Private Revisions of Contracts:

A Robustly Tractable Approach to Optimal Contracting

John Y. Zhu[1]

November 15, 2014

Abstract

I study optimal contracting in a general repeated moral hazard setting where the noisy signals of the agent's hidden efforts are either privately observed by the principal or observed by both agent and principal but not verifiable. In this setting, there is a subset of contract renegotiations that cannot be detected by the public. I call these renegotiations *private revisions* and I show that optimal private-revision-proof contracts are robustly tractable: Each date, if the signal exceeds some threshold then the agent is retained. Otherwise the agent is randomly terminated. This simple structure does not rely on any strong assumptions about the utility functions or the nature of the underlying uncertainty.

JEL Codes:

[1]zhuyiran@wharton.upenn.edu, The Wharton School, University of Pennsylvania. I thank Steve Matthews and Daniel Gottlieb for valuable comments.

1 Introduction

Think about a standard, dynamic moral hazard model where the agent exerts hidden effort, generating noisy informative signals, and the principal can pay the agent a contingent wage and possibly terminate him. Unless one assumes a special functional form for utility, or a simple signal distribution, or, more likely, some combination of both, the resulting optimal contract will simply not be tractable.

This is difficult to square with real-life contracts and makes integrating the theory with the rest of economics challenging. The lack of tractability means doing comparative statics can be challenging, particularly over two of the most important facets: how the agent trades consumption across states (risk-aversion) and across time (inter-temporal substitution). Moreover, the lack of tractability also means that it is typically difficult to introduce an optimal contracting component to a model of something else on which one suspects contracts have an important effect. And lastly, the lack of tractability is directly or indirectly the reason for a number of significant contracting puzzles: Why do contracts use promotion as an incentive device? What do contracts with savings look like? Why are contracts arms-length?

In this paper, I introduce a general, dynamic moral hazard framework to bridge the gap between theory and practice. At every date, the principal pays the agent some amount, the agent applies a hidden effort which generates a noisy informative signal, and the principal may choose to terminate the agent. Everything is governed by a long-term contract signed at date zero. Players have general recursive preferences and the underlying uncertainty satisfies a weak first-order stochastic dominance property. Thus, the model is in every way a standard, general, dynamic moral hazard model, except that I assume signals are either privately observed by the principal or observable but not verifiable. This is in contrast to the most common assumption that the signals are public - something like stock price or cars sold. But as the two public signals examples highlight, most people do not work in jobs with that kind of objective performance measure. The more prevalent setting is one where the performance measure is subjective. In such settings, the non-public signals approach of this paper is more relevant.

I then formulate and solve the optimal contracting problem. I find that optimal contracts are tractable across states and time. At each date, if the signal is above a threshold then the agent is retained, otherwise he is randomly terminated. If the agent is retained then he faces a threshold spot contract the next date as well. Moreover, the spot contract the agent faces each date is independent of the history leading up to that date. In particular, the threshold for date t is a constant of date t and optimal contracts are memoryless.

To understand the key insight of the paper that delivers robustly tractable optimal contracts, first think about contracting with public signals. Broadly speaking, there are two classes of contracts of interest: incentive-compatible contracts and the subset of renegotiation-proof (incentive-compatible) contracts. Which class one formulates the optimal contracting problem over depends on one's assumptions about commitment. In this paper, I argue that, in the non-public setting, there is a natural assumption about the degree of commitment that is between full-commitment and no commitment:

In a public signals setting, any alteration of the contract is observable by outside parties.

Thus, whether the principal and agent distort the payment amount or change how the signals get mapped to decisions doesn't really make a difference - they are all just simply renegotiations. In the non-public signals setting, *there is a big difference* between distorting a decision into something that didn't exist before and changing how the signals get mapped to existing decisions. For example, suppose the contract originally stipulates that the agent receives \$1 for a "poor" or "fair" signal and \$2 for a "good" signal. Changing the good signal payment to \$4 is a verifiable alteration of the contract even if the signals themselves are not verifiable. In contrast, changing the payment mapping so that the contract pays \$1 for poor and \$2 for fair or good is a *private revision* which cannot be verified by the public. In this paper, I assume that the principal and agent can commit to not make verifiable alterations but cannot commit to abstain from private revisions. As a result, the optimal contracting problem is formulated over the set of private-revision-proof contracts, which lies between the incentive-compatible and renegotiation-proof contracts. The resulting optimal private-revision-proof contracts are my robustly tractable optimal contracts.

1.1 Related Literature

2 The Benchmark K-Model

In this section, I consider a benchmark risk-neutral model where the principal privately observes the noisy informative signals.

The complete dynamic model is a static segment repeated K times. I call it the K-model. A segment spans the beginning of date t to the beginning of date $t+1$ where $t = 0, 1, \ldots K-1$. The following figure shows a generic segment:

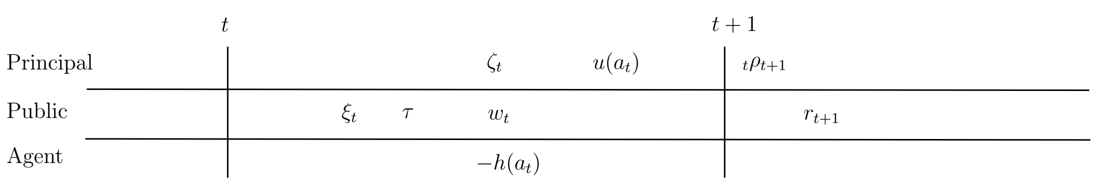

Figure 1: The three lanes that run across the figure denote what is observable to the principal only, to the public, and to the agent only.

At the beginning of date t, a public randomizing device ξ_t is realized and then a termination decision τ is made. If the agent is terminated then both parties exercise their outside options at date t with normalized values equal to zero. If the agent is retained then the agent is paid $w_t \in \mathbb{R}$ and chooses effort $a_t \in [0, \infty)$ with cost $h(a_t)$ satisfying $h'' > 0$ and $h'(0) = 0$. The principal observes a private randomizing device ζ_t and then receives a random utility with expectation $u(a_t)$ where u is a strictly increasing function. Next, at the beginning of date $t+1$, the principal observes a private informative signal ${}_t\rho_{t+1}$ with density $f(\cdot, a_t)$. I assume the distribution of ${}_t\rho_{t+1}$ has a common support across all effort levels and that $f(\cdot, a)$ first-order stochastic dominates $f(\cdot, a')$ whenever $a > a'$. The principal then makes a report r_{t+1} taking values in some sufficiently rich message space.

The date t termination decision, effort choice, and payment can depend on all prior public information: $\{r_i\}_{i=1}^t \cup \{\xi_i\}_{i=0}^t$. The date $t+1$ principal report can depend on all prior information private to the principal as well as the public randomizing devices: $\{\zeta_i, {}_i\rho_{i+1}, \xi_i\}_{i=0}^t$.

A contract specifies τ, and w_t, a_t, and r_{t+1} for $t = 0, 1, \ldots K-1$. A date s continuation contract specifies these objects for $t = s, s+1, \ldots K-1$. In general, a date s continuation contract is not a contract in the $(K-s)$-model in which the static segment is repeated $K-s$ times. This is because the principal's report strategy in a continuation contract can depend on his private randomizing devices and private informative signals that occur before the continuation contract. However, this mismatch can be easily fixed by introducing the notion of an *alternate version* of the $(K-s)$-model. In the alternate version, the initial private randomizing device is no longer just ζ_s but also includes all previous private randomizing devices and private informative signals of the K-model conditional on the public information at the beginning of the date s continuation contract.

Remark. *A date s continuation contract is a contract in its alternate version of the $(K-s)$-model.*

Given a contract, the total payoffs of the agent and principal are

$$W_0 = \mathbf{E}_{a,r}\left[\sum_{t=0}^{\tau-1} \beta^t (w_t - h(a_t))\right] \tag{1}$$

$$V_0 = \mathbf{E}_{a,r}\left[\sum_{t=0}^{\tau-1} \beta^t (-w_t + u(a_t))\right] \tag{2}$$

Furthermore, define the date t continuation payoffs (W_t, V_t) to be the total payoffs of the date t continuation contract viewed as a contract in its alternate version of the $(K-t)$-model. Thus, the totals payoffs of the agent and principal admit the following recursive representation:

$$W_0 = \mathbf{E}_{a,r}[w_0 - h(a_0) + \beta W_1] \tag{3}$$

$$V_0 = \mathbf{E}_{a,r}[-w_0 + u(a_0) + \beta V_1] \tag{4}$$

Finally, I assume that either player can choose to walk away from the contract. This imposes an interim participation constraint on both sides. Since outside options are normalized to zero, the interim constraint ensures that the continuation payoff process (W_t, V_t) is always nonnegative.

A contract is incentive-compatible if interim participation constraints are satisfied, the agent's effort is a best response to the principal's report strategy, and the following principal's truth-telling constraint is satisfied: given any public history h through the end of date $t-1$ and any two date t messages m and m', $V_t(h \cup m) = V_t(h \cup m')$. From now on all contracts are assumed to be incentive-compatible.

2.1 What are Private Revisions?

As an example, suppose the principal and agent initially agreed to a contract that contained the following date t continuation contract: At the beginning of date $t + 1$, the principal reports whether or not his private signal ${}_t\rho_{t+1}$ exceeds an *extremely high* threshold $\overline{\rho}$. If so, then a "*good*" continuation contract with payoff (W^g_{t+1}, V_{t+1}) is enacted. Otherwise a "*bad*" continuation contract with payoff (W^b_{t+1}, V_{t+1}) is enacted where $W^b_{t+1} < W^g_{t+1}$. Such a seemingly ex-post inefficient clause may appear in an optimal contract for the sake of ex-ante efficiency. It is precisely the type of unreasonable arrangement I want to eliminate endogenously.

Since the *good* payoff is extremely hard to achieve, this date t continuation contract provides poor incentives for effort a_t and would benefit from revision: The principal is better off lowering $\overline{\rho}$ and inducing a higher effort a_t. Of course, the agent will also be better off. Formally, changing the principal's report strategy is a renegotiation of the contract, just like a change in, say, how much w_t to pay the agent. In practice, however, these two types of renegotiations can be very different.

A change in w_t is verifiable by the public even if the signal that maps to w_t is not. A change in r_{t+1} need not be verifiable. The principal's report strategy depends in part on the principal's private history. Changes like the aforementioned lowering of $\overline{\rho}$ cannot be verified by the public. This is true of any *private revision* of r_{t+1} which leaves its dependence on the public history unchanged.[2]

In a public signals setting, there is no such thing as a private revision. Every revision is a verifiable contract renegotiation. Thus, as long as there is a third party such as a court that has perfect commitment and can enforce contracts, the principal and agent have no opportunity to change the contract after it is signed. Therefore, the space of relevant contracts is all incentive-compatible contracts. However, when signals are not public, even if there is a perfect court, the principal and agent still have a little room to change the contract in ways that remain under the court's radar. These changes are precisely the private revisions. Thus, the natural space of contracts in the non-public setting is not the entire set of incentive-compatible contracts, but rather the subset of private-revision-proof contracts.

I now recursively define private-revision-proof contracts. I first define private-revision-proof contracts for models of length one (the reduced-form model) and then inductively define private-revision-proof contracts for the K-model.

2.2 The Reduced-Form Model

The reduced-form model is the first segment of the K-model with continuation payoffs $(W_1, V_1) \geq 0$ appended to the end. Unlike in the K-model where (W_1, V_1) are quantities endogenous to the choice of contract, here (W_1, V_1) are simply exogenous functions of the

[2]Not every change in r_{t+1} is a private revision. For example, suppose there is a third message, *okay*, which, given the public history leading up to the date t continuation contract, is never reported. Then if the principal changes his strategy and reports *okay*, his revision is clearly not private.

public information available at the end of the first segment: ξ_0 and r_1. In the reduced-form model, the agent's and principal's utilities naturally take the forms of (3) and (4) respectively.

Without loss of generality, I focus on contracts that do not involve any chance of initial termination. A contract in the reduced-form model specifies a contract game $G = \{w_0\}$ and a Nash equilibrium $E = \{a_0, r_1\}$. On the principal's side, the truth-telling constraint implies that the principal's continuation payoff is constant over all equilibrium path messages. Without loss of generality, I may assume that the entire message space is on the equilibrium path and that V_1 is constant over the entire space.

Fix a contract $\{G, E = (a_0, r_1)\}$. Notice E is not the only equilibrium of G. Pick any other report strategy $\tilde{r}$ and let $\tilde{a}$ be the best response. Then $\{\tilde{a}, \tilde{r}\}$ is also an equilibrium of G. In this paper, I restrict attention to a subset of contracts by introducing an equilibrium selection:

Definition. *A reduced-form contract $\{G, E\}$ is private-revision-proof if $E \in \mathcal{P}(G)$, the Pareto-frontier of the equilibrium set of G.*

Suppose the principal and agent signed a contract $\{G, E = (a_0, r_1)\}$ that did not respect this selection: $E \notin \mathcal{P}(G)$. Let the principal choose any alternative $(\tilde{a}, \tilde{r}) \in \mathcal{P}(G)$ that Pareto dominates (a_0, r_1). Then before the agent acts, the principal can privately inform the agent that he plans on playing $\tilde{r}$ instead of r_1 and that therefore it is in the agent's best interest to play $\tilde{a}$ instead of a_0. I call such a change a **private revision**. In the previous subsection, I argued that there is scope for such private revisions. Anticipating such credible private revisions, the agent, ex-ante, demands a private-revision-proof contract.

The key property of private-revision-proof (PRP) contracts in the reduced-form model is that the principal uses a simple threshold report strategy:

Lemma 1. *Let $(a_0, r_1) \in \mathcal{P}(G)$. There is a threshold ${}_0\rho_1^*$ such that if ${}_0\rho_1 \geq {}_0\rho_1^*$ the principal reports the message that maximizes W_1. Otherwise the principal reports the message that minimizes W_1. Thus, it suffices to assume that the message space consists of only two messages: $\{good, bad\}$ with $W_1(good) \geq W_1(bad)$.*

This is easy to see. Suppose r_1 does not satisfy the threshold property. Then consider the following alternative report strategy $r_{improve}$: for all signals above ρ^* report the message that maximizes W_1 and below ρ^* report the message that minimizes W_1. Here ρ^* is selected so that under the measure generated by a_0, the agent's payoff under $r_{improve}$ is the same as under r_1. Let $a_{improve}$ be the agent's best response to $r_{improve}$. Clearly, the agent is better off in the equilibrium $\{a_{improve}, r_{improve}\}$. Moreover, due to the FOSD property of ${}_0\rho_1$, the agent's benefit as a function of the effort exerted is steeper than under r_1. This means that the best response effort $a_{improve} > a_0$, which means the principal is also better off.

Remark. *The principal does not use his private randomizing device ζ_0.*

A much stronger version of this important property holds once I move beyond the reduced-form model and introduce private-revision-proofness in the full K-model. Any dependence of the principal's report strategy r_{t+1} on any private component of the principal's history other than ${}_t\rho_{t+1}$ will be privately revised away before the agent chooses a_t. This is because any such dependence is not influenced by a_t, and will drag it down. Thus, private-revision-proofness severely limits the scope of history dependence. As a result, optimal contracts are forward rather than backward looking.

2.3 PRP Contracts in the K-model

Given the definition of PRP contracts in the reduced-form model and given the recursive structure of contracts in the K-model, the definition of PRP contracts in the K-model is naturally:

Definition. *A K-model contract is PRP if it is PRP as a reduced-form contract and if each continuation contract is PRP in the $(K-1)$-model.*

This definition implicitly assumes that a continuation contract is a contract in the $(K-1)$-model. But recall that such a continuation contract is only a contract in its alternate version of the $(K-1)$-model. Thus, technically speaking, the definition is not well-defined. However, the different versions only differ by the initial private randomizing devices which, as noted in the previous remark, PRP contracts ignore. Thus, with respect to PRP contracts, I can speak of models without reference to which version it is.

From now on contract will mean PRP contract. I now characterize the Pareto-optimal contracts. Since there is risk neutrality and equal discounting, the timing of pay is partially irrelevant and so I will just characterize one particular implementation of the agent-optimal contract that gives all of the Pareto-optimal surplus to the agent. Any other Pareto-optimal allocation can be achieved by simply taking an agent-optimal contract and shifting down the initial transfer to the agent. Of course, the pay structure of the particular contract I characterize may not be the robust one with respect to certain perturbations of the model parameters (e.g. making the agent slightly more impatient than the principal). However, the optimal effort levels and report thresholds are independent of the choice of implementation or allocation, as is the recursive structure of the contract. More importantly, I have chosen to start with the risk-neutral, equal discounting case mostly because of the ease of exposition. I will eventually extend the analysis to more general utility functions where the timing of pay matters.

Let S_t^* denote the surplus of Pareto-optimal contracts in the $(K-t)$-model.

Consider the agent-optimal contract with payoff $(W_0 = S_0^*, V_0 = 0)$. Either it is trivial (i.e. immediate termination) or it does not involve any chance of immediate termination. Assume the latter. Then let $(W_1(good), V_1(good))$ be the total payoff of the *good*-continuation contract and let $(W_1(bad), V_1(bad))$ be the total payoff of the *bad*-continuation contract. The truth-telling constraint implies $V_1(good) = V_1(bad) = V_1$. Moreover, due to the partial irrelevance of the timing of pay, it is without loss of generality to assume that $V_1 = 0$.

The total surplus S_0^*, which is also the agent's total payoff W_0, can be computed using a simple formula:

$$W_0 = S_0^* = -h(a(\Delta_1^*)) + u(a(\Delta_1^*)) + \beta(W_1(good) - F(\rho(\Delta_1^*), a(\Delta_1^*))\Delta_1^*) \qquad (5)$$

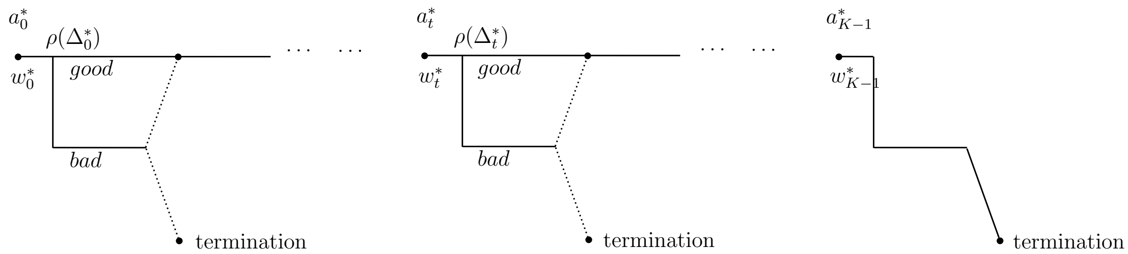

Figure 2

where $\Delta_1^* := W_1(good) - W_1(bad)$ is the date 1 "bonus" for good performance that provides the agent with incentives to exert effort a_0.

Here, $a(\Delta_1^*)$ is the agent's date 0 effort choice and $\rho(\Delta_1^*)$ is the principal's date 1 report threshold. Both are completely determined by Δ_1^*. To see this, suppose at some arbitrary date the agent faces some threshold ρ and bonus Δ at the beginning of the next date. Then the agent's effort that date is completely determined: $a = a(\Delta, \rho)$. Then by the PRP condition, the principal will select the ρ that maximizes $a(\Delta, \rho)$. Thus, ultimately, $\rho = \rho(\Delta)$ and $a = a(\Delta, \rho(\Delta)) = a(\Delta)$.

Since I am characterizing a Pareto-optimal contract, S_0^* must be maximized. Equation (5) then implies $W_1(good) = S_1^*$ and

$$\Delta_1^* := \underset{\Delta \in [0, S_1^*]}{\arg\max} -h(a(\Delta)) + u(a(\Delta)) - \beta F(\rho(\Delta), a(\Delta))\Delta \tag{6}$$

Since $W_1(bad) = W_1(good) - \Delta_1^* \in [0, S_1^*]$ and $V_1(bad) = V_1(good) = 0$, the *bad*-continuation contract can be achieved as an appropriately calibrated randomization between termination and the *good*-continuation contract. The probability of termination should be Δ_1^*/S_1^*.

Thus I have shown that in the agent-optimal contract, the agent applies effort $a_0^* := a(\Delta_1^*)$. The principal then observes ${}_0\rho_1$. If it exceeds $\rho(\Delta_1^*)$ then the principal reports *good* and the agent is retained. Otherwise the principal reports bad and the agent is potentially terminated. If the agent is retained then the agent-optimal contract of the $(K-1)$-model is enacted. Of course, by recursion, the agent-optimal contract of the $(K-1)$-model starts the exact same way as the agent-optimal contract of the K-model. Thus, the agent-optimal contract has a simple recursive backbone: At any date t, the agent-optimal contract has the same recursive threshold structure as in the first segment. See Figure 2.

To fully characterize the agent-optimal contract, it remains to compute the constants Δ_{t+1}^* and S_{t+1}^* for $t = 1, 2 \ldots K-1$ and to compute the payments w_t^*. This can be done using backwards induction. In the base case, $S_K^* = 0$. In the inductive step, using the arguments made earlier, given S_{t+1}^*, $\Delta_{t+1}^* = \arg\max_{\Delta \in [0, S_{t+1}^*]} -h(a(\Delta)) + u(a(\Delta)) - \beta F(\rho(\Delta), a(\Delta))\Delta$ and $S_t^* = h(a(\Delta_{t+1}^*)) + u(a(\Delta_{t+1}^*)) + \beta(S_{t+1}^* - F(\rho(\Delta_{t+1}^*), a(\Delta_{t+1}^*))\Delta_{t+1}^*)$. To compute w_t^*, simply note that $S_t^* = W_t^* = w_t^* - h(a(\Delta_{t+1}^*)) + \beta(S_{t+1}^* - F(\rho(\Delta_{t+1}^*), a(\Delta_{t+1}^*))\Delta_t^*) \Rightarrow w_t^* = u(a(\Delta_{t+1}^*))$. The backwards induction is now complete and I can state the optimality theorem.

Theorem 1. *The agent-optimal contract has the following recursive structure:*

- *At date t, the agent puts in effort $a_t^* = a(\Delta_t^*)$ and is paid $u(a_t^*)$.*
- *At the beginning of date $t+1$, if ${}_t\rho_{t+1} \geq \rho(\Delta_t^*)$ then the principal reports good and the agent is retained. Otherwise the principal reports bad and the agent is terminated with probability Δ_{t+1}^*/S_{t+1}^*.*
- *If the agent is terminated, both parties exercise their outside options at date $t+1$.*
- *If the agent is retained then the agent-optimal contract starting at date $t+1$ is enacted.*

The optimality theorem implies that Pareto-optimal contracts can be easily re-written in a way that satisfies the full-review property, which requires distinct realizations of ${}_t\rho_{t+1}$ to lead to distinct reports r_{t+1} for all t. See Levin (2003). The full-review property implies that the principal does not keep any private information from the agent.[3] Moreover, Pareto-optimal contracts are memoryless: the induced action sequence and report threshold sequence are both deterministic. It turns out that when the evaluation period is "short," the memoryless property can be somewhat of a curse. This is explored in Section 5, where I take the model to continuous-time and talk about arms-length relationships. In Section 3, I characterize Pareto-optimal contracts under a general class of recursive preferences. Pareto-optimal contracts still satisfy the full-review property and are essentially memoryless.

3 Contracts under General Recursive Preferences

Definition. *Define W_t^+ to be the date t ex-post agent continuation payoff which is a function of the public history up through the date t principal report plus the public randomizing device ξ_t. In general, I will add a "+" superscript to denote the ex-post version of a contract parameter (e.g. a_t^+, w_t^+).*

Recall, W_t is a function of the public history up through the date t principal report. Thus, W_t is the expectation of W_t^+ with respect to the date t public randomizing device. In the risk-neutral analysis, following a *good* report, $W_t \equiv W_t^+$. However, following a *bad* report, $W_t^+ \in \{W_t(\textit{good}), 0\}$.

The reason I now introduce extra notation for ex-post values is because the analysis of Pareto-optimal contracts with general recursive preferences requires some switching back-and-forth between ex-ante and ex-post parameters, beginning with the definition of the recursive preferences themselves.

Define the agent's recursive preference to be:

$$W_t = f_A^{-1}\mathbf{E}(f_A(W_t^+))) \qquad\qquad W_t^+ = u_A(w_t^+, a_t^+) + \beta_A W_{t+1}$$

u_A is strictly increasing in the first argument and $u_A(w, \cdot)$ satisfies $u_A''(w, \cdot) < 0$ and $u_A'(w, 0) = 0$ for all w. f_A is a strictly increasing continuous function. $\beta_A > 0$ is the agent discount

[3] The set of PRP contracts is a strict subset of the set of contracts satisfying the full review property. In general, requiring only the full-review property leads to higher payoffs - the principal can commit to inefficient report strategies that may be value enhancing ex-ante.

factor. Note, if the effort component were removed, $u_A(w) := w^\rho$, and $f_A(W) := W^{\alpha/\rho}$ then the agent would have Epstein-Zin preferences over wage.

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

$$V_t = f_P^{-1}\mathbf{E}(f_P(V_t^+)) \qquad\qquad V_t^+ = u_P(-w_t^+, a_t^+) + \beta_P V_{t+1}$$

u_P is strictly increasing in both arguments. f_P is a strictly increasing continuous function. $\beta_A > 0$ is the principal discount factor.

There are two differences between the risk-neutral case and the recursive case that are worth pointing out. The more minor difference is that under general recursive preferences, public randomization plays a more important role. Without certain concavity restrictions on u_A, u_P, f_A, and f_P, the Pareto-frontier is a priori not public randomization proof. For some agent payoffs, the Pareto-optimal contract must be a public randomization of other Pareto-optimal contracts. This was not true in the risk-neutral case.[4]

The more significant departure from the risk-neutral case is that the parameter values of a Pareto-optimal contract now depend on which Pareto optimal allocation is chosen. In the risk neutral case, every Pareto optimal allocation can be achieved by taking the agent-optimal contract and simply increasing the initial payment w_0. This means that without loss of generality, one can assume that all parameter values of all Pareto-optimal contracts except w_0 are independent of the choice of initial allocation. With recursive preferences, parameter values will, in general, depend on the initial allocation.

Thus, the entire set of Pareto-optimal contracts does not possess the structural uniformity that is a hallmark of the risk-neutral case. However, on the individual level, a Pareto-optimal contract under recursive preferences still has virtually all of the structure of a risk-neutral Pareto-optimal contract: the principal still uses a simple threshold report strategy; the full-review property is still satisfied; and, modulo the public randomization history, all parameters are still memoryless except that $w_t^+(bad) < w_t^+(good)$: That is, the date t ex-post wage depends on the most recent report. Recall, in the risk-neutral case, $w_t^+(bad) = w_t^+(good)$. This bit of history dependence means that following a *bad* report, if the agent is not terminated, his continuation contract is not actually Pareto-optimal. But the inefficiency is only transitory: After the initial lower wage, the slate is wiped clean and the agent's prospects become identical to what he would've faced if the report were *good* and the Pareto-optimal continuation contract were enacted.

I now retrace the steps that led to Theorem 1. To ease the exposition, I will conduct the analysis with the assumption that all Pareto-optimal continuation contracts are a priori public randomization proof. Afterwards, I relax this assumption and explain how the statement of the results change.

First, the definition of PRP contracts in the reduced-form model is unchanged. Second, the proof of Lemma 1 still holds. This means that the spot contract at each date still exhibits the simple threshold structure with two messages. Moving to the K-model, the recursive definition of PRP contracts is unchanged. The key difference is that now the timing of pay

[4] With general recursive preferences, having a concave Pareto-frontier and being public randomization proof are no longer equivalent. This equivalence holds only when f_P and f_A are linear and the preferences are expected utility.

is relevant. This means that if one verifies that the agent-optimal contract satisfies some desirable structural properties, one can no longer conclude that all other Pareto-optimal contracts also satisfy these desirable properties. I now characterize the structure of a generic Pareto-optimal contract.

So fix a Pareto-optimal allocation (W_0, V_0). Let $(W_1(good), V_1(good))$ and $(W_1(bad), V_1(bad))$ be its 1 continuation payoffs. It is still true that $V_1(good) = V_1(bad) = V_1$. It is still true that the payoff $(W_1(good), V_1(good))$ is on the date 1 Pareto-frontier although I can no longer assume that it is the agent-optimal one. That is, I can no longer assume that $V_1 = 0$

The question is how to deliver the *bad*-continuation payoff $(W_1(bad), V_1(bad))$. The continuity of $f_A^{-1}\mathbf{E}f_A$ implies that there is some mixture mix_0 of the *good*-continuation contract and termination that will deliver a payoff $W_1(bad)$ to the agent. If $V_1 = 0$ then we are done, because mix_0 will also deliver payoff $V_1(bad) = V_1 = 0$ to the principal. However, in general, $V_1 > 0$, which means mix_0 will actually deliver a payoff $\hat{V} < V_1$ to the principal. To fix this, consider the alternate continuation contract where the initial transfer is lessened by some constant amount D where D is defined to be the amount so that the agent's continuation payoff is $W_1(bad)$. Call this contract the D-continuation contract. The principal's payoff under the D-continuation contract obviously exceeds V_1. The trivial mixture mix_D between the D-continuation contract and termination that puts all the weight on the D-continuation contract generates an agent payoff equal to $W_1(bad)$. mix_D generates a principal's payoff $> V_1$. Now consider the family of continuation contracts that are identical to the *good*-continuation contract except the initial transfer is decreased by an amount $d \in [0, D]$. Call such a continuation contract a d-continuation contract. Let mix_d denote the mixture between the d-continuation contract and termination that delivers $W_1(bad)$ to the agent. When $d = 0$, mix_d delivers principal payoff $< V_1$. When $d = D$, mix_d delivers principal payoff $> V_1$. So by continuity, there is a d^* such that mix_{d^*} delivers payoff V_1 to the principal.

Thus, to achieve the *bad*-continuation payoff $(W_1(good), V_1(good))$ simply randomize between the d^*-continuation contract and termination according to mix_{d^*}.

Theorem 2. *Pareto-optimal contracts are memoryless and satisfy the full-review property. The date t Pareto-optimal continuation contract with agent payoff W_t has the following recursive structure:*

- *The agent puts in effort $a_t(W_t)$ and is paid $w_t(W_t, good)$.*
- *At the beginning of date $t+1$, if ${}_t\rho_{t+1} \geq \rho_{t+1}(W_t)$ then the principal reports good and the agent is retained. Otherwise the principal reports bad and the agent is terminated with some probability $p_{t+1}(W_t)$.*
- *If the agent is terminated, both parties exercise their outside options at date t.*
- *If the agent is retained following a good report, the Pareto-optimal continuation contract with agent payoff $W_{t+1}(W_t)$ is enacted.*
- *If the agent is retained following a bad report, then the continuation contract enacted is the same as the good-continuation contract except the initial payment is $w_{t+1}(W_{t+1}, bad)$ which is lower than $w_{t+1}(W_{t+1}, good)$.*

Once the initial payoff W_0 is fixed, the entire payoff process W_t is known. Thus, a Pareto-optimal contract's continuation payoff process is memoryless. Similarly, the action process $a_t(\cdot)$, the *good*- and *bad*- payment processes $w_t(\cdot, good)$ and $w_t(\cdot, bad)$, and the threshold process $\rho_{t+1}(\cdot)$ are all memoryless.

If a Pareto-optimal continuation contract is a public randomization of other Pareto-optimal continuation contracts, then its parameters will of course depend on the realization of the public randomizing device. Thus, modulo the history of public randomizations, a Pareto-optimal contract's parameters are still memoryless. In particular, the full-review property still holds. Theorem 2 would have to be slightly modified in the following way: Instead of there being a function $W_{t+1}(W_t)$, there would be a public random function $W_{t+1}^+(W_t^+)$. For each realization of W_{t+1}^+, there would be a corresponding *bad* continuation contract which would be identical to the *good* continuation contract except the initial payment smaller by some amount d^*. This d^* can be chosen to be independent of the realization of the *bad* continuation contract. The date t ex-post action $a_t^+(\cdot)$, ex-post *good*- and *bad*- payments $w_t^+(\cdot, good)$ and $w_t^+(\cdot, bad)$, and the date $t+1$ ex-post threshold $\rho_{t+1}^+(\cdot)$ would depend on W_t^+.

4 Observable but Not Verifiable Signals

The characterization of Pareto-optimal contracts can be easily extended to a setting where the informative signals are observed by both the agent and the principal but is still not verifiable. This is the same level of observability assumed in the relational contracts literature.

Formally, each segment of the K-model still begins with a public randomizing device ξ_t followed by a termination decision τ. There is no need for a private randomizing device. The agent exerts hidden effort a_t with cost $h(a_t)$ and receives payment w_t. The principal receives utility $u(a_t)$. Then at the beginning of date $t+1$, the informative signal ${}_t\rho_{t+1}$ is observed by the principal and agent but not by the public.

A contract specifies τ, and w_t and a_t for $t = 0, 1, \ldots K-1$. These objects can depend on the history of public randomizing devices and non-public informative signals leading up through the beginning of date t. A date s continuation contract specifies these objects for $t = s, s+1, \ldots K-1$. A date s continuation contract is a contract is the $(K-s)$-model. An incentive-compatible contract is one where the agent's effort process a is a best response to τ and w.

Just like before I can define continuation payoffs W_t and V_t for all t which depend on both the public and non-public history leading up through the realization of the informative signal at the beginning of date t.

Assumption. *Both the principal and agent can freely dispose of output.*

Free disposal is typically applied to generate a monotonicity property in contracts. Innes (1990) assumes this for the principal and numerous papers assume this for the agent. This assumption implies that the principal and agent's payoffs are weakly positively correlated across histories holding the public component fixed: In particular, let h be a history leading up to date $t+1$ and let ${}_t\rho'_{t+1}$ and ${}_t\rho''_{t+1}$ be two informative signals with ${}_t\rho'_{t+1} >_t \rho''_{t+1}$. Then

$$W_{t+1}(h\ {}_t\rho'_{t+1}) \geq W_{t+1}(h\ {}_t\rho''_{t+1}) \text{ and } V_{t+1}(h\ {}_t\rho'_{t+1}) \geq V_{t+1}(h\ {}_t\rho''_{t+1}) \tag{7}$$

Moreover, if it were the case that inequalities were strict, then whenever ${}_t\rho''_{t+1}$ occurred, the principal and agent could always surreptitiously choose the $h\ {}_t\rho'_{t+1}$-continuation contract leading both players to be strictly better off. Thus, it must be that either $W_{t+1}(h\ {}_t\rho_{t+1})$ or $V_{t+1}(h\ {}_t\rho_{t+1})$ is constant over all ${}_t\rho_{t+1}$. If $W_{t+1}(h\ {}_t\rho_{t+1})$ is the one that is constant, it also doesn't hurt to simply let $V_{t+1}(h\ {}_t\rho_{t+1})$ be constant as well. Thus, without loss of generality, I can focus on contracts where $V_{t+1}(h\ {}_t\rho_{t+1})$ is constant over all ${}_t\rho_{t+1}$.

This creates a contracting space that is equivalent to the space of incentive-compatible contracts in the private signals setting. Thus, I can now further refine the space by imposing private-revision-proofness. All the results and proofs go through unchanged.

5 Arms-Length Contracting

For this section I specialize the discrete time model so that ${}_t\rho_{t+1} \sim N(1, a_t)$. There is a natural embedding of the discrete time model in a continuous-time Brownian framework running from 0 to K. In the Brownian framework there is a Brownian motion $(Z_t, \mathcal{F}_t)$ along with a continuum of public randomizing devices $(\xi_t, \mathcal{G}_t)$. In the embedding, the continuum of public randomizing devices is sampled at integer dates. The agent of the discrete time model selects a continuous effort process $\{a_t\}_{t\in[0,K]}$ but can only change his effort level at integer dates. The chosen effort process generates a Brownian motion with drift: $dP_t = a_t dt + dZ_t$. The principal of the discrete time model 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 principal pays the agent via a continuous payment process w_t which can only change at integer dates. The agent's utility is $\mathbf{E}[\int_0^\tau \beta^t(w_t - h(a_t))dt]$ and the principal's utility is $\mathbf{E}[\int_0^\tau \beta^t(u(a_t) - w_t)dt]$. Here τ is the integer valued stopping time that equals the date when the contract is terminated.

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, at date $t+dt$, the principal observes ${}_t\rho_{t+dt}|a_t dt \sim N(dt, a_t dt)$. Factoring out a dt, it is equivalent to assume that the principal observes a signal $\sim N(\frac{1}{dt}, a_t)$.

Thus, as the length of each segment goes to zero, each ${}_t\rho_{t+dt}$ becomes increasingly noisy, and it becomes increasingly harder to provide incentives. There are two potential ways the principal can 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, and so is not feasible in my setting.

The other method exploits the efficiency of cross-pledging. The principal does not evaluate the agent over an interval of time. During this interval, he accrues a flow of signals, and then at the end of the time interval, he makes a termination decision based on some aggregation of all of the accrued signals. This arrangement is similar to the one proposed by Fuchs (2007) and can be formally implemented by assuming that the principal's message space is a singleton which always leads to continuation for all dates before the last date of the interval. Then at the very last date, the principal is able to send one of a number of messages, some of which lead to continuation, and others to termination. This method of providing incentives does not require a large continuation surplus to be put at risk. Unfortunately, this method is not private-revision-proof. On the second to last date of the interval, the principal will privately revise his report strategy so that it only depends on the last signal. The agent, expecting this, will choose zero effort except for the last instance.

Proposition 1. *In continuous-time, contracts induce zero effort.*

Proof. See Appendix. □

Proposition 1 implies that to induce effort, one must somehow expand the contracting space. The question is how to expand the space in way that is economically reasonable and leads to nontrivial effort. There are two impediments to inducing effort in the continuous-time setting. One, the principal receives a constant flow of extremely noisy information. Two, the principal constantly uses the noisy information he just received to privately revise the contract.

Thus, any useful expansion of the contract space must target at least one of these impediments. Targeting the second impediment would be difficult in practice as contracts would need to effectively ban private communication between the principal and agent. Thus I will focus on limiting the flow of information. In the current setup, the principal constantly samples the Brownian motion. A very natural expansion of the contract space is to allow the contract to stipulate when the principal can sample the Brownian motion. In practice, this can be achieved by having pre-specified evaluation times such as annual performance reviews, having the principal be physically distant from the agent, increasing the the principal's responsibilities, and creating a culture of independence.

Definition. *An arms-length contract is a contract that specifies a finite set of random evaluation times $0 < t_1 < t_2 < \dots t_N = K$ on which the principal can sample P_t and make a report. Each $t_i \in \mathcal{F}_{t_{i-1}} \times \mathcal{G}_{t_{i-1}}$ and $t_0 := 0$.*

Note, even though the principal only samples at discrete dates in an arms-length contract, he is still able to constantly communicate with the agent. In particular, the principal can still constantly privately revise the contract. However,

Lemma 2. *It is sufficient to check that an arms-length contract is private-revision-proof at evaluation dates.*

Proof. See Appendix. □

This lemma is not as trivial as it sounds. While it is true that in between evaluation dates the principal is not receiving any new information with which to privately revise, the action set he can affect with private revision is changing. During an evaluation period, the agent is still constantly selecting his action process. The concern is that halfway through an evaluation period when the agent has already sunk some of his effort costs, the principal can privately revise the contract to induce a higher effort process for the rest of the evaluation period. Lemma 2 would be trivial only if for some reason the agent can only update his effort choice at evaluation dates.

In general, the evaluation times t_i of an arms-length contract can be random. However, Pareto-optimality and the PRP condition together imply that without loss of generality each evaluation time t_i is constant.

Lemma 3. *Every Pareto-optimal allocation can be achieved by an arms-length contract whose evaluation times are non-random.*

To see this, consider the evaluation time t_1. Define V_0^+ to be the principal's date 0 *ex-post* payoff, which is realized after the date 0 public randomizing determining t_1 is realized. Pick a realization t_1^* of t_1 that maximizes V_0^+. Then the contract with $t_1 \equiv t_1^*$ delivers a higher principal payoff. Thus Pareto-optimality implies it is without loss of generality to assume t_1 is constant.

Recycling the analysis of the discrete model, at date t_1, the principal will employ a threshold report strategy. Following the report the contract will either terminate or a Pareto-optimal continuation contract is enacted. Therefore, by recursion, t_2 and in general t_i are also deterministic.

Unlike in the discrete time case where for a single evaluation period the agent chooses a single effort, in arms-length contracts the agent chooses an effort sequence. However, the maximization problem for the optimal effort sequence is equivalent to the maximization problem for a single effort. To see this, suppose the agent is a facing an evaluation period of length k, a principal report threshold ρ, and a bonus Δ. Then the optimal effort sequence solves the following problem:

$$a(k,\Delta,\rho) := \{a_s(k,\Delta,\rho)\}_{s\in[0,k)} = \underset{\{a_s\}_{s\in[0,k)}}{\arg\max}\, \beta^k \Delta \left[1 - \Phi^k\left(\rho - \int_0^k a_s ds\right)\right] - \int_0^k \beta^s h(a_s) ds \tag{8}$$

Here Φ^k is the cdf of a normal random variable with mean 0 and variance k. Optimality implies that for all $s \in [0,k)$, $\beta^s h'(a_s(k,\Delta,\rho)) = h'(a_0(k,\Delta,\rho))$. Thus, once a_0 is chosen, all the other efforts are pinned down and the maximization problem reduces to one only over the initial effort level.

Lemma 4. *In arms-length contracts, the induced effort sequence in an evaluation period is fully determined by a single effort, say, the initial one. Let $a(k,\Delta,\rho)$ and $\hat{a}(\hat{k},\hat{\Delta},\hat{\rho})$ be two induced effort sequences. If $a_t(k,\Delta,\rho) > \hat{a}_t(\hat{k},\hat{\Delta},\hat{\rho})$ for some $t \in [0, k \wedge \hat{k})$, then $a_s(k,\Delta,\rho) > \hat{a}_s(\hat{k},\hat{\Delta},\hat{\rho})$ for all $s \in [0, k\wedge\hat{k})$.*

To characterize Pareto-optimal contracts, it suffices to find for each $t \in [0, K)$, the surplus $S^*(t)$ of a Pareto-optimal arms length contract that starts at date t. The function $S^*(t)$ satisfies the following equation:

$$S^*(t) = \max_{k \in (0, K-t],\ \Delta \in [0, S(t+k)]} u(a(k,\Delta)) - h(a(k,\Delta)) + \beta^k (S^*(t+k) - F(k, \rho(k,\Delta), a(k,\Delta))\Delta) \tag{9}$$

with the boundary condition $S^*(K) = 0$.

Here, $\rho(k,\Delta)$ is the PRP report threshold the principal uses for an evaluation period of length k and incentives Δ. The corresponding action sequence chosen by the agent is $a(k,\Delta)$. $u(a(k,\Delta))$ and $h(a(k,\Delta))$ are the present discounted benefit and cost of the action sequence $a(k,\Delta)$. $F(k, \rho(k,\Delta), a(k,\Delta))$ is the probability that the principal reports *bad* at the end of an evaluation period of length k given threshold $\rho(k,\Delta)$ and action sequence $a(k,\Delta)$.

Computing $S^*(t)$ via Equation (9) is quite straightforward. The reason is that all of the auxiliary functions - $a(k,\Delta)$, $\rho(k,\Delta)$, $u(a(k,\Delta))$, and $h(a(k,\Delta))$ - can be computed beforehand and moreover, all of these functions are defined in a direct, non-recursive way. Once they are all computed, input them into Equation (9), which now becomes a standard Bellman equation for a bounded univariate function with compact domain.

I now summarize the computations of the auxiliary functions. First, for each ρ, k, Δ, we have already seen how computing the induced effort process $a(k,\Delta,\rho)$ is as easy as computing the single induced effort in a period of the discrete model. Once $a(k,\rho,\Delta)$ is determined, I can define the associated expected present discounted benefit and cost:

$$u(a(k,\rho,\Delta)) := \int_0^k \beta^s u(a_s(k,\rho,\Delta))ds \qquad h(a(k,\rho,\Delta)) := \int_0^k \beta^s h(a_s(k,\rho,\Delta))ds \tag{10}$$

And now by definition, the auxiliary function $\rho(k,\Delta) = \arg\max_\rho u(a(k,\rho,\Delta))$. Even though this maximization involves evaluating a function defined over continuous time action sequences, it is a quite simple maximization: Lemma 4 implies that to find $\rho(k,\Delta)$ it suffices to compute $\arg\max_\rho a_0(k,\rho,\Delta)$.

And now all the other auxiliary functions are determined: $a(k,\Delta) = a(k,\rho(k,\Delta),\Delta)$, $u(a(k,\Delta)) = u(a(k,\rho(k,\Delta),\Delta))$, $h(a(k,\Delta)) = h(a(k,\rho(k,\Delta),\Delta))$ and $F(k,\rho(k,\Delta),\Delta) = \Phi^k(\rho(k,\Delta) - \int_0^k a_s(k,\Delta)ds)$.

Proposition 2. *Let $S^*(t)$ be the solution to (9), and let $k^*(t), \Delta^*(t)$ be the associated argmaxes. An arms-length contract specifying evaluation dates $0 < t_1 < t_2 < \ldots < t_N = K$ and action sequence $\{a_s\}_{s=0}^K$ is Pareto-optimal if and only if the following conditions are satisfied:*

- *For each $i = 1, 2, \ldots K$, $t_i - t_{i-1} = k^*(t_{i-1})$, and for any $s \in [t_{i-1}, t_i)$, $a_s = a_{s-t_{i-1}}(k^*(t_{i-1}), \Delta^*(t_{i-1}))$*
- *At each evaluation date t_i, if $P_{t_i} - P_{t_{i-1}}$ exceeds $\rho(k^*(t_{i-1}), \Delta^*(t_{i-1}))$ the principal reports good and the agent is retained. Otherwise the principal reports bad and the agent is terminated with probability $\Delta^*(t_{i-1})/S^*(t_i)$.*

- *If the agent is retained after the evaluation at date t_i then a Pareto-optimal arms-length contract starting at date t_i is enacted.*

In the special case $K = \infty$, characterizing Pareto-optimal contracts is even more straightforward. By self-similarity, solving the Bellman equation for the surplus function $S^*(t)$ simplifies to solving for a single surplus S^*.

Corollary 1. *In any infinite horizon Pareto-optimal arms-length contract, the surplus S^* solves the following equation:*

$$S^* = \max_{k>0,\ \Delta\in[0,S^*]} u(a(k,\Delta^*)) - h(a(k,\Delta^*)) + \beta^k(S^* - F(\rho(k,\Delta^*), a(k,\Delta^*))\Delta^*) \tag{11}$$

Let k^* be the corresponding argmax. The agent is evaluated every k^* dates. At the end of any evaluation period if the agent's performance that period exceeds $\rho(k^*,\Delta^*)$ then the principal reports good and the agent is retained. Otherwise the principal reports bad and the agent is terminated with probability Δ^*/S^*.

One of the interesting properties of arms-length contracts relates to how the effort sequence within an evaluation period evolves. Suppose the agent faces some Δ and a principal report threshold ρ in an evaluation period of length k. Then the agent solves the maximization problem in Equation (8) where $a_s(k,\Delta,\rho)$ is the agent's optimal effort choice s units of time into the evaluation period.

Notice that the agent's efforts are perfect substitutes in the benefit component but are imperfect substitutes in the cost component due to discounting. This implies that the agent's chosen effort sequence within an evaluation period starts small, monotonically increases, and approaches its maximum right before the evaluation. For example, if $h(a) = \frac{1}{2}a^2$ then $a_s(k,\Delta,\rho) = a_0(k,\Delta,\rho)/\beta^s$.

Corollary 2. *The agent procrastinates in arms-length contracts.*

6 Conclusion

This paper looks at optimal contracting in a general dynamic moral hazard model where the noisy signal is either privately observed by the principal or is observable but not verifiable. I highlight a class of contract renegotiations called private revisions that ought to survive even in the presence of a public third party with perfect commitment and ability to enforce contracts. I then characterize optimal private-revision-proof contracts. These contracts have simple threshold spot contracts and are essentially memoryless. This tractability across states and time is robust to different utility functions and signal distributions. As an application, I extend the model to continuous time and show how private-revision concerns can rationalize arms-lengths contracting.

7 Appendix

Proof of Proposition 1. Let $a(dt, K)$ denote the initial induced effort in the model with time length dt and terminal date K. Since the initial effort is always the largest, it suffices to show that

$$\lim_{dt \to 0} \lim_{K \to \infty} a(dt, K) = 0$$

Let $a(dt) := \lim_{K\to\infty} a(dt, K)$. 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. It must be that

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

for some $\lambda \in (0, \infty)$. 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 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 if $\lim_{dt\to 0} F\left(\rho(dt), 0, \frac{1}{dt}\right)/dt = \lambda$ then $\lim_{dt\to 0} F\left(\rho(dt), k, \frac{1}{dt}\right)/dt = \lambda$ for all k. Together, a change-of-variable and the standard Gaussian tail estimate imply

$$F\left(\rho(dt), 0, \frac{1}{dt}\right) = F\left(\rho(dt)\sqrt{dt}, 0, 1\right) \approx \frac{\phi(|\rho(dt)\sqrt{dt}|)}{|\rho(dt)\sqrt{dt}|}$$

Similarly,

$$F\left(\rho(dt), k, \frac{1}{dt}\right) = F\left((\rho(dt) + k)\sqrt{dt}, 0, 1\right) \approx \frac{\phi(|(\rho(dt) + k)\sqrt{dt}|)}{|(\rho(dt) + k)\sqrt{dt}|} =$$
$$\frac{\phi(|\rho(dt)\sqrt{dt}|)}{|\rho(dt)\sqrt{dt}|} \cdot \frac{|\rho(dt)\sqrt{dt}|}{|(\rho(dt) + k)\sqrt{dt}|} \cdot e^{-\rho(dt)kdt - k^2 dt/2}$$

It suffices to show that $\rho(dt)(dt)^r = 0$ for all $r > \frac{1}{2}$. Suppose not, then $F\left(\rho(dt)\sqrt{dt}, 0, 1\right) < F\left((dt)^\varepsilon, 0, 1\right)$ for some ε. But the Gaussian tail estimate implies

$$F\left((dt)^\varepsilon, 0, 1\right)/dt \sim \phi((dt)^\varepsilon)/(dt)^{1+\varepsilon} \to 0$$

Contradiction. □

Proof of Lemma 2. It suffices to prove that a contract which is PRP at evaluation dates is PRP.

Without loss of generality, consider the initial evaluation period $[0, t_1]$. Let Δ be the bonus and ρ be the report threshold at the evaluation date t_1. The induced action sequence is $a(t_1, \Delta, \rho)$ from (8). From now on I will refer to this sequence as a and use a_s denote the element of a at date s. Let $\{w_s\}_{s\in[0,t_1)}$ be the wage sequence. The agent's date t continuation payoff for any $t \in [0, t_1)$ is

$$W_t = \beta^{t_1-t}\mathbf{E}W_{t_1}\left(\rho - \int_0^{t_1} a_s ds\right) - \int_t^{t_1} \beta^{s-t}h(a_s)ds + \int_t^{t_1} \beta^{s-t}w_s ds \tag{12}$$

where $W_{t_1}(\rho - \int_0^{t_1} a_s ds)$ is the random variable $\in \{W_{t_1}(good), W_{t_1}(bad)\}$ whose distribution depends on $\rho - \int_0^{t_1} a_s ds$.

Suppose at some interim date $\tilde{t} \in (0, t_1)$, the contract is not private-revision proof. Let $\tilde{\rho}$ be a Pareto-improving report threshold. Then the agent's revised action at date $s \in [\tilde{t}, t_1)$ is

$$a'_s := a_{s-\tilde{t}}\left(t_1 - \tilde{t}, \Delta, \tilde{\rho} - \int_0^{\tilde{t}} a_s ds\right)$$

By Lemma 4, $a'_s > a_s$ for all $s \in [\tilde{t}, t_1)$. By assumption, the agent's revised continuation payoff $W'_{\tilde{t}} :=$

$$\beta^{t_1-\tilde{t}}\mathbf{E}W_{t_1}\left(\tilde{\rho} - \int_0^{\tilde{t}} a_s ds - \int_{\tilde{t}}^{t_1} a'_s ds\right) - \int_{\tilde{t}}^{t_1} \beta^{s-t}h\left(a'_s\right)ds + \int_{\tilde{t}}^{t_1} \beta^{s-t}w_s ds \geq W_{\tilde{t}} \tag{13}$$

I now show that the contract isn't private-revision-proof at date 0. Suppose the principal privately revises to $\tilde{\rho}$ at date 0. The agent's revised action sequence is $a(t_1, \Delta, \tilde{\rho})$ which I will call $\tilde{a}$ for short. Let $\tilde{W}_0$ be the agent's new date 0 payoff. By definition, $\tilde{W}_0$ is weakly larger than his payoff under the action sequence $\{a_s\}_{s\in[0,\tilde{t})} \cup \{a'_s\}_{s\in[\tilde{t},t_1)}$ which is

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

To show that the principal is also better off under $\tilde{\rho}$, consider the function $f(x) := \mathbf{E}W_{t_1}(\tilde{\rho} - x - \int_0^{t_1} \tilde{a}_s ds)$ defined over $\mathbb{R}$. Since the agent chooses $\tilde{a}$ in response to $\tilde{\rho}$, it must be that

$$f'(0) = \beta^{-t_1+s}h'(\tilde{a}_s) \quad \forall s \in [0, t_1) \tag{14}$$

Here, $f'(0)$ is the normalized marginal benefit of the action sequence $\tilde{a}$ and the constant $\beta^{-t_1+s}h'(\tilde{a}_s)$ is the normalized marginal cost. f is an increasing logistic-shaped function with a convex lower half and a concave upper half. $f(0)$ must be in the concave region. If not, then there exists a unique $x^* > 0$ such that $f'(x^*) = f'(0)$ and $f'(x) > f'(0)$ for all $x \in (0, x^*)$. The principal can achieve a Pareto-improvement by revising $\tilde{\rho}$ to $\tilde{\rho} - x^*$: the agent will still apply the same effort, but the probability of reporting *bad* is decreased.

Consider the scenario where the agent has chosen a up to date $\tilde{t}$ and faces $\tilde{\rho}$. I previously defined a' to be the agent's best response starting from date $\tilde{t}$. Suppose instead the agent chooses the action sequence $\tilde{a}$ starting from date $\tilde{t}$: At date $s \geq \tilde{t}$, he chooses $\tilde{a}_s$. The normalized marginal benefit is $f'(\int_0^{\tilde{t}}(a_s - \tilde{a}_s)ds)$ and by (14) the normalized marginal cost is $f'(0)$. If $a > \tilde{a}$, then $f'(\int_0^{\tilde{t}}(a_s - \tilde{a}_s)ds) < f'(0)$. That is, the marginal cost exceeds the marginal benefit at $\tilde{a}$. Moreover, since f is concave for all values of $x > 0$, the only way the agent can adjust effort so as to equate marginal cost and benefit is to decrease effort, which implies $a' < \tilde{a}$. Contradiction. So $a \leq \tilde{a}$ and the principal is also better off under $\tilde{\rho}$. $\square$

References

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

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

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

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