The Optimality of Arm's Length, Memoryless Contracts

John Y. Zhu[1]

June 19, 2015

Abstract

For most workers, the primary source of incentives is managerial feedback. In this environment, I show why long-term employment relationships are optimally governed by a series of arm's length, memoryless shorter-term contracts: The worker is evaluated infrequently and his future prospects depend on current evaluations, but not on evaluations from long ago. The setting is a general repeated moral hazard model with costless monitoring and complete contracts, where the manager could, in theory, employ a highly complex, history-dependent long-term contract. The intuition for the optimality of arm's length, memoryless contracts is robust, applicable, in particular, to settings with risk aversion.

JEL Codes: C73, D86, J41, M5
Keywords: labor contracts, evaluation, principal-agent, arm's length, memoryless, renegotiation, robust contracts, continuous-time.

[1]zhuyiran@wharton.upenn.edu, The Wharton School, University of Pennsylvania. I thank Steve Matthews, Bruno Biais, Daniel Gottlieb and Mike Waldman for valuable comments. I also thank the seminar participants at University of Pennsylvania, 2014 Shanghai Microeconomics Workshop, Cornell, London School of Economics, London Business School, Toulouse School of Economics, Peking University Guanghua School of Management and Southern Methodist University.

1 Introduction

While profit may be the ultimate measure of a firm's value, it is often a poor signal of individual worker performance. Therefore, tying a worker's welfare to firm profit can be a bad decision, providing very little motivation for effort and, yet, exposing the worker to unnecessary risk. Instead, it is better for most workers to draw incentives from managerial feedback that can provide information much more attuned to an individual's performance. How should feedback-based employment contracts be designed? How frequently should feedback occur? How should feedback map into future prospects? In this paper, I provide some robust answers to these questions by studying optimal contracting with renegotiation-proof equilibria in a repeated moral hazard model where a manager evaluates and observes private signals of worker performance.[2]

In many employment relationships, the manager refrains from constantly monitoring the worker. Instead, occasional evaluations help maintain a coarse incentive system that determines future prospects. Moreover, these prospects depend on current evaluations in a straightforward way, but not on evaluations from long ago. Thus, overall, the worker is governed by a series of arm's length, memoryless contracts of finite duration. This paper provides a robust intuition for why such basic arrangements are optimal for inducing worker effort, even in the presence of risk aversion, complete contracts and costless monitoring. Such an intuition has so far been lacking in the theory of optimal contracts.

The infrequent evaluations feature of the optimal contract reflects the wisdom of hands-off management, that recognizes how too much monitoring can lead to worker disengagement. For example, Google maintains a sparse management structure that is purposefully designed so as to prevent the temptation to micromanage (Garvin, 2013). On the other hand, at modern call centers, the ability of managers to provide frequent and almost instant feedback has been shown to lead to worker disengagement (Morse, 2004).[3] In this paper, I describe an incentive channel through which frequent evaluations lead to worker disengagement. This provides a causal link from too much feedback to poor performance that does not rely on cognitive limitations and biases.

Despite the optimality of infrequent evaluations, a manager could still, in theory, employ a highly history-dependent, long-term contract. In practice, however, a contract has a definite termination date that separates it from the next contract and serves as a barrier to long-term history-dependence.[4] For example, a worker's contract after receiving a promotion is

[2]To the best of my knowledge, my paper is the first to systematically study optimal contracting when the equilibrium is required to be renegotiation-proof. As I will argue, this is a natural approach to optimal contracting in the feedback-based moral hazard setting I consider. The large literature on renegotiation-proof contracting is related but distinct. There, the focus is on contracts that involve renegotiation-proof *mechanisms*. There is also a related literature in game theory that studies renegotiation-proof equilibria but the focus is not on optimal game design. See Related Literature.

[3]Evidence from recent Gallup research suggests that disengagement is a substantial drain on firm profitability: A 2012 analysis of 21,213 business units in 75 organizations across the world found that top-quartile units average 22% greater profitability than bottom-quartile units. See Harter et al (2013).

[4]There has been some theoretical work, e.g. Fudenberg et al (1990), showing how certain optimal long-term contracts can be implemented as a series of short-term contracts. However, the implementation relies on the ability for the choice of today's short-term contract to depend on yesterday's performance in yesterday's

typically no longer sensitive to pre-promotion performance.[5] In my model, I show how the optimally infrequent evaluations break the formal long-term contract faced by the worker into a series of memoryless shorter-term contracts. Moreover, each component contract is a structurally simple pass/fail-type contract. At its end, the component contract determines whether the agent is above or below the bar and makes no finer distinctions. If the agent is above the bar, he is moved on to the next memoryless component contract. Otherwise, he is terminated.

The setting I consider is a general repeated moral hazard model where each period a worker exerts hidden effort and the manager observes private signals of worker effort. One can think of such signals as the manager's private judgement of the worker's performance when he conducts an evaluation. After forming a judgement, the manager makes a report and then the worker receives some pay and is either retained for the next period or terminated.

To understand why optimal contracts are memoryless in this setting, consider the manager deciding whether or not to report that the worker should be retained. Instead of using a strategy where the report decision is based only on the private judgement from this evaluation, the manager could, in theory, use a highly history-dependent strategy. For example, he may decide to report that the worker be retained if and only if his private judgement this time around plus at least 50% of all the private judgements from previous evaluations exceed a common threshold. The crux of my argument for why optimal contracts are memoryless is that such a strategy would invariably get revised into something that did not depend on the outcomes of previous evaluations by the time it was used: Upon entering the current evaluation period, all of the worker's previous efforts are now sunk. To continue to tie the worker's prospects to private judgements based on previous work will only dull the worker's incentives today. This not only hurts the worker but also the manager. Thus, it is in the interest of both parties upon entering the current evaluation period to revise the manager's report strategy in such a way so that it depends only on the evaluation of current work.

The memoryless feature of the optimal contract depends crucially on the manager and worker being able to revise their equilibrium play in the middle of the relationship if it is in their mutual interest to do so. But, of course, what's best for the manager and worker in the middle of the relationship may not be what's best for them originally. Thus, ex-ante, both parties may want to commit not to make such revisions in the middle of the relationship. But assuming the ability to commit not to make such revisions is unrealistic: The manager's report strategy is *private*, as is the worker's effort strategy, and a contractual clause dictating that the manager cannot revise his private report strategy and the worker cannot revise his private effort strategy is, in practice, unenforceable.

Notice this intuition for memoryless contracts differs from the one used in the relational contracts setting of Levin (2003). There, the idea is that one can always take a history-dependent optimal contract and change the pay structure in a payoff-neutral way so that the incentives provided by variations in promised value are replaced by incentives provided by transitory bonuses. Such an argument only works when both parties are risk-neutral and

short-term contract. Thus, the particular details of today's contract will still depend on the entire history of performance leading up to today.

[5] Baker et al (1988) notes the prevalence of horizontal equity systems in which workers at the same level are paid equally, thereby ignoring the differing performance histories of the workers.

discount future payoffs in the same way. In contrast, my intuition essentially boils down to an appeal for focusing on contracts that involve renegotiation-proof equilibrium strategies.

What about a rationale for infrequent evaluations? To fix ideas, suppose every day the worker can choose an effort that then generates a "score" at the end of the day. The higher is the effort, the more likely the score is high. Suppose whenever the manager evaluates the worker, he observes the sum of the scores generated between the previous evaluation and the current one. This is the manager's private signal. Thus, if the manager evaluates the agent every day, he sees a score for every day's effort choice. On the other hand, if he evaluates the agent every year, then he only sees yearly aggregate scores.

When I say that infrequent evaluations are optimal, I mean that the optimal contract calls for the manager to only periodically collect aggregate scores of the worker even though it is costless for him to constantly collect the more detailed, day-by-day scores.

When dates are short, individual efforts are able to generate only noisy, incremental scores. A memoryless contract that incentivizes today's effort using only today's score will be unable to induce much effort. The result is similar to the limits to collusion findings of Abreu et al (1991) and Sannikov and Skrzypacz (2007). Therefore, to be able to effectively provide incentives, rewards must be linked to good performance over many days. If the manager could commit to a particular private strategy, the optimal thing to do would be to collect daily scores over a long-period of time and then make a recommendation with major consequences for the worker based on all the scores the manager has collected. For example, at the end of the long-period, the manager could decide whether to keep the worker or terminate him, and base his recommendation on whether or not on at least 50% of the days he collected a score that exceeds some reasonable threshold. This is similar to the strategy that Abreu et al (1991) suggests to overcome the cooperation problem they study. Unfortunately, this is also precisely the kind of strategy I considered and ruled out earlier: A rational worker realizes that as time passes in this evaluation period, he and the manager will constantly want to revise the recommendation strategy so that it only depends on work that's yet to be done. As a result, the worker will exert no effort except for the last day. The only way to overcome this incentive breakdown is for the manager to collect only an aggregate score across the entire evaluation period. This way, the worker and manager simply do not have the detailed day-by-day scores needed to implement their ex-interim beneficial but ex-ante harmful strategy revisions.

Thus, infrequent evaluations serve as a valuable commitment device in a game with renegotiable private strategies. The optimal degree of infrequency - that is, the optimal number of days between the times when the manager collects an aggregate score - is determined by an intuitive tradeoff. Space the evaluation dates too close to each other and the inability to commit to history-dependent private strategies breaks down incentives. Space the evaluation dates too far from each other and the poor quality of the aggregate information breaks down incentives. The optimal evaluation frequency balances the two opposing forces.

Related Literature

There is a literature in psychology that looks at feedback frequency and its effect on performance. A number of recent studies have challenged the longstanding idea that more feedback means better performance. For example, experimental evidence from Lam et al (2011) suggests "an inverted-U relationship between feedback frequency and task performance," which is also a main theoretical prediction of my paper. However, most explanations of the negative effects of high feedback frequency revolve around cognitive biases and limitations (e.g. Lurie and Swaminathan, 2009). This paper looks at how the frequent use of feedback within the firm can hinder a manager's ability to effectively provide incentives to a rational agent.

As explained earlier, the infrequent evaluations of the optimal contract serve to break the long-term arrangement into a series of memoryless contracts, each of finite duration. Dye (1985) studies the optimal length of labor contracts in a setting where contracts have rigidities and recontracting is costly. Harris and Holmstrom (1987) study a long-term borrower-lender relationship where the parties can pay to observe a costly informative signal of project riskiness which varies over time. Like my paper, they consider, formally, a long-term contracting setting and define contract duration to be the number of dates between consecutive information acquisitions in the long-term contract. They also find that each component contract is memoryless with respect to its predecessors. However, there has been little research done tying contract length to agency frictions.

More generally, this paper seeks to provide a robust intuition for a broad class of simple contracts. The disparity between the simple contracts frequently observed in practice and the typically complex dynamic contracts predicted by theory is puzzling. Empirically, Baker et al (1988) documents the paucity of transitory bonuses; Prendergast (1999) describes widespread wage and evaluation compression; and Posner et al (2004) points out how contracts typically "divide the future into very crude partitions." These are all salient features of the arm's length, memoryless optimal contracts that emerge in my theoretical analysis.

Methodologically, this paper is related to the literature on renegotiation-proof relationships. The term "renegotiation-proof" is somewhat ambiguous. In the optimal contracting world - where a contract is formally a contract mechanism plus an equilibrium of the mechanism - the renegotiation-proof condition is typically applied to the mechanism rather than the equilibrium. In contrast, game theory papers such as Benoît and Krishna (1993) apply the renegotiation-proof condition on the equilibria of the game. But there the focus is not so much on some notion of optimal game design but, rather, the limiting properties of the renegotiation-proof equilibrium payoff set for general repeated games.

My approach to renegotiation-proof relationships straddles the contracting and game-theoretic approaches. I focus on optimal mechanism design, but instead of applying the renegotiation-proof condition on the mechanism, I apply it to the equilibria of the mechanism like in game theory. In the feedback-based moral hazard setting, the choice to apply the renegotiation-proof condition to the equilibria is particularly compelling. Equilibria in my model involve private strategies: The worker applies hidden effort and the manager makes reports based on privately observed signals. Therefore any renegotiation of strategies can be done privately, making it difficult for the manager and worker to set up a commitment device against renegotiation ex-ante. This is in contrast to renegotiations of the mechanism (or renegotiations of an equilibrium involving public strategies) that can be prevented by contracts with third-parties or access to institutions with commitment power.

This paper is closely related to Fuchs (2007) and is part of the subjective evaluation literature (see also MacLeod, 2003, and Levin, 2003 Section IV). The key difference between

Fuchs (2007) and this paper is that Fuchs (2007) does not restrict the equilibrium strategy to be renegotiation-proof. In Section 2.1 I show how a Fuchs (2007)-type contract in my setting would violate renegotiation-proofness in a very strong way, leading the worker to exert zero effort at all dates except the last one. My model is similar to that of Levin (2003) Section IV. The setting is a repeated game with private monitoring and as Levin (2003) points out, in such a setting it is not even assured that the optimal contract is recursive, much less memoryless and arm's length. To make progress, Levin (2003) restricts attention to only those contracts that induce the manager to fully reveal his private judgement after each evaluation.

Finally, this paper is related to a small literature on robust contracting that aims to show how certain simple, tractable arrangements may be optimal or approximately optimal in some robust sense. Recent examples include Chassang (2013) and Carroll (2015).

2 Pareto-Optimal PRP Contracts

In this section I introduce the repeated principal-agent moral hazard model and prove that Pareto-optimal contracts subject to the principal and agent playing renegotiation-proof strategies are memoryless. Here, the agent represents the worker and the principal represents the manager and ownership. This paper does not consider any potential agency problems arising from the manager-ownership relationship. Rather the two parties are assumed to be aligned and treated as one entity.

I call the contracts with renegotiation-proof strategies *private-revision-proof* (PRP) contracts to emphasize that players can privately coordinate to revise their continuation equilibrium play and to distinguish them from the contracts of the renegotiation-proof contracting literature that apply the renegotiation-proof condition to the mechanism rather than to the equilibrium. In this section, the evaluation frequency is exogenously fixed. In the next section, I allow the contract to stipulate when evaluations occur and show that, in addition to being memoryless, Pareto-optimal PRP contracts are also arm's length.

I begin by considering a benchmark risk-neutral model before generalizing to recursive preferences.

2.1 The Benchmark Risk-Neutral K-Model

The complete dynamic model is a static segment repeated K times. I call it the K-model. A segment starts after the beginning of date t and ends at the beginning of date $t+1$ for $t = 0, 1, \dots K-1$. Figure 1 shows a generic date t segment.

The principal and agent are governed by a long-term contract spanning the entire K-model. In the segment starting from date t, the contract first makes a public transfer w_t from the principal to the agent.[6] The agent then chooses hidden effort $a_t \in [0, \infty)$ with a strictly convex, differentiable cost function $h(a_t)$ satisfying $h'(0) = 0$. This effort generates a private random utility ${}_t\rho_{t+1}$ for the principal. This utility is the principal's private informative signal

[6]There are no secret side-payments that circumvent the contract. Also, there is no cash burning for now. Later, I allow for cash burning and show that it is redundant. See Lemma 3.

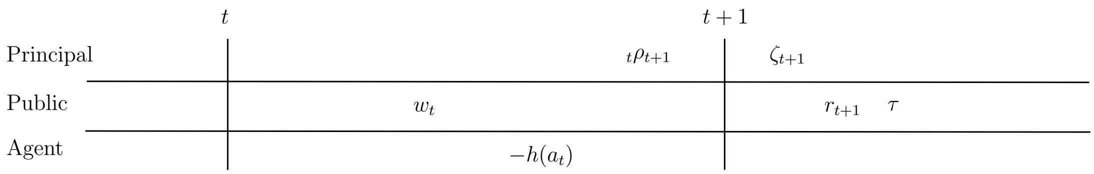

Figure 1: The three lanes that run across the figure denote what is observable to the principal only, to the public, and to the agent only.

and is distributed according to the density $f(\cdot, a_t)$. I assume that $f(\cdot, a_t)$ is continuously differentiable and satisfies the monotone-likelihood-ratio-property (MLRP) with respect to effort. This means, given any two effort levels $a < a'$, the ratio $f(\rho, a')/f(\rho, a)$ is strictly increasing in ρ - higher utility suggests higher effort. At the beginning of date $t+1$, the principal privately observes an uninformative randomizing device ζ_{t+1} and makes a public report r_{t+1} taking values in some contractual message space. This marks the end of the segment.

The contract now makes a termination decision represented by a stopping time τ and sets an upcoming transfer w_{t+1}. If the agent is retained, then the next segment starts with transfer w_{t+1}. Otherwise, the agent is terminated, $\tau = t+1$, and $w_\tau = w_{t+1}$ is the severance pay. Both parties exercise their outside options at date τ with values equal to $(O_A(w_\tau), O_P(-w_\tau))$, where $O_A(w_\tau) = O_A + w_\tau$ and $O_P(-w_\tau) = O_P - w_\tau$. At the end of the last segment, termination is automatic.

The date 0 transfer w_0 is the date 0 component of the contract's mechanism. For $t > 0$, the date t termination decision and transfer comprise the date t component of the contract's mechanism. These two functions can depend on all prior public information: $h_t := \{r_i\}_{i=1}^t$. For $t < K$, the agent's date t effort choice depends on h_t, and the principal's report at the beginning of date $t+1$ depends on the public history as well as the principal's private history: $H_{t+1} := h_t \cup \{_i\rho_{i+1}, \zeta_{i+1}\}_{i=0}^t$. Formally, a contract consists of a mechanism $\{\tau, \{w_t\}_{t=0}^K\}$ specifying the termination clause and pay structure, and a strategy profile $\{a, r\} := \{a_t, r_{t+1}\}_{t=0}^{K-1}$ specifying the agent's recommended effort process and the principal's recommended report rule.

Given such a contract, the agent's computed value of his own continuation payoff at the beginning of the date t segment for $t < K$ is the expected present discounted value of his future earnings less his future effort costs plus his future outside option value.[7]

$$W_t(h_t) = \mathbf{E}_{\{a,r\}}\left[\sum_{s=t}^{\tau-1} \beta^{s-t}(w_s - h(a_s)) + \beta^{\tau-t} O_A(w_\tau) \middle| h_t\right]$$

The principal's computed value of his own continuation payoff at his date $t+1$ decision

[7] I emphasize that $W_t(h_t)$ is what the agent thinks he will receive because, in general, the principal will have a better idea of what the agent should expect to get due to his payoff relevant private information. The principal's computed value of the agent's continuation payoff is the same expectation but conditioned over the finer information set $H_t \cup r_t$.

node, occurring right before the start of the date $t+1$ segment, is

$$V_{t+1}(H_{t+1}) = \mathbf{E}_{\{a,r\}}\left[\sum_{s=t+1}^{\tau-1} \beta^{s-(t+1)}({}_s\rho_{s+1} - w_s) + \beta^{\tau-(t+1)} O_P(-w_\tau) \middle| H_{t+1}\right]$$

Unlike typical moral hazard problems where the incentive-compatibility (IC) constraint is purely on the agent side, here, the principal also has an IC constraint. To see this, suppose at the beginning of some date t, the principal can report one of two messages, a or b. If his continuation payoff is strictly larger following a then he will always report a. He can do this even if it violates the contract's recommended report strategy because his strategy is private and deviations cannot be detected. Therefore, in order for the principal to be willing to report different messages following different private histories, he must be indifferent between reporting any of them. This is an insight of MacLeod (2003). The following lemma formalizes the intuition.

Lemma 1. *Let $\{a, r\}$ be a perfect Bayesian equilibrium. Then $V_{t+1}(H_{t+1})$ is h_t-measurable.*

Proof. See appendix. □

Thus, when the strategy profile is a perfect Bayesian equilibrium, I can write $V_{t+1}(h_t)$ for $V_{t+1}(H_{t+1})$. $V_{t+1}(h_t)$ is also the principal's computed value of his own continuation payoff at the beginning of the date $t+1$ segment. However, in order to conceptually distinguish between the two continuation payoffs, I will call the latter $V_{t+1}(h_{t+1})$.

Definition. *A contract is incentive-compatible if its strategy profile is a perfect Bayesian equilibrium.*

The key implication of Lemma 1 is that in an incentive-compatible contract the principal's computed value of his own continuation payoff is common knowledge. This paves the way for introducing a notion of renegotiation-proof equilibrium play.

Definition. *A perfect Bayesian equilibrium is private-revision-proof (PRP) at the beginning of the date t segment if the continuation play is weakly Pareto-optimal among all continuation equilibria that are PRP at the beginning of the date $t+1$ segment. A contract is PRP if its perfect Bayesian equilibrium is PRP at the beginning of the initial segment.*

I also introduce an interim participation constraint that requires there to not exist, for any t, a w such that $O_A(w) > W_t(h_t)$ and $O_P(-w) > V_t(h_t)$. The assumption being, if the condition were violated, there would be some date t and some payment w such that both parties could profitably separate after the beginning of date t with severance pay w.

From now on, the term "contract" will denote a PRP contract that satisfies the interim participation constraint.

The PRP definition is essentially the standard one for a renegotiation-proof equilibrium. There are some technical issues that have been swept under the rug due to the fact that the principal has private information, and can, in general, form more accurate expectations about future payoffs than the agent. Similar issues arise in the definition of the interim participation constraint - e.g. what if $W_t(h_t) < O_A(w)$ but $W_t(H_t \cup r_t) > O_A(w)$? It turns out, under the natural definition of renegotiation-proof equilibrium, the principal's computed value of the agent's continuation payoff will actually coincide with the agent's computed value, and so the potential thorny technically issues never materialize. See Appendix A for the precise formulation of PRP.

The PRP condition is a natural condition to impose on the play of a game with private strategies. Its importance comes from the fact that it predicts that the principal provides feedback-based incentives in a min/max threshold form.

Lemma 2. *Fix a perfect Bayesian equilibrium and a public history h_t. If the equilibrium is PRP at the beginning of the date $t+1$ segment, then it is PRP at the beginning of the date t segment only if on the set of compatible H_{t+1} the principal employs a min/max threshold report strategy: There is some threshold ρ_{t+1} such that if ${}_t\rho_{t+1} > \rho_{t+1}$ then the principal reports the message that maximizes $W_{t+1}(h_{t+1})$. Otherwise, the principal reports the message that minimizes $W_{t+1}(h_{t+1})$.*

Proof. See appendix. □

The idea is that the principal can always replace a non min/max threshold report strategy with the unique min/max threshold report strategy that leaves the agent's continuation payoff unchanged under his original effort choice. Because the min/max threshold report strategy utilizes the most extreme punishment/reward combination available given the contract mechanism, it provides the agent strictly more incentives to work than the original report strategy. This induces the agent to work harder which makes the principal better off. The agent is better off as well since he could always attain his original payoff by choosing to exert his original effort choice.

While the PRP condition implies the min/max threshold report structure, it does not determine what the min and max outcomes should be for the agent at any particular point in time. The nature of these outcomes are determined once attention is restricted to Pareto-optimal PRP contracts.

To see this, consider the first evaluation in a Pareto-optimal PRP contract, occurring at the beginning of date 1. Intuitively, since the contract is Pareto-optimal, if the agent is above the bar, he should be rewarded as much as possible. This means, starting at date 1, the principal and agent should be governed by a Pareto-optimal PRP contract over the remaining $K-1$ segments of the model. On the other hand, if the agent is below the bar, the punishment should be as low as possible because the more painful is the punishment, the lower the bar has to be set which means the less the punishment actually has to occur on the equilibrium path. This means a more efficient relationship ex-ante. The interim participation constraint implies that the worst possible punishment the contract can inflict is termination.

Thus, I have now shown that in the first segment of a Pareto-optimal PRP contract, the agent puts in some effort a_0^* and the principal sets a threshold ρ_1^*. If the agent is above the bar, a Pareto-optimal PRP contract over the remaining $K-1$ segments is enacted, otherwise the agent is terminated. It is worth emphasizing, conditional on the agent being above the bar, the evaluation makes no finer distinctions. This means that the Pareto-optimal PRP contract over the remaining $K-1$ segments that is enacted is the same one no matter how much the agent is above the bar.[8] But now, since the continuation contract starting at the second segment is itself Pareto-optimal PRP, we may conclude, by recursion, that the relationship in the second segment has the same threshold structure as in the first segment. Thus, in the second segment, the agent again exerts some effort a_1^* and the principal again sets some threshold ρ_2^*. If the agent is above the bar, then starting at date 2, a Pareto-optimal PRP contract over the remaining $K-2$ segments is enacted and so on. Note, a_1^* and ρ_2^* are constants just like a_0^* and ρ_1^* because there is only one continuation contract starting at the second segment. Following this recursive process to its end, I obtain the sequence of constants $\{a_t^*, \rho_{t+1}^*\}_{t=0,1,...K-1}$ of efforts and performance thresholds characterizing the original Pareto-optimal PRP contract. Moreover, since the current analysis is for the special risk-neutral, equal discounting case, the Pareto-frontier is linear and can be generated by taking one particular Pareto-optimal PRP contract and simply shifting the initial wage. This means that the optimal sequence of constants $\{a_t^*, \rho_{t+1}^*\}_{t=0,1,...K-1}$ is actually shared by all Pareto-optimal PRP contracts.

Proposition 1. *There exists a unique sequence of constants* $\{a_t^*, \rho_{t+1}^*\}_{t=0,1,...K-1}$ *such that all Pareto-optimal PRP contracts share the following memoryless, pass/fail structure: For each date* $t < \tau$*, the agent puts in effort* a_t^**. At the beginning of date* $t+1$*, the principal reports one of two possible messages - pass or fail. The principal reports pass if and only if* ${}_t\rho_{t+1} > \rho_{t+1}^*$*. If the agent passes then the relationship continues into date* $t+1$*. Otherwise the principal reports fail which leads to termination.*

Fix a relationship governed by a Pareto-optimal PRP contract and pick any date t*. The unique continuation contract governing the relationship starting at date* t *is a Pareto-optimal PRP contract over the remaining* $K-t$ *segments.*

Proof. See appendix. □

Notice, the only time risk-neutrality is invoked during the arguments leading to Proposition 1 is when it is claimed that all Pareto-optimal contracts share the same sequence of constants $\{a_t^*, \rho_{t+1}^*\}_{t=0,1,...K-1}$. On the other hand, the logic for why each Pareto-optimal PRP contract is characterized by *some* sequence of constants $\{a_t^*, \rho_{t+1}^*\}_{t=0,1,...K-1}$ is driven purely by the PRP assumption on contracts and the MLRP assumption about ${}_t\rho_{t+1}$ (used to prove Lemma 2). One therefore expects that the memoryless, pass/fail structure of Pareto-optimal PRP contracts should survive under much more general settings, including non risk-neutral settings. We will see shortly that this will, indeed, be the case.

In Proposition 1, the sequence of efforts and thresholds that characterize all Pareto-optimal PRP contracts can be calculated in a straightforward way via backwards induction. At date t, Proposition 1 tell us that the agent is facing a pass/fail evaluation where what's at stake is the future relationship starting at date $t+1$. Proposition 1 also tells us that this future relationship is Pareto-optimal, governed by a Pareto-optimal PRP contract over the remaining $K-(t+1)$ segments. Let S_{t+1}^* denote the surplus value of this relationship

[8]Recall, by Lemma 1, the principal's continuation payoff at date 1 is the constant V_1. So the enacted Pareto-optimal PRP contract over the remaining $K-1$ segments is the unique one that delivers total payoff V_1 to the principal.

- the Pareto-optimal date $t+1$ continuation surplus. This is the difference in payoff for the agent between being retained and being terminated. Given S^*_{t+1} and an arbitrary pass/fail threshold ρ_{t+1}, the agent's best response is to choose effort $a(S^*_{t+1}, \rho_{t+1})$ at date t. The relationship starting at date t then generates surplus

$$u(a(S^*_{t+1}, \rho_{t+1})) - h(a(S^*_{t+1}, \rho_{t+1})) + \beta(1 - F(\rho_{t+1}, a(S^*_{t+1}, \rho_{t+1})))S^*_{t+1} - (1-\beta)(O_A + O_P). \tag{1}$$

Here $F(\cdot, a)$ is the cumulative distribution function associated to the density $f(\cdot, a)$. By Proposition 1, we know that the relationship starting at date t is also Pareto-optimal. Therefore, the principal chooses ρ^*_{t+1} to maximize (1). This then determines optimal effort $a^*_t = a(S^*_{t+1}, \rho^*_{t+1})$, as well as the Pareto-optimal date t continuation surplus S^*_t, which is equal to (1) evaluated at ρ^*_{t+1}.

I have now shown how to compute the Pareto-optimal threshold ρ^*_{t+1}, effort a^*_t and date t continuation surplus S^*_t given the Pareto-optimal date $t+1$ continuation surplus S^*_{t+1}. Since at the very end of the relationship there is no surplus, $S^*_K = 0$. Now the entire Pareto-optimal effort and threshold sequence can be computed using backwards induction.

As for the pay structure, risk-neutrality and equal discounting imply that it is largely indeterminate. The only restrictions come from the principal's IC constraint and the interim participation constraints.

In Proposition 1's implementation of the Pareto-optimal PRP contracts, the principal only reveals whether the agent is above or below the bar but not by how much. Formally, the message space of the contract consists of $\{pass, fail\}$. However, in PRP contracts, whatever information about the agent's performance that is not revealed will not play a role in future incentive provision. Thus, these contracts can be equivalently implemented so that the message space is the entire range of the principal's private signal and the principal, at each date, fully reveals his private judgement of the agent but acts based only on whether the agent's performance is above or below the bar. Such a fully-revealing implementation satisfies the *full-review property* of Levin (2003) Section IV although most full-review contracts are not PRP. The PRP condition provides a game-theoretic motivation for focusing on a subset of full-review contracts.[9]

The mechanics of the Pareto-optimal contracts are intuitive. The incentive structure is not burdened in any way by history and is, instead, forward-looking. At each evaluation, what's at stake for the agent is simply the surplus value of the future relationship. Nothing more, less, or different is in play. The nature of this incentive structure has a number of implications.

First, it means that the longer is the relationship the greater is the effort that can

[9] For an example of a contract that cannot be implemented as a full-review contract, consider one where the principal reports *pass* at the end of the first segment no matter what but at the end of the second segment reports *pass* if and only if both private signals are positive. Such a report strategy cannot be implemented in a way that satisfies the full-review property. This is because if the principal fully reveals his private signal at the end of the first segment, then in the case that the first signal is negative, the agent knows that there is no chance that he can pass at the end of the second segment which means he has no incentive to work in the second segment.

be induced. In fact, the forward-looking incentive structure means that, given a shorter relationship lasting K_1 segments and a longer relationship lasting K_2 segments, the effort sequence induced over the last K_1 segments of the longer employment relationship matches the effort sequence of the shorter employment relationship. Thus, the value-added for the longer relationship is coming entirely from the initial phase of the relationship. This initial phase is when the surplus value of the future relationship is largest and therefore when the induced effort is highest. Thus, there are increasing returns of effort to relationship length.

A companion implication is that, over time, as the relationship winds down and the remaining surplus value of the relationship declines, the amount of effort that can be induced declines as well. In the current model, this means that the principal faces an agent that becomes less productive over time.

However, in general, Proposition 1 need not be interpreted as predicting that the agent is less productive over time. The model can be generalized so that each segment of the K-model is unique, with its own particular signal distribution $f_t(\cdot, a_t)$, effort cost function h_t etc. For example, one can assume that for all $t > 0$, $f_t(\cdot, a) = f_0(\cdot - E_t, a)$ where $E_1, E_2, \ldots E_{K-1}$ is a sequence of increasing, positive constants. In this case, the agent gains experience over time and the same level of effort generates higher and higher returns over time as measured by E_t. It is easy to show that Proposition 1 continues to hold. In this case, despite declining incentives and effort levels, productivity can still increase.

Changing States

More generally, the optimality analysis can accommodate settings where the nature of the employment relationship at any time depends on some evolving state variable which captures conditions in the economy or within the firm. For example, productivity may depend on other inputs of the firm which may vary over time. Or the outside options of the principal and agent may depend on the overall health of the economy which also may be changing and may not change in the same way as the state of the firm.

Formally, suppose f, h, O_A and O_P all depend on a (possibly time-inhomogenous) Markov state-of-the-world variable s. Then Proposition 1 still holds except that for each date t, the optimal effort $a_t^*(s)$ and optimal threshold $\rho_t^*(s)$ are state dependent.

Cash-Burning

The K-model assumes any w taken from the principal is given to the agent. A more general approach is to allow the mechanism to "burn cash" by giving the agent only a portion of the w and burning the rest. If an amount $b > w$ is to be burned, then the agent receives nothing and must, in addition, contribute $b - w$ from his own pocket to be burned. Relaxing the ex-post budget balance constraint by allowing cash-burning gives the mechanism the added flexibility to destroy arbitrarily large amounts of surplus unlike termination. In theory, this is useful because incentives are provided by putting surplus at risk. However, the interim participation constraint prevents a contract from destroying any more surplus than termination can destroy. As a result,

Lemma 3. *Every Pareto-optimal PRP contract with cash-burning is payoff equivalent to one without cash-burning.*

More generally, the logic behind Lemma 3 can be used to show that no matter what type of surplus destruction technology one introduces into the model, any Pareto-optimal PRP contract can be rewritten so that the only source of surplus destruction is termination.

Without the interim participation constraint, cash burning would be strictly Pareto-improving. Indeed, by threatening to burn very large amounts of cash conditional on very low realizations of the private signals, the contract can approach first-best efficiency.

Yet, cash-burning does not play a prominent incentive role in real life contracts. Lemma 3 provides a simple reason why this is the case based on the assumption that the principal and agent can always agree to walk away from their relationship if it is in their mutual interest to do so.

The seminal paper on the cash-burning approach is MacLeod (2003). It is essentially the $K = 1$-model with cash burning except for two differences: ${}_0\rho_1$ is finite valued and there is no interim participation constraint at date 1. Additionally, MacLeod (2003) does not impose PRP, but in a static setting, the renegotiation-proof condition is degenerate. The lack of a date 1 interim participation constraint allows for mechanisms where, at date 1, following a "*good*" report, the principal and agent exercise outside options $(O_A(w_1), O_P(-w_1))$ while following a "*bad*" report, the principal still exercises $O_P(-w_1)$ but the agent exercises $O_A(w_1 - b)$ where $b > 0$ is cash burned. In fact, the optimal contract always has the aforementioned *good*/*bad* structure for some $b > 0$ where the principal reports *bad* if and only if the lowest value of ${}_0\rho_1$ is realized. If ${}_0\rho_1$ has a density like in my paper, then it is optimal to set $b = \infty$ and report *bad* on the probability zero event when ${}_0\rho_1$ realizes its lowest value.

Fuchs (2007) then considers a repeated version of the MacLeod (2003) static model. It is essentially the K-model with cash burning except, again, ${}_t\rho_{t+1}$ is finite valued and there is no PRP condition. The key insight of Fuchs (2007) is that one can treat each date as a different task, in which case the single-task dynamic model becomes a multi-task static model. The optimal contract is the multi-task version of the MacLeod (2003) optimal contract: the principal reports *bad* if and only if the lowest value of ${}_t\rho_{t+1}$ is realized for all $t < K$. This multi-task contract is implemented dynamically as follows: at each date $t < K$, there is a singleton message space $\{good\}$. On the final date, the principal can report *good* or *bad*, and *bad* is reported if and only if the worst history of signals is realized.

In the static setting, the key difference between MacLeod (2003) and my model is the interim participation constraint. In the dynamic setting, the most important difference between Fuchs (2007) and my model is the PRP condition.

To see why, forget about the interim participation constraint and consider the following contract in the K-model with cash burning which is in the spirit of Fuchs (2007): The principal sets a low threshold $\underline{\rho}$ and reports *good* at every date $t < K$. At date K, if ${}_t\rho_{t+1} < \underline{\rho}$ for all $t < K$, then the principal reports *bad* and $(O_A(w - b), O_P(-w))$ is exercised for some $b > 0$. Otherwise the principal reports *good* and $(O_A(w), O_P(-w))$ is exercised. Arguments in Fuchs (2007) can be used to show that any PRP contract can be Pareto-dominated by a contract of this type.

However, such contracts do not have renegotiation-proof strategies. Once the final agent decision node is reached, all previous efforts are sunk. At this point, tying the agent's payoff to signal realizations of previous efforts is equivalent to tying his payoff to uninformative random variables - this dulls incentives and reduces welfare. Following the memoryless alteration outlined in the proof of Lemma 2, any such dependence can be privately revised away in a strictly Pareto-dominant way so that, ultimately, the principal's report strategy depends only on the private judgement of the final effort. The agent, expecting this private revision, will exert zero effort at all his decision nodes except the final one.

2.2 Recursive Preferences

The memoryless pass/fail structure of Pareto-optimal contracts is robust to settings where players have general recursive preferences. Under such preferences, without certain concavity restrictions on the utility functions, the Pareto-frontier is a priori not public-randomization-proof unlike in the risk-neutral case. Thus, for each $t < K$, I add a public randomizing device ξ_t before the beginning of the date t segment. The public history is now $h_t := \xi_0 \cup \{r_i, \xi_i\}_{i=1}^t$ for $0 < t < K$ and $h_0 := \xi_0$.

Define the agent's recursive preference to be:

$$W_t(h_t) = \begin{cases} u_A(w_t, a_t) + W_t(h_t)^+ & t < \tau \\ O_A(w_\tau) & t = \tau \end{cases}$$
$$W_t(h_t)^+ = \beta_A f_A^{-1} \mathbf{E}_{r_{t+1},\xi_{t+1}} \left[f_A(W_{t+1}(h_{t+1})) \mid h_t \right].$$

$W_t(h_t)^+$ is the agent's date t ex-post continuation payoff, computed after he has sunk his date t effort cost and consumed his date t wage. u_A is continuous, strictly increasing in w with range $(-\infty, \infty)$, and, as a function of effort, $u_A''(w, \cdot) < 0$ and $u_A'(w, 0) = 0$. f_A is continuously differentiable and strictly increasing with range $(-\infty, \infty)$. $\beta_A > 0$ is the agent discount factor.

Similarly, define the principal's recursive preference to be:

$$V_t(h_t) = \begin{cases} f_P^{-1} \mathbf{E}_{a_t} \left[f_P(u_P(-w_t, {}_t\rho_{t+1}) + \beta_P V_{t+1}(h_t)) \mid h_t \right] & t < \tau \\ O_P(-w_\tau) & t = \tau \end{cases}$$
$$V_{t+1}(h_t) = f_P^{-1} \mathbf{E}_{r_{t+1},\xi_{t+1}} \left[f_P(V_{t+1}(h_{t+1})) \mid H_{t+1} \right]$$

Here, $V_t(h_t)$ is the principal's continuation payoff at the beginning of the date t segment and $V_{t+1}(h_t)$ is the principal's continuation payoff at his date $t+1$ decision node, occurring right before the date $t+1$ segment. u_P is continuous and strictly increasing in both arguments. For every ${}_t\rho_{t+1}$, $u_P(\cdot, {}_t\rho_{t+1})$ has range $(-\infty, \infty)$. f_P is continuous and strictly increasing with range $(-\infty, \infty)$. $\beta_A > 0$ is the principal discount factor.

O_A and O_P are both continuous and strictly increasing functions with range $(-\infty, \infty)$.

I can now state the theorem characterizing Pareto-optimal contracts. For simplicity, I will assume that there is no need for public randomization of the Pareto-frontier.

Theorem 1. *For each $t < K$, there exists a function $W_{t+1}(\,\cdot\,)$ such that for any non-terminal Pareto-optimal PRP contract over the last $K - t$-segments that delivers agent payoff W_t, the* *agent puts in effort $a_t^*(W_t)$ and is paid $w_t^*(W_t)$ at date t. At the beginning of date $t+1$, if ${}_t\rho_{t+1} > \rho_{t+1}^*(W_t)$ then the agent is retained and the Pareto-optimal PRP contract over the remaining $K-(t+1)$-segments that delivers agent payoff $W_{t+1}(W_t)$ is enacted. Otherwise the agent is terminated.*

Theorem 1 is the natural generalization of Proposition 1 except that different Pareto-optimal PRP contracts need not have the same action and threshold sequences. Moreover, as the statement of Theorem 1 implies, unlike the risk-neutral case, while some Pareto-optimal PRP contracts are non-trivial, others may lead to immediate termination.

The functions $W_{t+1}(\,\cdot\,)$ link, via the agent's continuation payoff, the date t's Pareto-optimal PRP allocation to date $t+1$'s for all $t < K$. This is unlike the risk-neutral case where the links are indeterminate because the Pareto-frontier is linear. Since the links $W_{t+1}(\,\cdot\,)$ do not depend on the principal's reports r_{t+1}, Pareto-optimal PRP contracts are memoryless.

The proof of Theorem 1 follows the same logic as that of Proposition 1. In particular, Lemma 2 still holds: The basic strategy for the proof is still to replace a non min/max threshold report strategy with the unique one that generates the same $W_t(h_t)^+$ under the original induced effort $a_t(h_t)$. The next step is, again, to show that for all $a < a_t(h_t)$ ($a > a_t(h_t)$), $W_t(h_t)^+$ is strictly smaller (greater) under the new strategy than under the original strategy. Since f_A is strictly monotonic, it suffices to show that for all $a < a_t(h_t)$ ($a > a_t(h_t)$), $f_A(W_t(h_t)^+/\beta_A)$ is strictly smaller (greater) under the new strategy than under the original strategy. Since $f_A(W_t(h_t)^+/\beta_A)$ is of the form $\mathbf{E} f_A(W_{t+1}(h_{t+1}))$, the original argument works here if the new strategy involves reporting the equilibrium message that maximizes (minimizes) $f_A(W_{t+1}(h_{t+1}))$ when ${}_t\rho_{t+1}$ is above (below) the threshold. Since f_A is strictly monotonic, maximizing/minimizing $f_A(W_{t+1}(h_{t+1}))$ is equivalent to maximizing/minimizing $W_{t+1}(h_{t+1})$ which is, by definition, what the new report strategy does. So the original argument works and the rest of the proof of Lemma 2 goes through exactly like in the risk-neutral case.

Lemma 2 is a min/max result that says it is optimal to give the agent as little as possible below a certain threshold and as much as possible above it. This type of result is well-known and has appeared before, most notably, as an optimality of debt result in Innes (1990). The difference between Lemma 2 and the version of the min/max lemma appearing in Innes (1990) and related papers lies in the assumptions. In addition to the MLRP assumption, those versions of the min/max lemma typically require bilateral risk-neutrality. Lemma 2 does not require the risk-neutrality assumption because the maximization problem has an extra, simplifying constraint driven by Lemma 1: The principal's continuation payoff needs to be constant over all equilibrium reports. This means that in my model, the increasing and decreasing of the agent's continuation payoff as a function of the principal's reports has to be done in a way that does not affect the principal's own continuation payoff. Thus, the potentially complex tradeoff between ex-ante incentive-provision and ex-post risk-sharing does not exist in my model.

After Lemma 2 is proved, the next step is to consider an arbitrary non-terminal date t Pareto-optimal PRP allocation with agent payoff W_t. Lemma 2 says there are at most two possible date $t+1$ continuation PRP contracts that can follow. Let V_{t+1} be the principal's continuation payoff that is common to both date $t+1$ continuation PRP contracts. Pareto-optimality and the proof of Lemma 2 imply that these two PRP contracts are the termination contract with principal payoff V_{t+1} and the Pareto-optimal PRP contract over the remaining $K-(t+1)$ dates with principal payoff V_{t+1} (which may also be the termination contract). Pareto-optimality and the PRP condition imply that V_{t+1} depends only on W_t. So $V_{t+1} = V_{t+1}(W_t)$ and similarly $a_t^* = a_t^*(W_t)$, $w_t^* = w_t^*(W_t)$, and $\rho_{t+1}^* = \rho_{t+1}^*(W_t)$. The only part of the risk-neutral analysis that can't be imported is the surplus-based algorithm for computing a_t^*, w_t^*, and ρ_{t+1}^*.

The Pareto-optimal PRP contract over the remaining $K-(t+1)$ dates with principal payoff $V_{t+1}(W_t)$ implicitly defines via the date $t+1$ PRP Pareto-frontier a $W_{t+1}(W_t)$. $W_{t+1}(W_t)$ is the agent payoff under the date $t+1$ Pareto-optimal PRP allocation linked to the date t Pareto-optimal PRP allocation with agent payoff W_t. This concludes the proof of Theorem 1.

PRP and Three Notions of Renegotiation

In my moral hazard model, I commit to a contract mechanism but apply a renegotiation-proof condition, PRP, to the dynamic strategies of the principal and agent. The idea is that if at some point the continuation play is not weakly Pareto-optimal, then the principal and agent can privately coordinate on a strictly Pareto-improving continuation equilibrium. I think of such a coordinated adjustment in strategies as a kind of "private revision," hence the term private-revision-proof contract. The PRP condition is basically the renegotiation-proof condition of Benoît and Krishna (1993) applied to a special class of contract games involving only private strategies. Benoît and Krishna (1993) study the limiting properties of the renegotiation-proof equilibrium payoff set for general repeated games. In contrast, I am concerned with optimal mechanism design under the PRP condition. Due to the private nature of the contract games I consider, in my setting the need to focus on renegotiation-proof equilibria as opposed to perfect equilibria is particularly compelling. Arguments appealing to third party contracts or reputation-based market disciplining forces that may discourage renegotiation of public strategies or renegotiation of mechanisms do not apply to renegotiation of private strategies.

In contrast, the standard approach to renegotiation-proof contracting is to apply the condition to the contract mechanism at the beginning of each date rather than to the contract equilibrium. See, for example, Wang (2000) and Malcomson and Spinnewyn (1988). Finally, there is a notion of contract renegotiation between the time effort is exerted and when the outcome is revealed. Such a renegotiation procedure allows the initial contract offer to focus on providing incentives and the renegotiated contract to focus on sharing risk. This disentangling of incentives and risk is a fruitful way of generating intuitive, simple contracts. For example, Matthews (1995), Matthews (2001), and Dewatripont, Legros, and Matthews (2003) show how sales and debt arrangements can emerge as optimal initial contract offers. Hermalin and Katz (1991) show how such renegotiations can increase ex-ante efficiency if the effort is observable but not verifiable. However, if effort is unobservable, then renegotiation typically lowers ex-ante efficiency. Fudenberg and Tirole (1990) and Ma (1991) argue that it may be optimal for the agent to randomize effort, thereby creating asymmetric information at the renegotiation stage. In general, withholding information from the principal diminishes the scope for renegotiation, which lessens the ex-ante efficiency loss caused by renegotiation. This idea is relevant to the optimality of arm's length relationships, a topic I take up in the following section.

3 Arm's Length, Memoryless Contracts

This paper set out to demonstrate the optimality of arm's length, memoryless contracts. Proposition 1 and Theorem 1 showed how Pareto-optimal PRP contracts are composed of a series of memoryless, pass/fail short-term contracts. However, the model exogenously fixed the evaluation frequency. Thus, there was no notion of one contract being more or less arm's length than another contract. I now incorporate an arm's length dimension by considering a limiting continuous-time model where the information flow is captured by a stochastic process and the contract is allowed to stipulate how frequently the principal can sample the stochastic process.

In the model considered in the previous section, any increase in the informational content of the private signals would expand the PRP Pareto-frontier and could make both players better off. This demonstrates a pervasive intuition that the more information the principal can gain access to, the better he can overcome the agency problem, which ultimately benefits everyone. If one applies this intuition across time, then, barring any monitoring costs, the conclusion is that the principal ought to try and evaluate the agent as frequently as possible, ideally continuously. In reality, the costs associated with evaluating a worker would prevent continuous evaluation. However, as the hands-off management philosophy suggests, even if costs were not an issue, there are many situations where it seems a manager should refrain from acquiring as much information relevant to a worker's performance as possible. I now show that in a continuous-time version of my model where the principal could, in theory, costlessly and continuously sample the underlying informative stochastic process, he will not want to do so. Continuous sampling plus the temptation to privately revise strategies will cause incentives to break down completely. Instead, the principal and agent will want to agree on an *arm's length contract* that limits the frequency of evaluation. My analysis will produce a notion of optimal evaluation frequency. The resulting optimal arm's length contract is structurally isomorphic to a discrete-time optimal PRP contract except the discrete evaluation dates are now endogenously determined instead of exogenously given. This analysis will complete my stated goal of showing the optimality of arm's length, memoryless contracts.

For simplicity, I specialize to a Gaussian version of the benchmark risk-neutral model where ${}_t\rho_{t+1} \sim N(a_t, 1)$ and $K = \infty$. There is a natural embedding of this model in a continuous-time setting with Brownian motion $(Z_t, \mathcal{F}_t)$. In the embedding, the agent selects a continuous effort process $\{a_t\}$ that can change values only at integer dates. The effort process generates a Brownian motion with drift: $dP_t = a_t dt + dZ_t$. The principal samples P_t only at integer dates. For example, at date 2 the principal observes P_2. But since he has already observed P_1, observing P_2 is equivalent to observing $P_2 - P_1 =_d {}_1\rho_2$. The mechanism transfers payments from the principal to the agent via a continuous payment flow process $w_t dt$

that changes only at integer dates. The agent's utility is $\mathbf{E}[\int_0^\tau \beta^t(w_t - h(a_t))dt + \beta^\tau O_A(w_\tau)]$ and the principal's utility is $\mathbf{E}[\int_0^\tau \beta^t(dP_t - w_t dt) + \beta^\tau O_P(-w_\tau)]$. The termination time τ is an integer-valued stopping time.

By increasing, in lock-step, the frequency at which the agent can change his effort level, the principal can sample the Brownian motion, the payment level can be changed, and the contract can be terminated, the model approaches the continuous-time limit. In this limit, right before date $t + dt$, the principal observes ${}_t\rho_{t+dt}|a_t dt \sim N(a_t dt, dt)$. Factoring out a dt, it is equivalent to assume that the principal observes a signal $\sim N(a_t, \frac{1}{dt})$.

Lemma 4. *In the continuous-time limit, PRP contracts induce zero effort.*

Proof. See appendix. □

As the length of each segment goes to zero, each ${}_t\rho_{t+dt}$ becomes increasingly noisy, making it hard to provide incentives. There are two potential ways the principal might overcome this problem.

One way exploits the fact that normal random variables, no matter how noisy, have arbitrarily informative tails. So the principal can induce effort by setting the report threshold to be extremely low and punishing the agent extremely when ${}_t\rho_{t+dt}$ drops below the threshold. This arrangement is similar to the one proposed by MacLeod (2003). Unfortunately, this method requires an extremely large continuation surplus to be put at risk, violating the interim participation constraint. Therefore, it is not feasible in my setting.

The other method is to aggregate numerous noisy signals. Fix a nontrivial time interval and have the principal's message space be a singleton that leads to retention for all dates before the last date of the interval. Then at the very last date, the principal can send one of two messages, leading either to retention or termination, with the report strategy being that the principal picks a threshold and reports the message that leads to termination if and only if the sum of the accrued signals is at or below the threshold. This arrangement is similar to the one proposed in Radner (1985). It is incentive-compatible in my setting and does not require a large continuation surplus to be put at risk. However, it is not PRP. By Lemma 2, at the final agent decision node of the interval it will be strictly Pareto-dominant to privately revise away any dependence of the principal's final report on anything but the last informative signal. The agent, expecting this, will choose zero effort except for the last instant.

The proof of Lemma 4 is based on the idea that, as the evaluation time interval shrinks, events in the time interval that provide the principal with sufficient confidence that the agent is putting in effort or, conversely, that the agent is shirking, become vanishingly rare. Thus, Lemma 4 is not special to models with risk-neutral players or normally distributed information. Also, there is nothing special about continuous-time. The proof of Lemma 4 more generally makes the point that, as evaluation frequency increases, eventually, incentives and effort start declining.

With PRP contracts, the basic impediment to inducing effort is that the principal and agent have too many opportunities to privately revise their equilibrium play. Thus, to be able to induce effort, the contract must somehow limit these opportunities. The most direct way would be to write a clause in the contract that restricts the times when the parties can privately revise. But this is tantamount to restricting private communication and, as I have argued before, is hard to enforce in practice. A more realistic approach is to limit the flow of information to the principal so that the set of times when the principal and agent can make Pareto-improving private revisions is *endogenously* limited. I take this approach by allowing the contract to stipulate when the principal can sample the Brownian motion.

Definition. *An arm's length contract is a contract that specifies a discrete set of random evaluation times $0 < t_1 < t_2 < \dots$ at which the principal can sample P_t and make a report. For each i, t_i is measurable with respect to the σ-algebra generated by $\{P_1, \dots P_{i-1}\}$.*

Note, it is not enough to simply restrict the times when the principal can make a termination decision. There is nothing in the definition of PRP contracts that prevents the principal and agent from writing a contract where the message space is trivial except for discrete moments. The negative result of Lemma 4 applies to such contracts.

It is also not enough to restrict the principal to evaluate the agent at discrete dates but then assume that when he does conduct an evaluation he can recover all the detailed, instant-by-instant performance data. That is, suppose the principal only evaluates the agent at discrete dates $0 < t_1 < t_2 \dots$ but assume that when the principal evaluates at date t_k, he observes the entire process P_t from t_{k-1} to t_k. Abreu et al (1991) points out how this method of revealing information only in discrete blocks can improve efficiency.[10] However, staggered but full information revelation will not work in my setting. While technically not covered by Lemma 4, the proposed contractual setting still features the same negative result. The problem is, once again, renegotiation. Within an evaluation interval, the PRP condition wipes out effort incentives at all instances except for the last one.

Thus, what's truly important for preserving incentives is only providing the principal the type of aggregate information which he cannot decompose into information about instantaneous effort choices.

With the introduction of arm's length contracts, there is now a notion of optimal evaluation frequency. If evaluations are too frequent then the PRP condition unravels any incentives; if evaluations are not frequent enough then discounting and increasing variance erode any incentives.

Proposition 2. *There exist constants S^* and Δ^* such that all Pareto-optimal arm's length PRP contracts have surplus S^* and share the following memoryless structure: The principal evaluates the agent at $k\Delta^*$ for $k \in \mathbb{Z}^+$. In every evaluation period, the agent puts in the same effort process $a(S^*, \Delta^*)$. At date $k\Delta^*$ for $k \in \mathbb{Z}^+$, if $P_{k\Delta^*} - P_{(k-1)\Delta^*} \geq \rho(S^*, \Delta^*)$ then the agent is retained. Otherwise the agent is terminated.*

Proof. See appendix. □

[10] Rahman (2014) demonstrates the power of this approach in a repeated game model with frequent actions where the players see the noisy public information in real time. He shows that if a disinterested mediator can make secret recommendations to the players and periodically reveal to the public those recommendations, then the real-time public information is effectively encrypted and the players are able to interpret the information only periodically. Once again, efficiency increases and, in fact, a folk theorem can be proved.

Notice Pareto-optimal arm's length PRP contracts are structurally isomorphic to the discrete-time PRP contracts from before except instead of exogenously fixed evaluation dates, the evaluation frequency is now endogenously determined.

Proposition 2 shows that by withholding information from the principal through contractual restrictions on when he is allowed to evaluate the agent, the scope for renegotiation is reduced and the value of the relationship increases. Of course, if the principal can secretly monitor the agent at dates not stipulated by the arm's length contract then the contract is undermined and once again incentives break down. In practice, one way for a manager to commit not to monitor when he is not supposed to is to make monitoring costly. Thus, another way to interpret the zero effort result of Lemma 4 is that it highlights the value of creating monitoring costs.

Monitoring costs can be generated by introducing a bureaucracy, increasing the manager's span of control, or overloading the manager with other responsibilities. This has implications for the theory of the firm. When designing the organizational hierarchy, ownership should be cognizant of how the monitoring costs generated by the hierarchy affect the ability to induce effort. Make the monitoring costs too high and individually rational contracts would not be able to stipulate any monitoring, leading to zero effort. On the other hand, make the monitoring costs too low and the private-revision temptation would lead to rampant over-monitoring and, again, zero effort. The optimal hierarchy should lead to monitoring costs that balance these two opposing forces. The idea is reminiscent of a point made in Aghion and Tirole (1997). In that paper, one of many projects needs to be enacted. The principal and agent do not have perfectly aligned preferences over projects and both can exert effort to get information about project payoffs. There, the idea is that exogenously increasing the principal's effort cost function essentially commits the principal to put in less effort which, in turn, induces more effort from the agent. In a world where the principal retains formal authority over project selection but depends on the advice of the agent, this tradeoff can be beneficial for the principal.

The idea of limiting the information possessed by the principal in order to prevent inefficient renegotiation is also a theme of the renegotiation under asymmetric information literature. See Dewatripont and Maskin (1990) for an excellent summary. Dewatripont (1989) and Laffont and Tirole (1990) both argue it may be optimal for agent types to pool, thereby making a type-dependent ex-post efficient renegotiation difficult. In my paper, there is no adverse selection on the agent's side. Moreover, in the arm's length contracting setting, the harmful renegotiation is prevented not because the principal no longer knows when to renegotiate. Rather, the principal no longer observes the signal he needs to renegotiate. This approach to limiting renegotiation by limiting observability is similar to Dewatripont and Maskin (1989) and Crémer (1995).

4 Conclusion

In many employment settings, the worker has very little direct influence on objective performance measures of the firm. In these situations, the best way to induce hidden effort is for a manager to evaluate and form private judgements of the worker and then make reports. My paper made a simple point about this type of contracting situation: Since the worker and manager are effectively playing a game with private strategies, it is reasonable to assume that their equilibrium play is renegotiation-proof.

Under this assumption, I found that a Pareto-optimal long-term contract reduces to a sequence of arm's length, memoryless shorter-term contracts. Moreover, at each evaluation date, the manager uses a simple pass/fail rule to determine whether to retain the worker. This allowed me to provide a rationale for many features of real-life employment contracts that were not easily explained by the literature. I showed that the detail-free nature of the optimal contracts is robust to risk-averse players and does not depend on any strong assumptions about the underlying uncertainty.

Given the simplicity and robustness of the results, there are numerous fruitful extensions to explore. One weakness of this model and most contracting models is the lack of private savings. Introducing private savings should not compromise the intuition for renegotiation-proof strategies. Might the simple structure survive and provide a tractable framework to study optimal savings behavior in the presence of agency conflicts? Another natural extension is to introduce the possibility of promotion. If promotion can increase the surplus at stake in the relationship then incorporating a promotion clause into the current job's contract can improve efficiency. The threshold structure of optimal contracts may supply a rationale for the standard threshold-based promotion schemes up-or-stay and up-or-out.

Appendix A: Private-Revision-Proof Equilibrium

Call a strategy profile $\{a, r\}$ *natural* if every r_{t+1} depends on H_{t+1} only up to h_t and ${}_t\rho_{t+1}$. Notice, in a natural equilibrium, the principal's computed value $W_t(H_t \cup r_t)$ of the agent's continuation payoff at the beginning of the date t segment is measurable with respect to h_t and agrees with agent's own computation. That is, the agent's continuation payoff is common knowledge. And since we already know the principal's continuation payoff is common knowledge, this means that if attention is restricted only to the set of natural equilibria, there is no ambiguity in the definition of PRP. From now on, a natural equilibrium will be called natural-PRP if it satisfies the PRP definition applied over the set of all natural equilibria.

Now consider an arbitrary equilibrium that is natural starting from the date $t+1$ segment. That is, r_{s+1} depends on H_{s+1} only up to h_s and ${}_s\rho_{s+1}$ for all $s \geq t+1$. The continuation payoff process is then common knowledge starting from the date $t+1$ segment and it is well-defined to talk about whether the equilibrium is natural-PRP starting at the date $t+1$ segment. Assume that the equilibrium is natural-PRP starting at the date $t+1$ segment. Fix a history h_t and consider this equilibrium at the beginning of the date t segment. Here, the principal's computed value of the agent's continuation payoff need not be common knowledge because r_{t+1} may depend on aspects of the principal's private history that are not captured by h_t nor ${}_t\rho_{t+1}$. Let $a_t(h_t)$ be the induced agent effort and let $W_t(h_t)$ be the agent's computed value of his own continuation payoff.

I now describe a way to revise the principal's report strategy r_{t+1} in a way so that, combined with the agent's revised best-response effort, the equilibrium becomes natural-PRP starting from the date t segment. Moreover, unless the revision is trivial, in the revised equilibrium, the agent's common knowledge continuation payoff will be strictly greater than his computed value of his own continuation payoff under the original equilibrium. The principal's common knowledge continuation payoff will also strictly increase.

With this result and the fact that any equilibrium is natural-PRP starting at the last segment, it is natural to define PRP in the following way. An equilibrium is PRP if and only if it is a natural equilibrium that is natural-PRP.

I now describe the aforementioned revision.

Again, fix an h_t and consider all the possible common knowledge agent continuation payoffs $W_{t+1}(h_{t+1})$

that could occur in the beginning of the date $t+1$ segment. Typically, if one changes the principal's report strategy $r_{t+1}(h_t)$ then this will affect the agent's computed value of his own continuation payoff at the beginning of the date $t+1$ segment. That is, changes to $r_{t+1}(h_t)$ will affect the values $W_{t+1}(h_{t+1})$ for each h_{t+1}. However, due the fact that the equilibrium is natural starting from the date $t+1$ segment, each $W_{t+1}(h_{t+1})$ is actually unaffected by changes to $r_{t+1}(h_t)$.

For all $s < K$, define $W_s(h_s)^+$ to be the agent's date s ex-post continuation payoff after he exerts his effort $a_s(h_s)$. Replace the principal's report strategy $r_{t+1}(h_t)$ with the unique min/max threshold report strategy that delivers the same $W_t(h_t)^+$ to the agent under the original effort choice $a_t(h_t)$. Here, a min/max threshold report strategy is one where there is a threshold ρ_{t+1} such that if ${}_t\rho_{t+1} \leq \rho_{t+1}$, then the principal reports the message that minimizes $W_{t+1}(h_{t+1})$. Otherwise, he reports the message the maximizes $W_{t+1}(h_{t+1})$. Notice, here I used the property that the values $W_{t+1}(h_{t+1})$ are unaffected when the principal's report strategy is changed from $r_{t+1}(h_t)$ to something else.

Under this new report strategy, the agent is at least as well of as before at the beginning of the date t segment because he can always guarantee himself his original continuation payoff by choosing the original effort $a_t(h_t)$. Let $\hat{a}_t(h_t)$ be the new induced effort. Now, first suppose that $a_t(h_t) > 0$. Then the MLRP property implies that for all $a < a_t(h_t)$ ($a > a_t(h_t)$), $W_t(h_t)^+$ is strictly smaller (greater) under the new report strategy than under the original report strategy. Thus, $\hat{a}_t(h_t) \geq a_t(h_t)$. Moreover, since effort cost $h(\cdot)$ and $W_t(h_t)^+$ under any report strategy are differentiable with respect to effort, a necessary condition for optimality is the first-order condition. The MLRP condition implies that the slope of $W_t(h_t)^+$ with respect to effort at $a_t(h_t)$ strictly increases when moving from the original report strategy to the new report strategy. Since the first-order condition held at $a_t(h_t)$ under the original report strategy, it must not hold under the new report strategy. Thus, $\hat{a}_t(h_t) \neq a_t(h_t)$. Since I already showed $\hat{a}_t(h_t) \geq a_t(h_t)$, this must mean $\hat{a}_t(h_t) > a_t(h_t)$. Thus, $\hat{W}_t(h_t) > W_t(h_t)$ and $\hat{V}_t(h_t) > V_t(h_t)$. Lastly, suppose $a_t(h_t) = 0$. Under the min/max threshold report strategy, $W_t(h_t)^+$ as a function of effort is strictly increasing with strictly positive slope. Since $h'(0) = 0$, it must be that $\hat{a}_t(h_t) > 0 = a_t(h_t)$, and again $\hat{W}_t(h_t) > W_t(h_t)$ and $\hat{V}_t(h_t) > V_t(h_t)$.

This completes the proof that the proposed revision is strictly Pareto-improving. The equilibrium is now clearly natural-PRP starting at the date t segment.

Appendix B: Proofs

Proof of Lemma 1. Suppose H_{t+1} and H'_{t+1} share the same public history h_t but $V_{t+1}(H'_{t+1}) > V_{t+1}(H_{t+1})$. Then starting from decision node H_{t+1}, the principal can achieve payoff $V_{t+1}(H'_{t+1})$ by simply playing his continuation report strategy as if he were at H'_{t+1}. This argument goes through because the principal's preferences are independent of his private information (e.g. ${}_t\rho_{t+1}$). There is no single-crossing property. □

Proof of Lemma 2. This was proved as part of the development of the definition of PRP. See Appendix A. □

Proof of Proposition 1. Introduce a weaker notion of PRP called one-shot-PRP: A natural equilibrium (see definition in Appendix A) is one-shot-PRP at the beginning of the date t segment if it is one-shot-PRP at the beginning of the date $t+1$ segment, and there does not exist a natural equilibria that has the same continuation play at the beginning of the date $t+1$ segment, and is strictly Pareto-dominant at the beginning of the date t segment. A contract is one-shot-PRP if its natural equilibrium is one-shot-PRP at the beginning of the initial segment.

One-shot-PRP is essentially a weak renegotiation-proof condition that requires that players do not have an incentive to renegotiate their action choices in the current segment while keeping all future actions unchanged. In the 1-model, all contracts are Pareto-optimal one-shot-PRP. Suppose Pareto-optimal one-shot-PRP contracts have already been characterized for the $K-1$-model. Consider a Pareto-optimal one-shot-PRP contract in the K-model. The logic of Lemma 2 applies to one-shot-PRP contracts and implies that at the beginning of the second segment, there are at most two possible one-shot-PRP contracts of the $K-1$-model that are enacted. Pareto-optimality and the proof of Lemma 2 imply that one contract is the termination contract with payoff $(W_1 = O_A + O_P - V_1, V_1)$ and the other is the Pareto-optimal one-shot-PRP contract of the $(K-1)$-model with payoff $(W_1 = S_1^* + O_A + O_P - V_1, V_1)$. Here, V_1 is indeterminate and S_1^* denotes the surplus shared by all Pareto-optimal one-shot-PRP contracts of the $K-1$-model. Pareto-optimality also uniquely determines the min/max report threshold ρ_1^* and effort a_0^*. The induction is complete and I have now shown that Pareto-optimal one-shot PRP contracts satisfy the memoryless, recursive structure described in Proposition 1.

By definition, Pareto-optimal one-shot-PRP contracts are weakly better than Pareto-optimal PRP contracts. I now show that Pareto-optimal one-shot-PRP contracts are actually PRP. This proves the proposition. Induct on K. Clearly, the claim is true for $K = 1$. Now suppose it is true for all $K-1$-models and consider a Pareto-optimal one-shot-PRP contract in the K-model. Suppose it is not PRP. Then, at the beginning of some date t segment, after some history h_t, there is a profitable renegotiation of the continuation equilibrium play. h_t cannot be h_0, since this would imply that the Pareto-optimal one-shot-PRP contract is strictly Pareto-dominated by a PRP contract, contradiction. h_t also cannot be of the form $\{pass, pass, \ldots pass, fail\}$ because these histories lead to termination which is automatically renegotiation-proof since no actions are involved. So h_t must be of the form $\{pass, pass \ldots pass\}$. But, by the characterization of Pareto-optimal one-shot-PRP contracts, we know that the continuation contract is one-shot-PRP over the remaining segments and by induction, we know that any such contract is PRP. Contradiction. □

Proof of Lemma 4. Let $a(dt)$ denote the stationary effort induced in the infinite-horizon model with time length dt. Let $\rho(dt)$ denote the associated threshold. Let $F(\cdot, \mu, \sigma^2)$ denote the cdf of a normal random variable with mean μ and variance σ^2. Consider the limit

$$\lambda := \lim_{dt \to 0} F\left(\rho(dt), a(dt), \frac{1}{dt}\right) / dt$$

Since $F\left(\rho(dt), a(dt), \frac{1}{dt}\right)$ is the conditional probability of termination at each date, λ is the limiting rate of termination. If $\lambda = \infty$, then there is almost surely immediate termination and the result is trivially true. So suppose $\lambda < \infty$. I now show that for any $k \in (-\infty, \infty)$,

$$\lim_{dt \to 0} F\left(\rho(dt), a(dt) + k, \frac{1}{dt}\right) / dt = \lambda$$

This result means that in the limit, the agent faces the same rate of termination no matter what effort he chooses. This means the agent faces no incentives when $dt \to 0$ and therefore $\lim_{dt \to 0} a(dt) = 0$.

To prove the claim, I normalize the problem and show that

$$\lim_{dt \to 0} F\left(\rho(dt), 0, \frac{1}{dt}\right) / dt = \lim_{dt \to 0} F\left(\rho(dt), k, \frac{1}{dt}\right) / dt \tag{2}$$

for all $k > 0$. If the LHS of (2) is 0, then so is the RHS. So from now on, assume the LHS of (2) > 0.

Together, a change-of-variable and the standard normal tail estimate imply

$$\lim_{dt \to 0} F\left(\rho(dt), 0, \frac{1}{dt}\right) \Bigg/ \frac{\phi(|\rho(dt)\sqrt{dt}|)}{|\rho(dt)\sqrt{dt}|} = \lim_{dt \to 0} F\left(\rho(dt)\sqrt{dt}, 0, 1\right) \Bigg/ \frac{\phi(|\rho(dt)\sqrt{dt}|)}{|\rho(dt)\sqrt{dt}|} = 1$$

Similarly,

$$\lim_{dt \to 0} F\left(\rho(dt), k, \frac{1}{dt}\right) \Bigg/ \frac{\phi(|(\rho(dt)-k)\sqrt{dt}|)}{|(\rho(dt)-k)\sqrt{dt}|} = \lim_{dt \to 0} F\left((\rho(dt)-k)\sqrt{dt}, 0, 1\right) \Bigg/ \frac{\phi(|(\rho(dt)-k)\sqrt{dt}|)}{|(\rho(dt)-k)\sqrt{dt}|} = 1$$

The tail estimate (i.e. the denominator) can be rewritten as

$$\frac{\phi(|\rho(dt)\sqrt{dt}|)}{|\rho(dt)\sqrt{dt}|} \cdot \frac{|\rho(dt)\sqrt{dt}|}{|(\rho(dt)-k)\sqrt{dt}|} \cdot e^{\rho(dt)kdt - k^2 dt/2}$$

Since $\rho(dt)$ is tending to $-\infty$, the middle fraction tends to 1. To prove (2), it suffices to show that $e^{\rho(dt)kdt-k^2dt/2}$ also tends to 1. For this, it suffices to show that $\lim_{dt\to 0}\rho(dt)(dt)^r = 0$ for all $r > \frac{1}{2}$. Suppose not, then there is a $\delta > 0$ such that $\lim_{dt\to 0} -\rho(dt)(dt)^{\frac{1}{2}+\delta} > 0$ which means for any $\varepsilon \in (0,\delta)$, $\lim_{dt\to 0} -\rho(dt)(dt)^{\frac{1}{2}+\varepsilon} = \infty$. This implies

$$\lim_{dt\to 0} F\left(\rho(dt)\sqrt{dt}, 0, 1\right) \Big/ dt \leq \lim_{dt\to 0} F\left(-(dt)^{-\varepsilon}, 0, 1\right) \Big/ dt \tag{3}$$

I now show the RHS of (3) is 0, which contradicts the assumption that the LHS of (2) is greater than 0, thereby proving the lemma.

The normal tail estimate implies

$$1 = \lim_{dt\to 0} \frac{\phi((dt)^{-\varepsilon})}{(dt)^{-\varepsilon}} \Big/ F\left(-(dt)^{\varepsilon}, 0, 1\right) = \lim_{dt\to 0} \frac{\phi((dt)^{-\varepsilon})}{(dt)^{1-\varepsilon}} \Big/ \frac{F\left(-(dt)^{\varepsilon}, 0, 1\right)}{dt} \tag{4}$$

L'Hopital's rule implies that $\lim_{x\to\infty} e^{-x}x^r = 0$ for all r. This implies that $\lim_{dt\to 0}\phi((dt)^{-\varepsilon})/(dt)^{1-\varepsilon} = 0$. (4) now implies that $F\left(-(dt)^{\varepsilon}, 0, 1\right)/dt = 0$. □

Proof of Proposition 2. Begin by assuming that at each evaluation date t_i, the agent can commit to an effort process until t_{i+1}. This assumption restricts the agent's decision nodes to be at 0 and the evaluation dates instead of at all times. As a result, the PRP condition needs to be checked only at date 0 and the evaluation dates. I will eventually verify that the Pareto-optimal PRP contracts I derive under this assumption are in fact PRP without the assumption.

Lemma 5. *Every arm's length Pareto-optimal PRP contract is payoff equivalent to one with non-random evaluation times.*

Proof. Lemma 2 and Pareto-optimality imply that there is a (constant) optimal evaluation date t_1 following which either a unique Pareto-optimal PRP contract is played or the agent is terminated depending on whether P_{t_1} exceeds a certain threshold. By recursion, the unique Pareto-optimal PRP contract enacted at t_1 will also feature an optimal first evaluation date t_2 that is, given the uniqueness of the contract enacted at t_1, also a constant. Similarly, each evaluation time t_i is a constant. □

Self-similarity of the infinite horizon model and Lemma 5 imply that the evaluation times are of the form $k\Delta$ for $k \in \mathbb{Z}^+$ and some constant $\Delta > 0$. Next, given a threshold ρ, maximal continuation surplus S, and evaluation time length Δ, the best response effort process of the agent in an evaluation period is

$$\{a_s(S,\Delta,\rho)\}_{s\in[0,\Delta)} = \underset{\{a_s\}_{s\in[0,\Delta)}}{\arg\max}\ \beta^{\Delta} S\left[1 - F\left(\rho, \int_0^{\Delta} a_s ds, \Delta\right)\right] - \int_0^{\Delta} \beta^s h(a_s) ds$$

where $F(\cdot,\mu,\sigma^2)$ is the cdf of a normal random variable with mean μ and variance σ^2. Next, given S and Δ, the PRP condition requires that the Pareto-optimal threshold satisfy

$$\begin{aligned}\rho(S,\Delta) := \underset{\rho}{\arg\max} \int_0^{\Delta} \beta^s(u(a_s(S,\Delta,\rho)) - h(a_s(S,\Delta,\rho)))ds + \\ \beta^{\Delta}\left(1 - F\left(\rho, \int_0^{\Delta} a_s(S,\Delta,\rho)ds, \Delta\right)\right) S\end{aligned}$$

Define $a(S,\Delta) := \{a_s(S,\Delta,\rho(S,\Delta))\}_{s\in[0,\Delta)}$. For each Δ, define $S(\Delta)$ as the solution to

$$\begin{aligned}S = \int_0^{\Delta} \beta^s(u(a_s(S,\Delta)) - h(a_s(S,\Delta)))ds + \beta^{\Delta}\left(1 - F\left(\rho(S,\Delta), \int_0^{\Delta} a_s(S,\Delta)ds, \Delta\right)\right) S + \\ (1-\beta^{\Delta})(O_A + O_P)\end{aligned}$$

Finally, $\Delta^* = \arg\max_\Delta S(\Delta)$ and $S^* = S(\Delta^*)$.

That last thing to do is to relax the assumption that the agent commits to an effort sequence at the beginning of every evaluation period and check that the contract is still PRP at all times, including in between evaluation dates. See Lemma 6 below. □

Lemma 6. *It is sufficient to check that an arm's length contract is PRP at evaluation dates.*

Proof. Without loss of generality, consider the initial evaluation period $[0, \Delta^*]$. From now on I will refer to $\rho(S^*, \Delta^*)$ as ρ, the action sequence $a(S^*, \Delta^*)$ as a and use a_s denote the element of a at date s. The agent's date t continuation payoff for any $t \in [0, \Delta^*)$ is

$$W_t = \beta^{\Delta^*-t}\left[W_{\tau=\Delta^*} + \left(1 - F\left(\rho, \int_0^{\Delta^*} a_s ds, \Delta^*\right)\right) S^*\right] - \int_t^{\Delta^*} \beta^{s-t} h(a_s) ds + \int_t^{\Delta^*} \beta^{s-t} w_s ds$$

Suppose at some interim date $\tilde{t} \in (0, \Delta^*)$, the contract is not private-revision proof. Let $\tilde{\rho}$ be a strictly Pareto-improving report threshold. Then the agent's revised action at date $s \in [\tilde{t}, \Delta^*)$ is

$$a'_s := a_{s-\tilde{t}}\left(S^*, \Delta^* - \tilde{t}, \tilde{\rho} - \int_0^{\tilde{t}} a_s ds\right)$$

The first-order conditions for effort imply that $e^{s-\tilde{t}} h'(a'_s) = h'(a'_{\tilde{t}})$ and $e^{s-\tilde{t}} h'(a_s) = h'(a_{\tilde{t}})$ for all $s \in [\tilde{t}, \Delta^*)$. This means either the entire sequence a' is higher than the entire sequence a over the interval $[\tilde{t}, \Delta^*)$ or vice versa. Since $\tilde{\rho}$ is a strict Pareto-improvement, it must be that $a'_s > a_s$ for all $s \in [\tilde{t}, \Delta^*)$, and also the agent's revised continuation payoff $W'_{\tilde{t}} > W_{\tilde{t}}$.

I now show that the contract isn't private-revision-proof at date 0 which is a contradiction. Suppose the principal privately revises to $\tilde{\rho}$ at date 0. The agent's revised action sequence is $a(S^*, \Delta^*, \tilde{\rho})$ which I will call $\tilde{a}$ for short. Let $\tilde{W}_0$ be the agent's new date 0 payoff. $\tilde{W}_0$ is strictly larger than his payoff under the action sequence $\{a_s\}_{s\in[0,\tilde{t})} \cup \{a'_s\}_{s\in[\tilde{t},\Delta^*)}$ which is

$$\beta^{\tilde{t}} W'_{\tilde{t}} - \int_0^{\tilde{t}} \beta^s h(a_s) ds + \int_0^{\tilde{t}} \beta^s w_s ds > \beta^{\tilde{t}} W_{\tilde{t}} - \int_0^{\tilde{t}} \beta^s h(a_s) ds + \int_0^{\tilde{t}} \beta^s w_s ds = W_0$$

To show that the principal is also better off under $\tilde{\rho}$, consider the function

$$f(x) := \left(1 - F\left(\tilde{\rho}, x + \int_0^{\Delta^*} \tilde{a}_s ds, \Delta^*\right)\right) S^*$$

defined over $\mathbb{R}$. f is an increasing logistic-shaped function with a convex lower half and a concave upper half. Since $\tilde{a}$ is a best-response to $\tilde{\rho}$, it must be that

$$f'(0) = \beta^{-\Delta^*+s} h'(\tilde{a}_s) \quad \forall s \in [0, \Delta^*) \tag{5}$$

and $f(0)$ is in the concave upper half.

The first-order conditions for effort imply that either $a \leq \tilde{a}$ or $a > \tilde{a}$. First suppose $a > \tilde{a}$. Now consider the scenario where the agent chooses to follow a up to date $\tilde{t}$ but then switches to $\tilde{a}$ starting from date $\tilde{t}$. At date $\tilde{t}$, the marginal benefit of the date s effort at level $\tilde{a}_s$ is $\beta^{\Delta^*-\tilde{t}} f'(\int_0^{\tilde{t}}(a_s - \tilde{a}_s)ds)ds$ for any $s \in [\tilde{t}, \Delta^*)$. The marginal cost is $\beta^{s-\tilde{t}} h'(\tilde{a}_s)ds$.

Since $\int_0^{\tilde{t}}(a_s - \tilde{a}_s)ds > 0$ and $f(0)$ is in the concave upper half, it must be that $f'(\int_0^{\tilde{t}}(a_s - \tilde{a}_s)ds) < f'(0)$. Combining this with (5) implies that under the proposed scenario, starting at date $\tilde{t}$, the marginal benefit of effort going forward is less than the marginal cost. This means that the optimal effort sequence starting from $\tilde{t}$ - which I have already previously defined as a' - must be smaller than $\tilde{a}$. By assumption, this implies that $a' < \tilde{a} < a$. Contradiction. So $\tilde{a} \geq a$. But $\tilde{a} \neq a$, so $\tilde{a} > a$ and the principal is also strictly better off under $\tilde{\rho}$ at date 0. □

References

[1] Abreu, D., P. Milgrom and D. Pearce (1991) "Information and Timing in Repeated Partnerships," *Econometrica* Vol. 59, pp. 1713-1733

[2] Baker, G. P., M. C. Jensen and K. J. Murphy (1988) "Compensation and Incentives: Practice vs. Theory," *The Journal of Finance* Vol. 43, pp. 593-616

[3] Benoît, J.-P. and V. Krishna (1993) "Renegotiation in Finitely Repeated Games," *Econometrica* Vol. 61, pp. 303-323

[4] Carroll, G. (2015) "Robustness and Linear Contracts," *American Economic Review* Vol. 105, pp. 536-563

[5] Chassang, S. (2013) "Calibrated Incentive Contracts," *Econometrica* Vol. 81, pp. 1935-1971

[6] Crémer, J. (1995) "Arm's Length Relationships," *Quarterly Journal of Economics* Vol. 110, pp. 275-295

[7] Dewatripont, M. (1989) "Renegotiation and Information Revelation Over Time: The Case of Optimal Labor Contracts," *Quarterly Journal of Economics* Vol. 104, pp. 589-619

[8] Dewatripont, M., P. Legros and S. A. Matthews (2003) "Moral Hazard and Capital Structure Dynamics," *Journal of the European Economic Association* Vol. 1, pp. 890 - 930

[9] Dewatripont, M. and E. Maskin (1990) "Contract Renegotiation in Models of Asymmetric Information," *European Economic Review* Vol. 34, pp. 311-321

[10] Dewatripont, M. and E. Maskin (1995) "Contractual Contingencies and Renegotiation," *RAND Journal of Economics* Vol. 26, pp. 704-719

[11] Dye, R. A. (1985) "Optimal Length of Labor Contracts," *International Economic Review* Vol. 26, pp. 251-270

[12] Fudenberg, D., B. Holmstrom and P. Milgrom (1990) "Short-Term Contracts and Long-Term Agency Relationships," *Journal of Economic Theory* Vol. 51, pp. 1-31

[13] Fudenberg, D. and J. Tirole (1990) "Moral Hazard and Renegotiation in Agency Contracts," *Econometrica* Vol. 58, pp. 1279-1319

[14] Fuchs, W. (2007) "Contracting with Repeated Moral Hazard and Private Evaluations," *American Economic Review* Vol. 97, pp. 1432-1448

[15] Garvin, D. A. (2013) "How Google Sold Its Engineers on Management," *Harvard Business Review* December 2013 Issue

[16] Harris, M. and B. Holmstrom (1987) "On the Duration of Agreements," *International Economic Review* Vol. 28, pp. 389-406

[17] Hermalin, B. E. and M. L. Katz (1991) "Moral Hazard and Verifiability: The Effects of Renegotiation in Agency," *Econometrica* Vol. 59, pp. 1735-1753

[18] Harter, J. K., F. L. Schmidt, S. Agarwal and S. K. Plowman (2013) "The Relationship Between Engagement at Work and Organizational Outcomes," *Gallup*

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

[20] Laffont, J.-J. and J. Tirole (1990) "Adverse Selection and Renegotiation in Procurement," *The Review of Economic Studies* Vol. 57, pp. 597-625

[21] Lam, C. F., D. S. DeRue, E. P. Karam and J. R. Hollenbeck (2011) "The Impact of Feedback Frequency on Learning and Task Performance: Challenging the ?More is Better? Assumption," *Organizational Behavior and Human Decision Processes* Vol. 116, pp. 217-228

[22] Levin, J. (2003) "Relational Incentive Contracts," *American Economic Review* Vol. 93, pp. 835-857

[23] Lurie, N. H. and J. M. Swaminathan (2009) "Is Timely Information Always Better? The Effect of Feedback Frequency on Decision Making," *Organizational Behavior and Human Decision Processes* Vol. 108, pp. 315-329

[24] Ma, C. (1991) "Adverse Selection in Dynamic Moral Hazard," *Quarterly Journal of Economics* Vol. 106, pp. 255-275

[25] MacLeod, W. B. (2003) "Optimal Contracting with Subjective Evaluation," *American Economic Review* Vol. 93, pp. 216-240

[26] Malcomson, J. M. and F. Spinnewyn (1988) "The Multiperiod Principal-Agent Problem," *Review of Economic Studies* Vol. 55, pp. 391-407

[27] Matthews, S. A. (1995) "Renegotiation of Sales Contracts," *Econometrica* Vol. 63, pp. 567-590

[28] Matthews, S. A. (2001) "Renegotiation of Moral Hazard Contracts under Limited Liability and Monotonicity, *Journal of Economic Theory* Vol. 97, pp. 1-29

[29] Morse, G. (2004) "Feedback Backlash," *Harvard Business Review* October 2004 Issue

[30] Posner, E., K. Eggleston and R. Zeckhauser (2000) "The Design and Interpretation of Contracts: Why Complexity Matters," *Northwestern University Law Review* Vol. 95, pp. 91-132

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

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

[33] Rahman, D. (2014) "Frequent Actions with Infrequent Coordination," working paper

[34] Sannikov, Y. and A. Skrzypacz (2007) "Impossibility of Collusion under Imperfect Monitoring with Flexible Production," *American Economic Review* Vol. 97, pp. 1794 - 1823

[35] Wang, C. (2000) "Renegotiation-Proof Dynamic Contracts with Private Information," *Review of Economic Dynamics* Vol. 3, pp. 396-422