\left(h^u\right)+\varepsilon \frac{p}{1-p}\left(\prod_{\varnothing \leq h^u<\left(h^t, H, h^j\right)} a\left(h^u\right)\right)\right)}$$

This choice of $\tilde{a}\left(h^t, L, h^j\right)$ also leaves the continuation value $V\left(h^t\right)$ constant assuring that IC constraints for $\tau \leq t$ are unaffected. Incentive constraints between $t+2$ and $t+j+1$ are not affected either since the changes realized are only contingent on period $t+1$ action.

Finally, $V\left(h^t, H\right)$ has been increased and $V\left(h^t, L\right)$ decreased as a result of these changes so incentives in period $t+1$ have been relaxed. This implies at least a weak improvement can be achieved.

(iii)Follows from the proof of (ii)

(iv) Since $\forall h^t \times h^j \; a\left(h^t, H, h^j\right) \geq a\left(h^t, L, h^j\right)$ then since:

$$V\left(h^t\right)=a\left(h^t\right) \sum_{j=1}^{\infty} s\left(\delta^j\left(\sum_{h^j} P\left(h^t, h^j\right) \prod_{h^\tau \in\left(h^t, h^j\right)} a\left(h^\tau\right)\right)\right)$$

it follows that $V\left(h^t, H\right) \geq V\left(h^t, L\right)$.

(v) Suppose $a\left(h^t, L\right)=a\left(h^t, H\right)=0$ and $a\left(h^t\right)>0$. Let $\tilde{a}\left(h^t\right)=a\left(h^t\right)-\varepsilon$ for some small $\varepsilon>0$. Also with probability $\gamma\left(\tilde{h}^t\right)$ after h^t let the Principal pretend that all other histories $\tilde{h}^t$ occurred where the relative weights are given according to the likelihood of each alternative history $\tilde{h}^t$ occurring on the equilibrium path. γ must also be chosen so that the value after

h^t is kept constant. This change has kept values constant so past IC are still satisfied. The only change that we have made is to increase the probability that when deciding effort at $t+1$ the Agent is somewhere in the outcome tree where the current outcome realization affects his continuation value. Therefore, the incentive constraint for effort at $t+1$ has been relaxed. Incentives in all future periods are unaffected since the Agent's beliefs are the same that under the original contract. ■

Proof of Lemma 4. It follows from noting that as $\delta \to 1$ and rearranging (1) becomes $s > \frac{c(1-p)}{(p-q)} \frac{1}{T}$. Since the left hand side is a positive number we can always find a large enough T so that the condition is satisfied. Note as well that for any fixed T as $\delta \to 1$, $\beta \to 0$ so there is a feasible β^* to implement the contract. ■

Proof of Theorem 2 (Folk Theorem). The inefficiency λ is given by:

$$\lambda_T = \frac{c(1-p)}{(p-q)} \frac{1}{S_T}$$

Taking the limit as $\delta \to 1$

$$\lim_{\delta \to 1} \lambda_T = \frac{c(1-p)}{(p-q)\, T(\varepsilon, s)\, s}$$

Now let $T(\varepsilon, s)$ be greater than $\bar{T}$, where:

$$\bar{T} = \frac{c(1-p)}{\varepsilon (p-q)\, s}$$

Therefore the inefficiency is:

$$\lambda_T = \frac{c(1-p)}{T(\varepsilon, s)(p-q)\, s} < \frac{c(1-p)}{\frac{c(1-p)}{\varepsilon(p-q)s}(p-q)\, s} = \varepsilon$$

■

References

[1] Abreu, Dilip; Milgrom, Paul and Pearce, David. “Information and Timing in Repeated Partnerships.” *Econometrica*, 1991, 59(6), pp.1713-33.

[2] Abreu, Dilip; Pearce, David and Stacchetti, Ennio. “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring.” *Econometrica*, 1990, 58(5), pp. 1041-63.

[3] Amarante, Massimiliano. “Recursive Structure and Equilibria in Games with Private Monitoring.” *Economic Theory*, 2003, 22(2), pp. 353–74.

[4] Aoyagi, Masaki. “Collusion in Dynamic Bertrand Oligopoly with Correlated Private Signals and Communication.” *Journal of Economic Theory*, 2002, 102 (1), pp. 229–248.

[5] Baker, George; Gibbons, Robert and Murphy, Kevin J. “Subjective Performance Measures in Optimal Incentive Contracts.” *Quarterly Journal of Economics*, 1994, 109(4), pp. 1125-56.

[6] Baker, George; Gibbons, Robert and Murphy, Kevin J. “Relational Contracts and the Theory of the Firm.” *Quarterly Journal of Economics*, 2002, 117(1), pp. 39-84.

[7] Baker, George P. “Incentive Contracts and Performance Measurement.” *Journal of Political Economy*, 1992, 100(3), pp. 598-614.

[8] Bhaskar, V. and van Damme, Eric. “Moral Hazard and Private Monitoring.” *Journal of Economic Theory*, 2002, 102(1), pp. 16-39.

[9] Bewley, Truman F. *Why wages don't fall during a recession*. Cambridge, MA: Harvard University Press, 1999.

[10] Brown, Martin; Falk, Armin and Fehr, Ernst. “Contractual Incompleteness and the Nature of Market Interactions.” CEPR, Discussion Papers: No. 3272, 2002.

[11] Cappelli, Peter and Chauvin, Keith. "An Interplant Test of the Efficiency Wage Hypothesis." *The Quaterly Journal of Economics*, Vol. 106, No. 3. 1991,pp. 769-787.

[12] Che, Yeon-Koo and Yoo, Seung-Weon. “Optimal Incentives for Teams.” *American Economic Review*, 2001, 91(3), pp. 525-41.

[13] Compte, Oliver. “Communication in Repeated Games with Imperfect Private Monitoring.” *Econometrica*, 1998, 66(3), pp. 597-626.

[14] Compte, Oliver. "On Failing to Cooperate when Monitoring is Private." *Journal of Economic Theory*, 2002, 102(1), pp.151-88.

[15] Compte, Oliver. "On Sustaining Cooperation without Public Observations." *Journal of Economic Theory*, 2002, 102 (1), pp. 106-50.

[16] Doornik, Katherine. "Relational Contracting in Partnerships." *Journal of Economics and Management Strategy* (forthcoming).

[17] Fehr, Ernst and Falk, Armin. "Psychological Foundations of Incentives." *European Economic Review*, 2002, 46(4/5), pp. 687-724.

[18] Fuchs, William. "Subjective Evaluations: The Bonus as a Signal of Performance." Unpublished Paper, 2004.

[19] Garvey, Gerald T. "Why Reputation Favors Joint Ventures over Vertical and Horizontal Integration: A Simple Model." *Journal of Economic Behavior and Organization*, 1995, 28(3), pp. 387-97.

[20] Gibbons, Robert and Waldman, Michael. "Careers in Organizations: Theory and Evidence," in Orley C. Ashenfelter and David Card, eds., *Handbook of Labor Economics.* Vol. 3. Amsterdam: Elsevier North Holland, 1999, pp. 2373-2437.

[21] Holmstrom, Bengt. "Moral Hazard in Teams." The Bell Journal of Economics, Vol 13, No.2. 1982, pp.324-340.

[22] Holmstrom, Bengt and Milgrom, Paul. "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership and Job Design." *Journal of Law, Economics and Organization*, 1991, 7, pp. 24-52.

[23] Kahn, Charles and Mookherjee, Dilip. "A Competitive Efficiency Wage Model with Keynesian Features." *Quarterly Journal of Economics*, 1998, 103(4), pp. 609-45.

[24] Kandori, Michihiro. "Introduction to Repeated Games with Private Monitoring." *Journal of Economic Theory*, 2002, 102(1), pp.1–15.

[25] Kandori, Michihiro and Matsushima, Hitoshi. "Private Observation, Communication and Collusion." *Econometrica*, 1998, 66(3), pp. 627-52.

[26] Krueger, Alan B. and Summers Lawrence H. "Efficiency Wages and the Inter-Industry Wage Structure." *Econometrica,* Vol 56, No. 2. 1988, pp.259-293.

[27] Lazear, Edward and Rosen, Sherwin. "Rank-Order Tournaments as Optimum Labor Contracts." *Journal of Political Economy*, 1981, 89(5), pp.841-64.

[28] Lazear, Edward. "Output-Based Pay: Incentives or Sorting?" in Solomon W. Polachek, eds., *Research in Labor Economics*. Vol.23. Greenwich, CT: JAI Press, 2005, pp.1-25.

[29] Levin, Jonathan. "Multilateral Contracting and the Employment Relationship." *Quarterly Journal of Economics*, 2002, 117(3), pp. 1075-1103.

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

[31] MacLeod, W. Bentley. "Optimal Contracting with Subjective Evaluation." *American Economic Review*, 2003, 93(1), pp. 216-40.

[32] MacLeod, W. Bentley and Malcomson, James M. "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica*, 1989, 57(2), pp. 447-80.

[33] MacLeod, W. Bentley and Parent, Daniel. "Job Characteristics and the Form of Compensation," in Solomon W. Polachek, eds., *Research in Labor Economics*. Vol. 18. Greenwich, CT: JAI Press, 1999, pp.177-242.

[34] Malcomson, James M. "Work Incentives, Hierarchy, and Internal Labor Markets." *Journal of Political Economy*, 1984, 92(3), pp. 486-507.

[35] Malcomson, James M. "Individual Employment Contracts," in Orley Ashenfelter and David Card, eds., *Handbook of labor economics*. Vol. 3B. Amsterdam: Elsevier North Holland, 1999, pp.2291-372.

[36] Millward, Neil; Stevens, Mark; Smart, David and Hawes, W. R. *Workplace industrial relations in transition: The ED/ESRC/PSI/ACAS surveys*. Aldershot, U.K.: Dartmouth, 1992.

[37] Pearce, David and Stacchetti, Ennio. "The Interaction of Implicit and Explicit Contracts in Repeated Agency." *Games and Economic Behavior*, 1998, 23(1), pp. 75-96.

[38] Prendergast, Canice. "The Provision of Incentives in Firms." *Journal of Economic Literature*, 1999, 37(1), pp. 7-63.

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

[40] Rosen, Sherwin. "Prizes and Incentives in Elimination Tournaments." *American Economic Review*, 1986, 76(4), pp. 701-15.

[41] Sannikov, Yuliy. "A Continuous-Time Version of the Principal-Agent Problem." Unpublished Paper, 2004.

[42] Shapiro, Carl and Stiglitz, Joseph E. "Equilibrium Unemployment as a Worker Discipline Device." *American Economic Review*, 1984, 74(3), pp. 433-44.

[43] Shavell, Steven. "Risk Sharing and Incentives in the Principal and Agent Relationship." *Bell Journal of Economics*, 1979, 10(1), pp. 55-73.