Draft

SUFFICIENCY OF LOCAL DEVIATIONS FOR DYNAMIC MORAL HAZARD WITH PRIVATE COSTS

Guy Arie[1]

Abstract

The marginal cost of effort often increases as effort is exerted. In a dynamic moral hazard setting, dynamically increasing cost create information assymetry that breaks the sufficiency of one-shot-deviation incentive compatibility (OSD-IC), preventing any analysis. This paper recovers OSD-IC and characterizes the optimal contract. The result is obtained by developing an alternative formulation for dynamic moral hazard that is based on duality theory. The stronger result utilizes a critical feature of the new formulation – a history's value is monotonic in the dual state space and satisfies a single crossing condition. The optimal contract is consistent with the popular yet thus far puzzling use of non-linear incentives, for example in sales-force compensation.

Keywords: Dynamic moral hazard, nonlinear incentives, private information, dynamic mechanism design, duality, linear programming, stochastic programming, dynamic programming.

1. INTRODUCTION

Increasing marginal costs is a standard component of economic analysis. In organizational settings, the increase in cost often has a dynamic motivation. A worker picking fruits, for example, gets tired as the day progresses. In other settings the task itself becomes harder over time. Sales performance, for example, is measured over a period, typically quarter or year. As the quarter progresses, the agent depletes all the "easy" sales leads and must exert more effort to generate later sales. Sales effort is inherently hard to monitor and pay is often performance based. If the firm would

[1]University of Rochester, Simon School of Business.
First version November 2011. This version June 2013.
This paper is based on the first chapter of my dissertation at Kellogg. I am indebted to Bill Rogerson, Mark Satterthwaite and especially Jeroen Swinkels and Michael Whinston for their advice and encouragement throughout this work. I thank many seminar participants and especially Eddie Dekel, Matthias Fahn, Paul Grieco, Jin Li, Michael Raith and Rakesh Vohra for helpful suggestions. Financial support from the GM Center at Kellogg is gratefully acknowledged.

know the agent's true cost, it may want to increase incentives towards the end of the quarter.

The existing literature only solves two period dynamic moral hazard problems in the presence of increasing marginal costs (Mukoyama and Şahin (2005) is a direct application and Ábrahám, Koehne, and Pavoni (2011) consider additional two period settings). For more periods, however, Fernandes and Phelan (2000) show that some incentive compatible (IC) contracts may violate the one-shot-deviation (OSD) condition.

This paper introduces a reformulation of the dynamic contract problem that is based on duality. This reformulation proves that while some IC contracts may violate OSD, the *optimal contract* does not. The formal optimal mechanism can be quite complicated to describe. However, section 5.3 provides an example in which the optimal contract can be implemented using a known twist on a quota contract and a second example in which the optimal contract can be implemented by a convex incentive scheme.

Joseph and Kalwani (1998) document the popularity of convex and quota based incentive schemes in sales related settings. However, as Prendergast (1999) summarizes: "rather remarkably, the theoretical literature has made little progress in understanding the observed (nonlinear) shape of compensation contracts, despite costs associated with nonlinearities." The same conclusion is echoed in more recent studies, see e.g. Misra and Nair (2009) and Larkin (2007). Thus, the analysis here shows that increasing marginal cost can provide a relatively simple micro-economic foundation for the popularity of these schemes.

The model is the simplest possible to capture the problem of privately increasing costs. A risk neutral agent decides every day whether to exert costly effort or not. The probability of success (a sale) in the day increases with effort. The cost of effort today is a convex function of past *effort*. Effort is unobserved and the principal can commit to a contract at the outset.

To see the incentive problem, suppose that the probability of a sale each period is $\frac{1}{2}$ if the agent exerts effort on a customer and zero otherwise, and that the agent's cost for making the n-th effort is n. If both the principal and the agent consider only current period incentives, a contract paying the agent $2n$ for a sale in day n is incentive compatible and provides the agent zero expected utility – clearly first best. However, if the agent considers future payoffs, this contract is no longer incentive compatible.

Shirking in the first period and then working whenever asked obtains the agent an expected utility of 1 *each* period. By shirking today the agent increases his rents from future work. The optimal contract must account for this additional incentive to shirk: the agent's utility difference between success and failure must increase.

The optimal contract can be informally described as a *dynamic quota:* the agent starts in an evaluation stage and eventually moves to a compensation stage. In the compensation stage the agent is paid a *fixed* piece-rate for each sale and works for an additional *fixed* time that is *independent* of any new outcomes. In the evaluation stage the agent is rewarded *only* by changes to the expected fixed piece-rate, the length of the compensation stage, and the quota the agent must meet to enter the compensation stage. If the agent accumulates enough early successes, his compensation per sale later in the quarter will be high. If the agent did not accumulate enough early successes, the contract leads the agent to stop working.

The contract is consistent with the more general features found in the empirical sales literature that are considered difficult to explain. Once the agent meets his "dynamic quota", his reward is based only on his highest anticipated cost, generating excessive rewards for successful agents, as found in both Misra and Nair (2009) and Larkin (2007), and is generally assumed to hold in such settings. On the other hand, the only way to profitably provide such high rewards is to limit the work by unsuccessful agents, resulting in a higher volatility of the work decision towards the end of the work period, consistent with the finding in Oyer (1998).

The analysis extends the existing literature by formulating the original problem as a linear program, deriving its dual and then obtaining a recursive representation – the *dynamic dual*. The dynamic dual analysis focuses on the optimal contract and considers a change that increases the expected profits. This change must violate some incentive constraint. If the constraint that is violated is always a one-shot-deviation constraint, then the optimal contract subject only to OSD must be IC.

The value of the dual problem at each history is the increase in expected profits at the *outset* of the contract from the optimal continuation starting in the history. This accounts for the continuation profits from the history *and* the effect of the optimal continuation on the agent's incentives from the start of the contract. In contrast, the standard (primal) dynamic value for a history is the expected continuation profits (see e.g. Spear and Srivastava (1987)). Thus, while it is well known that the contract may continue even if the standard value of a history is negative (i.e. the principal optimally committed to providing the agent very high continuation utility), whenever the dual value of a history is strictly negative, the optimal contract terminates.

The dual state variables proxy for the *degree of agency frictions* generated by the contract at the history. The model has two agency frictions - one generated by the effect of current utility on past incentives and one by the information asymmetry. Each friction determines a state variable. An important technical feature of the dual representation is that the contract's dual value is monotonic in each state variable and that the state variables are substitutes. This allows proving stronger results than typical for the optimal contract, and in particular the OSD result.[1]

Following a short literature review, section 2 lays out the dynamic production model. Section 3 develops the basic optimal contract problem and it's dynamic dual formulation. Section 4 proves the sufficiency of local deviations. Section 6 concludes.

1.1. *Relation To Existing Literature*

This paper contributes to recent progress in dynamic agency theory with private history dependent technology. Fernandes and Phelan (2000) consider a agency settings in which today's information or effort affects tomorrow's productivity but limit the history dependence to one last period via a Markov assumption. Nevertheless, the result in Fernandes and Phelan (2000) for moral hazard settings is negative - whenever today's effort affects tomorrows productivity, the one-shot-deviation principle does not apply.

Several recent papers follow the modeling approach of Fernandes and Phelan (2000) to analyze dynamic moral hazard settings with payoff relevant private histories. In a recent working paper, DeMarzo and Sannikov (2008) extend the aggregation problem of Holmstrom and Milgrom (1987) to settings in which the agent also obtains private shocks to his productivity that are correlated with past effort. While DeMarzo and Sannikov (2008) is closest to the setting studied here, the additional aggregation

[1]Following Spear and Srivastava (1987), dynamic moral hazard analysis uses the agent's continuation utility as the state. If the agent's continuation utility is exactly his outside option, the contract must typically terminate and the principal obtains his outside option. If the agent's continuation utility is very high, the principal must either give away the firm (in limited liability settings, see e.g. Clementi and Hopenhayn (2006)), or provide the agent costly insurance. In all cases, the principal's expected continuation value is highest for some expected agent's continuation utility between the two extremes and is thus non-monotonic. In the extension to private information developed by Fernandes and Phelan (2000), it may well be that the state variables are in some cases complements and in others substitutes.

problem changes the analysis and results. Another working paper, Tchistyi (2006), maintains the Markov structure in Fernandes and Phelan (2000) and devises a transformation of the agent continuation payoffs under deviations to obtain sufficiency of local deviations in the presence of unobservable utility shocks to the agent. Williams (2011) focuses on dynamic adverse selection – the agent only reports his income, which privately follows a Brownian motion.

Bergemann and Hege (2005), Bonatti and Hörner (2011) and Halac, Kartik, and Liu (2012) consider the case that surplus depends on the private history but the principal only cares about the first success. This simplifies the agency problem, and restores the one-shot-deviation principle as there is only a single instance in which rewards need to be provided. As a result, more involved questions – alternative contracting frameworks and collaboration between multiple agents can be studied.

Duality based approaches are widely used in economic modeling, dating back to the classic text by Rockafellar (1970). Vohra (2011) extends the analysis of static adverse selection models by analyzing the dual of the classic adverse selection problem. Marcet and Marimon (2011) and Mele (2011) consider shadow multipliers in a dynamic setting that can be applied to moral hazard problems. The formal analysis however assumes the OSD assumption holds. Abraham and Pavoni (2008) use a mixture of a shadow variable and the promised utility to construct a recursive model of savings and consumption. Their approach however relies on a numerical procedure to verify ex-post that the first order approach is valid. Finally, Mukoyama and Şahin (2005); Ábrahám, Koehne, and Pavoni (2011) consider related production processes for two period settings.

This paper makes several contributions. First, a characterization is provided to the specific setting of interest, which is more consistent with empirical observations than previous models. Second, the recursive dual model developed here is tractable and the approach can be applied to other settings. In particular, settings without a proof for sufficiency of OSD. Finally, while the previous dual literature focused only on the utility cost, this paper introduces the use of the private information cost as a dual variable.

2. MODEL

2.1. *Setup and Primitives*

There is a principal and an agent, both risk neutral. Both have an outside option set to zero. The agent has limited liability – i.e. money can only be transferred to the agent. Time is discrete. In each period the agent either works or not. The agent's work is costly to the agent and unobservable to the principal. The cost of effort in a period is c_n for a commonly known function $c : N \to R_+$ where n denotes the number of actual periods of work. That is, for the first period of work, n is one. For the second period of work, n is two, and so on. If the agent shirks in the first period, n in the second period is still one. The analysis will focus on the case that c_n is an increasing and convex function. However, the dual methodology that is developed applies also if c is fixed or non-monotonic.

ASSUMPTION 1 *The agent's cost, c_n, is increasing and convex in the work period number, n.*

A period's production outcome is either success or failure, denoted by $y \in Y = \{0, 1\}$. The principal earns a revenue of v from each success ($y = 1$) and zero from a failure ($y = 0$). The probability of success (resp. failure) in a period in which the agent works is $p \in (0, 1)$ (resp. $1 - p$). If the agent does not work, p is replaced with $p_0 \in [0, p)$. The results extend directly if the outcome space is increased to any countable set. To prevent the principal from making free profits, assume the principal incurs a cost of $v \cdot p_0$ for every period in which the contract is still active.[2]

As costs are increasing, the surplus from working becomes negative after enough effort was exerted. Let N^{FB} denote the maximum number of periods in which consecutive work increases surplus:

$$N^{FB} = \max n : c_n \leq v\,(p - p_0)\ .$$

The increase in costs is sufficient to prevent an infinite contract from being optimal.[3] The exposition is simplified by assuming that the agent and principal do not discount the future. Section 5.5 explicitly considers discounting and shows that adding

[2]This assumption only simplifies the exposition and is without loss of generality.

[3]Nevertheless, the dual formulation developed here can be identically derived for infinite horizon settings.

a discount factor amounts to a simple accounting exercise.

Before defining the contract, the following observation simplifies the exposition and notation. The first two parts are standard and the last is a direct outcome of risk neutrality.

REMARK 1 There is an optimal contract in which:

1. The agent works for at most N^{FB} periods
2. The required work decision is a stopping decision: if the agent is ever asked not to work, the contract terminates.
3. The agent is never paid in a period without work or with failure.

Given remark 1, the space of contract relevant pubic histories H is the space of previous outcomes:

$$H = \bigcup_{n=0}^{N^{FB}} Y^n \quad .$$

A public history $h \in H$ denotes a sequence of outcomes. Part two of Remark 1 implies that if the contract is not yet terminated, the agent was asked to work in all past periods. However, only the agent knows in which periods he actually did work and in which periods he shirked. As the cost to the agent of working in a period is a function of the *number* of periods in which the agent actually worked in the past, the only information in the agent's private history that is payoff relevant is the number of past shirks:

DEFINITION 1 The agent's *private* history (h, s) is the public history h and the number of past shirks s.

Let n_h denote the number of the period just after history h. If the agent did not deviate in the past, his cost of work in this period will be c_{n_h}. However, cost depends on the *private* history. With a slight abuse of notation let $c_{h-s} \equiv c_{n_h - s}$ denote the cost for any history h with past deviations s and $c_h \equiv c_{h-0}$. As the difference in cost between two work periods will play an important role, let $d_{h-s} \equiv c_{n_h-s} - c_{n_h-s-1}$ denote the cost difference between the current and previous periods if the agent shirked s times in the past and $d_h \equiv d_{h-0}$. To simplify the notation later on, set $d_1 = c_1$.

The analysis makes extensive use of histories *following* and *preceding* other histories.

Let $h = \langle h^1, h^2 \rangle$ denote the history h^1 followed by the history h^2. That is, the sequence of outcomes h^1 happened and then the sequence h^2 happened. For example, if the current history is h, then the next history will be either $\langle h, 1 \rangle$ or $\langle h, 0 \rangle$. Say that the history $\langle h^1, h^2 \rangle$ *follows* history h^1 and denote the "follows" relation by $\succeq$. That is[4]

$$\tilde{h} \succeq h \iff \exists \hat{h} \in H : \tilde{h} = \left\langle h, \hat{h} \right\rangle .$$

2.2. *The Contract*

By the revelation principle, there is no loss in considering only contracts that specify for each period a work decision and a wage based on the period's history. This section defines the contract using a simple transformation of the standard decision variables. This transformation will allow a linear formulation of the problem without affecting the interpretation of the resulting contract.

Typically, the contract specifies for each period whether the contract terminates in the period and the resulting wage. Let $(1 - a_h)$ denote the probability that the contract is terminated in history h, if h is reached. Let W_h denote the wage paid for success in the period.

Using the standard variables, the ex-ante probability that the contract would still be active after a success is $a_\emptyset \cdot a_{\{1\}}$. To linearize the formulation, the contract will use the *ex-ante* probability that the principal did not decide to terminate the contract yet instead. This probability, denoted q_h, can be defined recursively from any a:

$$q_\emptyset \equiv a_\emptyset \quad \text{, and } q_{\langle h,y \rangle} \equiv q_h \cdot a_{\langle h,y \rangle} \tag{2.1}$$

The standard payment W_h is only paid if the contract was not terminated by period h and the agent succeeds. Thus, to determine the expected payment, we must multiply W_h by q_h. This would destroy the linear nature of the problem. However, by remark 1, there is no loss of generality in having the contract specify the wage for success in the history, conditional on the contract not terminating yet:

$$w_h \equiv q_h \cdot W_h \; . \tag{2.2}$$

Definition 2 A contract is a pair of functions $\langle q, w \rangle$, with q_h specifying for each history h the cumulative probability that the agent will be still be asked to work in the period (equations 2.1) and w_h specifying the wage conditional on activity (equation 2.2).

[4] As the set H includes the empty set, $h \succeq h$.

Remark 2 As the agent and principal are risk neutral, there are infinitely many equivalent ways for the optimal contract to pay the agent w dollars. Paying a dollar for success today is equivalent to paying $\frac{1}{p}$ more dollars for success tomorrow (assuming the agent is asked to work). To remove this technical duplication of the optimal contracts, the analysis assumes the optimal contract makes payments "as early as possible." That is, from all contracts that specify the same work plan, the optimal contract is the one that pays more to the agent as early as possible.[5]

If the agent complies with the contract (i.e. works when asked to), the ex-ante expected revenue for the principal from a history h is the probability that the history is reached and the contract did not terminate, q_h, multiplied by the expected revenue from work $v \cdot (p - p_0)$. The ex-ante expected payment in a period h for a compliant agent is simply $p \cdot w_h$. Thus, the principal's expected continuation profit starting at history h from a contract the agent complies with is

$$V^h = q_h (p - p_0) v - pw_h + pV^{\langle h,1\rangle} + (1-p) V^{\langle h,0\rangle} \tag{2.3}$$

A similar expected value can be defined for the agent. Letting U_s^h be the agent's expected continuation utility from complying with the contract in all remaining periods starting at private history h, s we have:

$$U_s^h \equiv pw_h - q_h c_{h-s} + pU_s^{\langle h,1\rangle} + (1-p) U_s^{\langle h,0\rangle}$$

The $'s'$ subscript is omitted if $s = 0$. The optimal contract chooses q, w that maximizes $V^{\emptyset}$ subject to incentive compatibility (IC) and individual rationality (IR). As the agent can always choose to stop working IR is implied by IC and thus will be subsequently ignored.

2.3. *Incentive Compatibility*

Rather than considering general deviation plans by the agent, we consider only "Final Deviation Incentive Constraints" (FDIC). That is, IC's in which the agent (who possibly shirked in the past) considers one final shirk, assuming he must follow the contract in all later periods.

For any private history (h, s), the agent's expected utility from complying is U_s^h. To construct the FDIC, determine the agent's expected utility from shirking. If the

[5]The selection is correct if the agent is just slightly more impatient or risk averse than the principal as the early payments condition implies that the payments depend on fewer outcomes. As there is no issue of inter-temporal consumption smoothing this reduces the variation in the agent's compensation and thus increases his expected utility if he is risk averse without affecting the principal's profit.

agent deviates in history h, his expected payment in the period is $p_0 w_h$. The agent expects to succeed and transition to the history $\langle h, 1\rangle$ with probability p_0. As the agent is considering a final deviation, his expected continuation utility after shirking and succeeding is $U_{s+1}^{\langle h,1\rangle}$. Similarly, with probability $(1-p_0)$ the agent expects to transition to the private history $(\langle h, 0\rangle, s+1)$.This yields the following definition:[6]

Definition 3 The set of Final Deviations Incentive Constraints (FDIC) is

$$\forall h, s: \quad U_s^h \geq p_0 w_h + p_0 U_{s+1}^{\langle h,1\rangle} + (1-p_0)\, U_{s+1}^{\langle h,0\rangle} \quad \text{(FDIC)} \tag{2.4}$$

The next lemma verifies the relation between FDIC and IC:

Lemma 1 *If a contract is FDIC it is IC*

The intuition for the lemma is simply that any profitable deviation plan must have a profitable last deviation. However, FDIC is a stricter condition than IC. For example, it may be that the first deviation was more costly to the agent than the gain from the second deviation. In contrast, the subset of FDIC that correspond to the one-shot-deviation principle is weaker than IC:

Definition 4 The set of Local Deviation Incentive Constraints (LDIC) is

$$\forall h: \quad U_0^h \geq p_0 w_h + p_0 U_1^{\langle h,1\rangle} + (1-p_0)\, U_1^{\langle h,0\rangle} \quad \text{(LDIC)} \tag{2.5}$$

The LDIC require that whenever the agent considers a *one and only* deviation, the expected continuation profit is lower than the expected continuation profit from complying with the contract in all future periods. It is clear that IC implies LDIC. However, as in Fernandes and Phelan (2000), in the current model, LDIC does not imply IC:

Lemma 2 *A contract may satisfy all LDIC but violate IC.*

The proof in the appendix is by an example contract in which all LDIC hold with equality while the agent's optimal plan is to shirk in the first two periods. As FDIC is stricter than IC which in turn is stricter than LDIC, the following corollary follows:

Corollary 1 *If every optimal contract subject to LDIC satisfies FDIC, then every optimal contract subject to LDIC is optimal subject to IC.*

[6] See the appendix for a derivation of the FDIC directly from the problem's primitives.

The proof for sufficiency of LDIC will show that the condition of corollary 1 does in fact hold. For this, we must analyze the FDIC problem.

3. THE FDIC PROBLEM

3.1. *Standard Form*

The FDIC problem can be stated recursively using the "promised" utility for all possible private histories. This is a natural extension of Fernandes and Phelan (2000). The FDIC in period n for a previous number of shirks s must consider the agent's continuation utilities if the agent shirked $s+1$ times. As a result, the recursive problem is period n considers not only the standard utility regeneration constraint, but also the regeneration constraint for all possible previous number of shirks. The resulting dynamic problem is provided below, with the shadow cost identified for each constraint:

Definition 5 The Dynamic FDIC problem is

$$\max_{U\geq 0} V\left(1,1,U\right)$$

With $V\left(\right)$defined by:

(3.1)

$$\begin{aligned} V\left(n,\overline{q},U_0,...,U_{n-1}\right) &= \max_{U_0^y,...,U_n^y,q,w\geq 0} && q\left(p-p_0\right)v-p\cdot w \\ & && +pV\left(n+1,q,U_0^1,...,U_n^1\right) \\ & && +\left(1-p\right)V\left(n+1,q,U_0^0,...,U_n^0\right) \\ \text{subject to} & && \\ \text{probability constraint} & \quad (\mu_s) && \overline{q}\geq q \\ FDIC\ \forall s\leq n-1 & \quad (\lambda_s) && U_s\geq p_0 w+p_0 U_{s+1}^1+\left(1-p_0\right)U_{s+1}^0 \\ \text{Regeneration}\ \forall s\leq n-1 & \quad (\gamma_s) && U_s = pw-qc_{n-s}+pU_s^1+\left(1-p\right)U_s^0 \end{aligned}$$

The objective formulation is standard and the FDIC are identical to 2.4.

However, there are two non-standard elements in problem 3.1. First, the variable $\overline{q}$ and the accompanying probability constraint. In the standard notation, the upper bound on the probability of work is $'1'$. Here, as q_h identifies the ex-ante probability, the upper bound is simply the previous period's probability, which may be lower than 1:

$$0\leq q_{\emptyset}\leq 1 \quad ,\text{and } 0\leq q_{\langle h,y\rangle}\leq q_h$$

The probability constraint, together with the determination of the next period's $\overline{q}$ enforces this change.

The second non-standard element are the extra regeneration constraints for the utilities in the various private histories. These are called the "threat constraints" in Fernandes and Phelan (2000). The derivation there details the need for these for $s = 1$ assuming the one-shot-deviation principle holds. If OSD does not hold, the problem must have these for all possible private histories, as is the case here. The intuition for the additional regeneration constraints is simple. As the problem must set some continuation values after deviation in the FDIC, the dynamic problem must recursively define those values and maintain these just as it does for the original utility in the standard formulation.

The Dynamic LDIC problem is the same as 3.1 with the FDIC only for $s = 0$ and the regeneration constraint only for $s = 0$ and $s = 1$.

There are several difficulties with the dynamic problem. First, it is not well defined for all values of $U_0, ..., U_{n-1}$. For example, there is no solution if $U_0 = 0$ and $U_1 > 0$. This complicates proving even "standard" results, such as concavity. Second, one cannot prove using this problem (at least I could not) that in the optimal LDIC contract all FDIC are slack, which is critical for the remainder of the analysis. Instead, the analysis will proceed by deriving the dual of the dynamic problem.

Before deriving the dynamic dual problem, it would help to separate the two dynamic effects of shirking for the agent. First, shirking affects the transition probabilities to future histories. As a result, if the agent expects different continuation utilities in these histories, shirking affects the agent's expected continuation utility. This is the standard effect in all dynamic moral hazard settings. However, in our setting shirking also generates a *private* reduction in the agent's future costs. Separating these two effects will allow the analysis to correctly account for the change in each effect as the contract progresses.

For this distinction, let D_s^h denote the increase in the agent's utility starting in history h if he would have made one more shirk at some point in the past. That is D_0^h is the agent's gain from making one shirk in the past, D_1^h is the agent's gain from making a second shirk in the past, and so on. By construction:

$$D_s^h \equiv q_h d_{n-s} + p D_s^{\langle h,1\rangle} + (1-p)\, D_s^{\langle h,0\rangle}$$

Placing the regeneration constraint in the FDIC 2.4, collecting terms and replacing $\left(U_{s+1}^h - U_s^h\right)$ with D_s^h obtains a version of the FDIC in history (h, s) that directly reflects these two effects:

$$w\left(p-p_0\right)-q\cdot c_{n-s}+\left(p-p_0\right)\left(U_s^1-U_s^0\right)-p_0D_s^1-\left(1-p_0\right)D_s^0\geq 0 \tag{3.2}$$

The FDIC 3.2 reflects the work/shirk tradeoff for an agent. The first two terms are the direct tradeoffs: working increases the probability of getting paid but has a cost. The third term reflects the tradeoff between the continuation utilities. Working increases the probability of transitioning to the continuation after success (and thus obtaining the associated utility) and decreases the probability of transitioning to the continuation after failure. The last two terms reflect the private cost reduction that is forgone by working. If the agent shirks, he increases his utility after success, for example, by D_s^1. Thus, working implies losing this utility (adjusted for the change in transition probabilities).

The dynamic FDIC can be rewritten using these additional variables. To save on notation, $\vec{U}$ and $\vec{D}$ stand for the vector for all relevant s values and y for the possible outcomes. The shadow cost for each constraint is provided as well:

$$\begin{aligned}
V\left(n,\overline{q},\vec{U},\vec{D}\right) &= \max_{\left(\vec{U}^y,\vec{D}^y,q,w\right)\geq 0} \quad q\left(p-p_0\right)v-p\cdot w \\
&\qquad +pV\left(n+1,q,\vec{U}^1,\vec{D}^1\right) \\
&\qquad +\left(1-p\right)V\left(n+1,q,\vec{U}^0,\vec{D}^0\right)
\end{aligned} \tag{3.3}$$

$$\begin{array}{rcl}
\text{subject to} & & \\
\text{Probability} & (\mu_s) & q\leq\overline{q} \\
FDIC\ \forall s\leq n-1 & (\lambda_s) & Constraint\ 3.2 \\
\text{Regeneration}\ \forall s\leq n-1 & (\gamma_s) & U_s=pw-qc_{n-s}+pU_s^1+\left(1-p\right)U_s^0 \\
\text{Regeneration -}D\ \forall s\leq n-1 & (\delta_s) & D_s=qd_{n-s}+pD_s^1+\left(1-p\right)D_s^0
\end{array}$$

Problems 3.3 and 3.1 are equivalent, and the optimal contract is the solution to

$$\max_{U\geq 0,D\geq 0} V\left(1,1,U,D\right)$$

3.2. *Dual Form*

This section derives the dynamic dual representation of problem 3.3. Proposition 1 establishes that this dual is problem 3.11. The derivation here is informally based on first order condition arguments. This makes explicit the underlying economics and the (limited) importance of the linearity of the problem. A formal derivation using linear programming duality is provided in the appendix.

In problem 3.3, each of the state variables (except n) appears in the left hand side of one constraint. This is the probability constraint for $\overline{q}$ and the applicable regeneration constraint for any U_s or D_s. Thus, the shadow cost for each of these constraints is the optimal contract value for a change in the applicable state variable. Formally, if all $V(n, \cdot)$ functions are differentiable at the optimal solution, then for every history h the optimal V^h, U_s^h, D_s^h and shadow costs μ^h, γ_s^h , δ_s^h must satisfy

$$\text{(3.4)} \qquad \frac{\partial V^h}{\partial \overline{q}} = \mu^h \quad ; \quad \frac{\partial V^h}{\partial U_s} = \gamma_s^h \quad ; \quad \frac{\partial V^h}{\partial D_s} = \delta_s^h$$

Consider first the probability constraint $q \leq \overline{q}$ in a specific history. Because the original problem is linear in all q variables, if the subgame starting in the history increases the overall profitability of the contract, the constraint should bind. Moreover, μ^h, the shadow cost on the constraint, should capture exactly the marginal value of the contract starting from this period.

Definition 6 The dual value of a history, denoted μ^h, is the ex-ante increase in profit from the subgame starting in the history h.

It should be stressed that the dual value is *not* the continuation value. As will become clear below, the ex-ante increase in profits from any sub-game starting at h must also account for any change (positive or negative) in the contract terms in histories that preceded h as a result of the subgame starting in h.

As the other two shadow costs are used in deriving the dual value of a history, we first derive those.

By equation 3.4, γ_s^h captures the optimal contract's ex-ante cost of providing utility to an agent starting in the private history (h, s). This is the standard dynamic moral hazard consideration. Utility in period h affects the agent's preferences over possible continuations in all the histories that *precede* h. If h follows a success in history $\tilde{h}$, for example, the additional utility from history h increases the agent's incentive to succeed in $\tilde{h}$. This should allow the contract to reduce the other incentives provided in the preceding histories and save some costs. While the contract can simply adjust the payment to the in preceding histories to compensate for the dynamic effect of the utility change in h, this may not be optimal. A more profitable way to provide the utility may be to add some more work (and thus utility) in some other history h'.

The dual analysis directly determines history h's marginal *utility cost*. We first define it and then provide an intuitive explanation.

DEFINITION 7 The utility cost for any private history h, s, denoted γ_s^h is the probability weighted sum of all preceding shadow costs on the relevant continuation utility terms:

$$\gamma_s^h = (p - p_0)\left[-\sum_{\tilde{h}:h\succeq\langle\tilde{h},1\rangle}\frac{\lambda_s^{\tilde{h}}}{p} + \sum_{\tilde{h}:h\succeq\langle\tilde{h},0\rangle}\frac{\lambda_s^{\tilde{h}}}{1-p}\right]. \tag{3.5}$$

Equation 3.5 defines the marginal utility cost in history h as the weighted sum of all the shadow costs that *preceded* the history. A detailed derivation of the utility cost γ is provided in the appendix using linear programming duality as part of the proof for proposition 1.

The key insight is that the FDIC shadow cost, $\lambda_s^{\tilde{h}}$, captures the marginal cost to the principal of a change in the incentives in any public history $\tilde{h}$. In particular, $\lambda_s^{\tilde{h}}$ can be used to capture the incentive cost effect in history $\tilde{h}$ from a change in utility in any later history h .

Observing the FDIC 3.2, it follows that history h's marginal utility cost in history $\tilde{h}$ is either $\left(-\lambda_s^{\tilde{h}}\right)(p - p_0)$ if h follows a success in $\tilde{h}$. That is, an increase in utility from history h relaxes the FDIC in history $\tilde{h}$ that precedes h by an amount worth to the principal $\lambda_s^{\tilde{h}}(p - p_0)$. The exact opposite holds if h follows a failure in $\tilde{h}$.

Aggregating the costs in all previous periods will provide the utility cost γ^h. However, this requires one additional accounting exercise. Given that history h was reached, the outcome in history $\tilde{h}$ is known and so must be accounted for. To see this, assume differentiability and write the first order condition for U_s^1 and U_s^0 for problem 3.3:[7]

$$(p - p_0)\lambda_s + p\gamma_s^1 - p\gamma_s = 0 \quad ; \quad \text{and} \quad (p - p_0)\lambda_s + p\gamma_s^1 - p\gamma_s = 0$$

[7]A superscript on the shadow cost identifies the value for the corresponding continuation history (i.e. γ_s^1 for the shadow cost γ_s in the continuation history after success).

Isolating γ_s^y in each yields:

$$(3.6)\qquad \gamma_s^1 = \gamma_s - \frac{p - p_0}{p}\lambda_s \quad ; \quad \text{and} \quad \gamma_s^0 = \gamma_s + \frac{p - p_0}{1 - p}\lambda_s$$

Aggregating these obtains equation 3.5.

The final shadow cost, δ_s^h identifies the second dynamic effect of adding work in history h. This is the agent's incentive to increase his private information rents. By making the $(s+1)$ shirk in any preceding history, the agent reduces his cost in history h by d_{h-s}. This cost reduction is valuable to the agent only if he is asked to work in history h. Thus, in any preceding history $\tilde{h}$, adding work in history h increases the agent's incentive to shirk through the increase of the potential cost difference. Again, translating the change in the agent's incentive to shirk to the change in the contract's optimal profit is done through the incentive constraint. The FDIC 3.2 identifies h's incentive cost in history $\tilde{h}$ with past shirks s as $p_0\lambda_s^{\tilde{h}}$ if h follows a success in $\tilde{h}$ and as $(1 - p_0)\,\lambda_s^{\tilde{h}}$ if h follows a failure. Applying the same probability adjustments as for γ yields the next definition:

Definition 8 The *information rent cost* for any private history h, s, denoted δ_s^h is the probability weighted sum of all preceding shadow costs on the future shirking gains terms:

$$(3.7)\qquad \delta_s^h = \frac{p_0}{p} \sum_{\tilde{h}: h \succeq \langle \tilde{h}, 1\rangle} \lambda_s^{\tilde{h}} + \frac{1 - p_0}{1 - p} \sum_{\tilde{h}: h \succeq \langle \tilde{h}, 0\rangle} \lambda_s^{\tilde{h}} \,.$$

As for the utility cost, the value is derived in the appendix using duality, but can be obtained also by rearranging the first order conditions for D_s^1 and D_s^0 for problem 3.3 as above for U_s^y:

$$(3.8)\qquad \delta_s^1 = \delta_s + \frac{p_0}{p}\lambda_s \quad \text{and} \quad \delta_s^0 = \delta_s + \frac{1 - p_0}{1 - p}\lambda_s$$

Recursively aggregating obtains the value for any specific history.

If we can formulate the optimal contract problem using definitions 7 and 8, then the marginal cost for the contract of the agent's utility and information rent at any history can be constructed using only the FDIC shadow costs (λ) of the *preceding* histories, making a recursive formulation possible.

The dual value of each period (μ) in the optimal contract can now be formally derived from the first order condition on q in problem 3.3:

$$\mu \geq v \cdot (p - p_0) - \sum_{s=0}^{n-1} (c_{n-s}\lambda_s - c_{n-s}\gamma_s + d_{n-s}\delta_s) + p\mu^1 + (1-p)\,\mu^0 \tag{3.9}$$

To understand the dual value, suppose there is an optimal incentive compatible contract that terminates at the start of history h (before the agent is asked to work). Consider asking the agent to work in the history h, while making all the required adjustments to maintain incentive compatibility at the lowest possible cost to the principal.

The direct effect for the principal is an increase in expected profit of $v\,(p - p_0)$. However, the agent must be incentivized for his effort. We do not know yet whether in the optimal contract the agent will be paid for his effort, or compensated in any other way (e.g. even more future work). Nevertheless, for each private history s the FDIC provides the cost of this incentive for the principal: $\lambda_s c_{n-s}$. Note that the cost depends on the agent's private information s. In any optimal contract, only one private history will 'bind', which will be the private history with a non-zero FDIC shadow cost. This will be the strategy to determine that only the 'honest' private history really requires incentives.

In a static setting, the two direct effects are all that should be considered. However, our setting has three dynamic effects. First, recall that γ_s is the cost for the optimal contract for providing the agent utility in the period. As work causes the agent disutility, the change is "saving" the optimal contract $\gamma_s c_{n-s}$ (the payment to the agent is captured separately). Again, the true cost depends on the agent's private information s.

For the second dynamic effect, recall that δ_s captures the cost to the optimal contract from the agent's extra incentive to shirk in preceding histories to generate the private cost difference d_{n-s}. If the agent is not asked to work, he cannot enjoy this private cost difference. Thus, asking the agent to work generates the additional cost $\delta_s \cdot d_{n-s}$.

The final dynamic effect is the continuation value: $p\mu^1 + (1-p)\,\mu^0$.

The dynamic dual problem in each period identifies the period FDIC shadow costs λ_s that minimize μ, given the utility and information costs γ and δ. As the current shadow costs λ determine the utility and information costs in the next period, equation 3.9 and the recursive definitions for γ and δ can be used to construct a recursive formulation of the μ minimization problem.

The final piece of the puzzle is the decision whether, and how much, to pay the agent in the period. This is determined using the first order condition for problem 3.3 with respect to w, which can be written as:

$$(p-p_0)\sum_{s=0}^{n-1}\lambda_s \leq p + p\sum_{s=0}^{n-1}\gamma_s \tag{3.10}$$

The left hand side of inequality 3.10 identifies the direct effect on the agent from increasing the period payment: incentives to work increase at a rate of $(p-p_0)$. This relaxes the FDIC and is thus worth $(p-p_0)\,\lambda_s$ for any private history s. The right hand side identifies the effect on the contract value. First, the contract expects to pay the wage at a rate p. Second, paying the agent increases his expected utility in the period at a rate p, which costs the optimal contract $p\cdot\gamma_s$ for any private history s. The *wage constraint* requires that in any optimal contract, a payment is made if and only if the gains outweigh the costs. Whenever the inequality is slack, no payment should be made.

Aggregating all of the above, the dynamic dual problem is formulated recursively (again using $\vec{}$ to denote vectors of length $n-1$):

(3.11)

$$\mu\left(n,\vec{\gamma},\vec{\delta}\right) = \max \left[0,\ \min_{\vec{\lambda}\geq 0}\mu\left(n,\vec{\gamma},\vec{\delta},\vec{\lambda}\right)\ \right]$$

$$s.t.$$

$$\mu\left(n,\vec{\gamma},\vec{\delta},\vec{\lambda}\right) = v\,(p-p_0) - \sum_{s=0}^{n-1}\left(c_{n-s}\lambda_s - c_{n-s}\gamma_s + d_{n-s}\delta_s\right)$$
$$+p\mu\left(n+1,\vec{\gamma}^1,\vec{\delta}^1\right) + (1-p)\,\mu\left(n+1,\vec{\gamma}^0,\vec{\delta}^0\right)$$

wage constraint: (3.10)

utility costs: (3.6)

information rent: (3.8)

stopping condition: $\mu\left(N^{FB}+1,\vec{\gamma},\vec{\delta}\right) = 0$

PROPOSITION 1 *Problem 3.11 is the dual of problem 3.3. In particular*

$$\mu(1,0,0) = \max_{U,D \geq 0} V(1,1,U,D) = \max_{U \geq 0} V(1,1,U) \quad \textit{and}$$

the dual choice variables in every history are the FDIC shadow costs λ_s in both standard problems

The dual formulation identifies the value of a history by the transformation of the state space. Instead of using the standard promised utility, the problem uses the marginal cost of providing utility to the agent (γ). In addition, instead of using the extra 'promised' utility for an agent with private information, the problem uses the marginal cost of the information rent (δ).

Because the state space reflects 'costs', this formulation has desirable properties – convexity, monotonicity and single-crossing – that are proved below. Intuitively, higher costs are always 'bad'. In contrast, it is well known that in the standard formulation higher utility for the agent may well be required to increase overall efficiency and the principal's profits. Thus, the standard dynamic moral hazard problem is non-monotonic in its state variable (the agent's promised utility). These properties are exploited below to prove the one-shot-deviation (OSD) result, and later characterize the optimal contract.

4. SUFFICIENCY OF LOCAL DEVIATIONS

This section establishes that it is sufficient to consider only local deviations, which is established in theorem 1. The dual problem can be thought of as considering, in each public history, which private history requires the strongest incentives. The contract identifies for any history h, s determines the binding "Final Deviation Incentive Constraint" (FDIC). This is similar to letting s be an agent's 'type' and identifying the 'type' that requires the strongest incentive to work. The result of this section is that, for the optimal contract, the honest type, for which $s = 0$, requires the strongest incentives at all periods. Thus, it is sufficient to consider the "Local Deviation Incentive Constraint" (LDIC). Formally:

LEMMA 3 *If every solution to the FDIC dual satisfies also the constraint 4.1, then any optimal solution subject only to LDIC is an optimal contract in the original* *problem.*

$$\forall h, \forall s>0: \ \lambda_s^h = 0 \quad . \tag{4.1}$$

PROOF: Lemma 3 is a direct implication of complementary slackness given corollary 1. *Q.E.D.*

The reason that only LDIC should bind in the optimal contract appears simple. If the agent did shirk in the past, his costs this period are lower (c_n is increasing) and the effect of his shirking on future costs is lower (d_n is increasing). Therefore, shirking in the past lowers the incentives to shirk and LDIC should be sufficient. However, this intuition ignores a potential complication: if the agent expects to work more after failing than after succeeding, he may have a stronger incentive to shirk and fail so to enjoy the future lower costs may be larger.

To show that the intuition survives this wrinkle, we implement a "summation by parts" exercise that allows qualitative comparison and makes the argument mathematically explicit. This is similar to the integration by parts exercise in Myerson (1981).

4.1. *Private Information and Summation By Parts*

The methodology is explained using a two period problem. The details extending the result to multiple periods are provided in the appendix. If there are only two periods, the resulting second period dual problem can be written as:[8]

(4.2)

$$\begin{aligned} \mu(2,\gamma,\delta) = \max \quad & \left[0, \min_{(\lambda_0,\lambda_1)\geq 0} v(p-p_0) + \gamma\cdot c_2 - \delta\cdot d_2 - \lambda_0 c_2 - \lambda_1 c_1\right] \\ s.t. \quad & (p-p_0)(\lambda_0+\lambda_1) \leq p(1+\gamma) \end{aligned}$$

As $c_2 > c_1$, it is immediate that $\lambda_1 = 0$ and the condition of lemma 3 holds. However, there are two drawbacks to extending this formulation to multi-period analysis. First, λ_0 and λ_1 have the same sign in the objective. This would complicate determining that $\lambda_0 > 0$ while $\lambda_1 = 0$. For the line of proof adopted below, the problem is that if the problem is super-modular in λ_0, it is also super-modular in λ_1. Second, the feasible

[8]In the second period, γ and δ are scalars as the first period only had one shadow cost.

space of $(\lambda_0, \lambda_1) \times \gamma$ is *not* a sub-lattice.[9] The sub-lattice structure is required to apply the monotone comparative static arguments made in the proof.

The two problems identified in the previous paragraph are averted by transforming the shadow variables λ_s^h in each period using partial sums. Let Λ^h be a vector of length n_h such that Λ_m^h for $m = \{0, ..., n_h\}$ is the sum of the *last* $n_h - m$ elements of λ_s^h. That is:

$$\Lambda_m^h \equiv \sum_{s=m}^{n} \lambda_s^h \ . \tag{4.3}$$

This means that Λ_0^h is the sum of all the shadow costs for the public history h, while $\Lambda_{n_h-1}^h$ is exactly the last shadow cost, $\lambda_{n_h-1}^h$. To guarantee that the requirement $\lambda_s^h \geq 0$ is met, the following conditions must be added to the dual:

$$\begin{aligned} \Lambda_m^h - \Lambda_{m+1}^h &\geq 0 \\ \Lambda_m^h &\geq 0 \ . \end{aligned} \tag{4.4}$$

The first inequality guarantees that each λ_s^h is positive, except for $\lambda_{n_h-1}^h$. The second inequality guarantees that $\lambda_{n_{h-1}}^h$ is positive.[10] Given (4.4), the condition of lemma 3 can be re-written as:

$$\forall h, \ \Lambda_1^h = 0 \quad . \tag{4.5}$$

Corollary 2 *If every solution to the FDIC dual satisfies also the constraints 4.4 and 4.5, then any optimal contract subject to LDIC is optimal in the original problem.*

---

[9]That is, it may very well be that

$$\lambda_0^h + \lambda_1^h \leq \frac{p}{p - p_0} (1 + \gamma)$$

and

$$\hat{\lambda}_0^h + \hat{\lambda}_1^h \leq \frac{p}{p - p_0} (1 + \hat{\gamma}) \quad ,$$

but

$$\max\left[\hat{\lambda}_0^h, \lambda_0^h\right] + \max\left[\hat{\lambda}_1^h, \lambda_1^h\right] > \frac{p}{p - p_0} (1 + \max[\hat{\gamma}, \gamma]) \quad .$$

[10]Requiring only that $\Lambda_{n_h-1}^h \geq 0$ in the second line would imply, together with the first line that all $\Lambda_m^h \geq 0$ are positive, but would be more cumbersome.

The second period problem using Λ instead of λ is:

(4.6)

$$\begin{aligned}\mu(2,\gamma,\delta) = \max \quad & \left[0, \min_{(\Lambda_0,\Lambda_1)\geq 0} v(p-p_0) + \gamma\cdot c_n - \delta\cdot d_n - \Lambda_0 c_n + \Lambda_1(c_n - c_{n-1})\right] \\ s.t. \quad & (p-p_0)\Lambda_0 \leq p(1+\gamma) \\ & \Lambda_0 - \Lambda_1 \geq 0\end{aligned}$$

Observe that the two challenges identified after problem 4.2 are resolved. Λ_0 decreases the objective while Λ_1 increases it, and setting $\Lambda_1 = 0$ does not impose a limit on Λ_0. In addition, the space of $(\Lambda_0, \Lambda_1) \times \gamma$ that are feasible is a lattice.

4.2. *The Multiple Periods FDIC Dual*

As summation by parts is preserved under summation, it may be verified that, using Λ as defined above obtains the following state variables. The utility cost γ^h is replaced by its partial sums Γ^h, and the information rent δ^h is replaced by its partial sums Δ^h

$$\Gamma^h_m \equiv \sum_{s=m}^{n} \gamma^h_s = (p-p_0)\left[-\sum_{\tilde{h}:h\succeq\langle\tilde{h},1\rangle} \frac{\Lambda^{\tilde{h}}_m}{p} + \sum_{\tilde{h}:h\succeq\langle\tilde{h},0\rangle} \frac{\Lambda^{\tilde{h}}_m}{1-p}\right] \tag{4.7}$$

$$\Delta^h_m \equiv \sum_{s=m}^{n} \delta^h_s = \frac{p_0}{p}\sum_{\tilde{h}:h\succeq\langle\tilde{h},1\rangle} \Lambda^{\tilde{h}}_m + \frac{1-p_0}{1-p}\sum_{\tilde{h}:h\succeq\langle\tilde{h},0\rangle} \Lambda^{\tilde{h}}_m\,. \tag{4.8}$$

Γ^h_m captures the marginal cost for the principal for any agent utility generated starting from history h, aggregated over all agent types that shirked at least m times in the past. Δ^h_m captures the marginal cost for the principal for any potential private information gains for the agent over all agent types that shirked at least m times in the past. The dual period return, which we denote $f(n,\Gamma,\Delta,\Lambda)$ is now:

(4.9)

$$\begin{aligned} f(n,\Gamma,\Delta,\Lambda) \equiv\; & v(p-p_0) - c_n\cdot\Lambda_0 + c_n\cdot\Gamma_0 - d_n\cdot\Delta_0 \\ & + \sum_{m=1}^{n_h}\left[d_{n-m+1}(\Lambda_m - \Gamma_m) + (d_{n-m+1} - d_{n-m})\Delta_m\right] \quad . \end{aligned}$$

The first line of (4.9) is exactly the same as the period return $\mu(n,\gamma,\delta,\lambda)$ in 3.11 for $s = 0$, with the single period variables replaced by the summation-by-parts variables.

The second line captures the differences when the agent considers the effect of the $m+1$ shirk, rather than the m'th shirk in a period before n. The first term in the summation captures the lower cost for an agent that shirked at least m times in the past. For example, an agent that shirked once in the past will pay c_{n-1} in effort cost in period n rather than c_n. Thus, the second shirk saves the agent d_n or $(c_n - c_{n-1})$ less. The second term in the summation captures the lower cost reduction gains for an agent that shirked at least m times in the past. For example, an agent that already shirked once and is considering shirking again to save costs in period n only saves d_{n-1} from his second shirk, rather than the d_n saved by the first shirk.

The resulting dynamic dual is a direct transformation of problem 3.11, using the summation-by-parts variables instead of the single period variables, with the addition of the non-negativity constraint (4.4):

$$
\begin{aligned}
(4.10)\qquad F\left(n,\vec{\Gamma},\vec{\Delta}\right) &= \max\left[0, \min_{\vec{\Lambda}\geq 0} \hat{F}\left(n,\vec{\Gamma},\vec{\Delta},\vec{\Lambda}\right)\right] \\
s.t.& \\
\hat{F}\left(n,\vec{\Gamma},\vec{\Delta},\vec{\Lambda}\right) &= pF\left(n+1,\vec{\Gamma}^1,\vec{\Delta}^1\right) \\
&\quad + (1-p)\,F\left(n+1,\vec{\Gamma}^0,\vec{\Delta}^0\right) \\
&\quad + f\left(n,\vec{\Gamma},\vec{\Delta},\vec{\Lambda}\right) \\
s.t.& \\
&\quad (p-p_0)\,\Lambda_0 \leq p\,(1+\Gamma_0) \\
&\quad \Lambda_m - \Lambda_{m+1} \geq 0 \\
&\quad \vec{\Gamma}^1 = \vec{\Gamma} - \frac{p-p_0}{p}\vec{\Lambda} \quad ; \quad \vec{\Gamma}^0 = \vec{\Gamma} + \frac{p-p_0}{1-p}\vec{\Lambda} \\
&\quad \vec{\Delta}^1 = \vec{\Delta} + \frac{p_0}{p}\vec{\Lambda} \quad ; \quad \vec{\Delta}^0 = \vec{\Delta} + \frac{1-p_0}{1-p}\vec{\Lambda}
\end{aligned}
$$

LEMMA 4 *For every h, $\mu^h = F\left(n_h, \Gamma^h, \Delta^h\right)$. In particular, $F(1,0,0) = \max_{U,D\geq 0} V(1,1,U,D)$ and the corresponding optimal Λ's define the shadow variables λ in the solution to problem (3.11).*

PROOF: Problem 4.10 is a summation by parts transformation of problem 3.11. Proposition 1 applies. *Q.E.D.*

The dynamic problem 4.10 provides the value of the sub-game starting at each history, as a function of the partial sums state variables – the cost of providing utility

(Γ) and the information rent (Δ), if the agent shirked at most m times.

The dynamic dual has some desirable properties that allow proving stronger results than usual for dynamic contracting problems, including the sufficiency of local deviations that follows.

PROPOSITION 2 *The following hold for the dual problem 4.10:*

- $F(n, \Gamma, \Delta)$ *is convex in* (Γ, Δ). $\hat{F}(n, \Gamma, \Delta, \Lambda)$ *is convex in* Λ *for every* Γ, Δ.
- $F(n, \Gamma, \Delta)$ *decreases in* Γ_0 *and* Δ_0 *for any* $m > 0$
- $F(n, \Gamma, \Delta)$ *increases in* Γ_m *and* Δ_m *for any* $m > 0$

The first property is the mirror image of the concavity property of the standard dynamic moral hazard problem. The last two properties do not have parallels in the standard formulation. The continuation value for any period is almost never monotonic in the agent's promised utility (the standard state variable). In contrast, the state variables Γ_0, Δ_0 reflect the entire costs and the dual value F reflects the full value of the sub-game starting at the history for the principal. As higher costs reduce the value, the second result is obtained. The final result captures the basic OSD intuition. Because the agent's direct gains from shirking are lower if he deviated more in the past (c_{n-s} and d_{n-s} are decreasing in s), increasing the share of costs that apply only if the agent shirked already in the past increases the history's value.

4.3. *Proof for Sufficiency Of Local Deviations*

This section proves that in the solution for the problem $F(1, 0, 0)$, $\Lambda_m = 0$ for all $m > 0$. Thus, corollary 2 applies and local deviations are sufficient.

THEOREM 1 *Any optimal contract subject to LDIC is an optimal contract.*

The detailed proof is in appendix A.7. The following sketch identifies the economics underlying the result and the main technical steps.

By corollary 2 , it is sufficient to show that $\Lambda_m = 0$ for every $m > 0$. Observe that each Λ_m has three effects on $F(\cdot)$ – the period return $f(\cdot)$, the law of motion for Γ and the law of motion for Δ. The proof considers each of these separately and shows that all effects choose the lowest possible Λ_m.

Consider a relaxed problem that allows, for every $m > 0$ to choose separately $\Lambda_m^f, \Lambda_m^\Gamma, \Lambda_m^\Delta$ (all non-negative) such that Λ_m^f affects the period return term $f(\cdot)$, Λ_m^Γ affects only the law of motion for Γ , and Λ_m^Δ affects only the law of motion for Δ.

Observing equation 4.9, $f(\cdot)$ increases in Λ_m^f for any $m > 0$. This reflects the main intuition discussed at the start of the section – if the agent shirked in the past, his costs are lower and his future gains from shirking are lower. Thus, $\Lambda_m^f = 0$ must be optimal.

Proposition 2 states that for any $m > 0$, $F(\cdot)$ is increasing in the information costs Δ_m. This again is a result of convex costs. The earliest shirk generates the biggest – and most costly – information rent. As the dual problem finds the "worst" feasible outcome, it is optimal to pile all possible information rent costs on $m = 0$ and so for $m > 0$, $\Lambda_m^\Delta = 0$ must be optimal.

Finally, increasing Λ has no effect on the *expected* utility cost continuations, Γ_m, but only on the difference between continuation after success or failure. As F is convex, this increases the continuation value and thus is sub-optimal.

It remains to show that increasing Λ_0 does not create some interaction effect that counters any of the previous three intuitions. The only possible concern is with in the third argument (for Λ_m^Γ). However, the concern is valid only if as the agent is provided more incentives (Λ_0 increases), the convexity of the continuation values decreases. This is proved to be false using monotone comparative statics.

5. THE OPTIMAL CONTRACT

5.1. *The Optimal Contract Problems*

Theorem 1 implies that the optimal contract can be derived considering only the LDIC. Both the standard problem (3.1 or 3.3) and the dual problem (3.11) can be simplified if only LDIC are considered. Both are useful to characterizing the optimal contract.

The standard (*primal*) LDIC problem is:

$$
\begin{aligned}
V(n,\overline{q},U,D) &= \max_{(U^y,D^y,q,w)\geq 0} v\left(n,q,w,U^1,U^0,D^1,D^0\right) \\
\text{subject to} & \\
v\left(n,q,w,U^1,U^0,D^1,D^0\right) &= q(p-p_0)v - p\cdot w + pV\left(n+1,q,U^1,D^1\right) \\
&\quad + (1-p)V\left(n+1,q,U^0,D^0\right) \\
\text{Probability const.}\quad (\mu): &\quad q \leq \overline{q} \\
LDIC \quad (\lambda): &\quad w(p-p_0) - q\cdot c_n + (p-p_0)\left(U^1-U^0\right) - p_0 D^1 - (1-p_0)D^0 \geq 0 \\
\text{Regeneration -}U \quad (\gamma): &\quad U = wp - qc_n + pU^1 + (1-p)U^0 \\
\text{Regeneration -}D \quad (\delta): &\quad D = qd_n + pD^1 + (1-p)D^0
\end{aligned}
\tag{5.1}
$$

The *dual* LDIC problem is:

$$
\begin{aligned}
\mu(n,\gamma,\delta) &= \max\left[0,\ \min\nolimits_{\lambda\geq 0}\mu(n,\gamma,\delta,\lambda)\ \right] \\
& s.t. \\
\mu(n,\gamma,\delta,\lambda) &= v(p-p_0) - c_n\lambda + c_n\gamma - d_n\delta \\
&\quad + p\mu\left(n+1,\gamma^1,\delta^1\right) + (1-p)\mu\left(n+1,\gamma^0,\delta^0\right)
\end{aligned}
\tag{5.2}
$$

$$
\begin{aligned}
\text{wage constraint:} &\quad (p-p_0)\lambda \leq p(1+\gamma) \\
\text{utility cost:} &\quad \gamma^1 = \gamma - \frac{p-p_0}{p}\lambda \quad ; \quad \gamma^0 = \gamma + \frac{p-p_0}{1-p}\lambda \\
\text{information rent:} &\quad \delta^1 = \delta + \frac{p_0}{p}\lambda \quad ; \quad \delta^0 = \delta + \frac{1-p_0}{1-p}\lambda \\
\text{stopping condition:} &\quad \mu\left(N^{FB}+1,\gamma,\delta\right) = 0
\end{aligned}
$$

The following "standard" results allow a qualitative comparison between the standard and dual problems:

Lemma 5 *Problem 5.2 is convex in (γ,δ). The optimal solution decreases in the costs (γ and δ), γ,δ are substitutes. If $\mu(n,\gamma,\delta,\lambda) < 0$, the contract is terminated.*

The dual problem has attractive features that the primal problem lacks. Because the dual state variables are "costs", an increase in costs is always bad, and the problem is monotonic. In addition, as one cost increases, the future becomes less attractive, and the other cost becomes less important. Formally, $\mu(n,\gamma,\delta)$ is sub-modular in

$(-\gamma, \delta)$.[11]

As is well known, the standard continuation value may increase or decrease in the promised utility to the agent: if utility is too low, the agent can't be asked to work, and if utility is too high, the agent must be provided all the remaining surplus.

In the standard model, the continuation value (V) does not account for the effect of the continuation on previous profits. As a result, it may very well have been optimal to commit to continue the contract despite a negative V, in order to increase profits in earlier periods. The dual value of a history, however, captures all the effects of the history on the ex-ante profit. Therefore, if it is negative, the contract terminates. Observing that the state variables γ and δ are both higher in the continuation after failure after success, obtains the following simple corollary, that is difficult to obtain using standard tools:

LEMMA 6 *If the contract terminates after success, it terminates after failure.*

PROOF: It is sufficient to show that μ after failure is no larger than μ after success. By lemma 5, this is the case if for every λ, $\gamma^0 \geq \gamma^1$ and $\delta^0 \geq \delta^1$. The first follows directly from $\lambda \geq 0 : \gamma^0 \geq \gamma \geq \gamma^1$. For the second, as $p > p_0$, we have that $\frac{1-p_0}{1-p} > 1 > \frac{p_0}{p}$ and so $\delta^0 \geq \delta^1$. *Q.E.D.*

Note that the dynamic dual can also be derived for the "standard" problem in which costs are known (changing with n or fixed). The dual procedure can then be applied and the result is the same problem without the δ variables. Thus lemma 5 (for γ only) and lemma 6, apply also to the standard case.

5.2. *Dynamic Quotas*

This section shows the main characterization of the contract. We start with a common tradeoff in dynamic moral hazard problems – paying in wages vs. paying in continuation utilities. The dual analysis provides a simple proof for a general result (note that the result also applies to the case that costs are known):

PROPOSITION 3 *If the agent is ever paid in a period, the work plan (q) and wage (w) in all periods after the payment do not depend on future outcomes.*

[11] As the problem is convex, single crossing results rely on sub-modularity rather than super-modularity (which is used for the case that problems are concave).

Proof: If the agent is paid in a period, $(w > 0)$, the dual wage constraint in that period must bind:

$$\lambda = \frac{p + p\gamma}{p - p_0}$$

Placing this in the continuation value for γ^1 yields:

$$\gamma^1 = \gamma - \frac{p - p_0}{p}\lambda = \gamma - \left(\frac{p + p\gamma}{p}\right) = -1$$

Thus, in the next period after payment, the wage constraint is

$$(p - p_0)\,\lambda \leq 0$$

As $\lambda \geq 0$, this constraint can be satisfied only by setting $\lambda = 0$ which implies that the dual state variables in all continuation periods are fixed and cannot depend on additional outcomes, and so the contract cannot depend on these as well. *Q.E.D.*

Proposition 3 clarifies the tradeoff between incentivizing the agent through continuation utilities (conditioning the contract terms on future outcomes) and payments. The dual utility cost γ, makes the distinction explicit. If the contract decides to pay, it must be that the utility cost in all future periods after payment is (-1). That is, all the utility given to the agent after payment was used to generate some work and is essentially "free" to the principal. The cost of paying the agent one util today is exactly compensated by the incentives this util generated for "free" in previous periods. Thus, the most profitable way to incentivize work from this point onwards is by simple payments.

Another implication of proposition 3 is that at the start of the contract, the agent may be rewarded *only* through the effect of outcomes on continuation utilities. Once the contract chooses to reward the agent via payment, there is no going back – all future rewards are solely provided through payments. [12]

The dual value (μ) for all remaining periods after payment is obtained by placing $\gamma = -1$ and $\lambda = 0$ in the dual objective:

$$\mu\,(n, -1, \delta) = \max\left[v\,(p - p_0) - c_n - \delta d_n + \mu\,(n + 1, -1, \delta)\,, 0\right]$$

[12] The distinction is stark because the agent is risk neutral.

The dual value for a period is the efficient value $v(p-p_0)-c_n$ less the private information cost $\delta \cdot d_n$.

If the costs are public information ($\delta = 0$), the optimal dynamic contract therefore eventually either fires the agent without pay or "sells the firm" to the agent. Here, "selling the firm" means having the agent work in all remaining periods until the first best. It is easy to see that, if costs are public information, the contract can be implemented by paying the agent $\frac{c_n}{p-p_0}$ for all remaining periods until the first best period N^{FB}. At this last, first-best period, $c_{N^{FB}} = v(p-p_0)$ the agent's utility is exactly the principal's fixed cost[13], and total surplus is zero. The dual value of the contract is exactly all the remaining surplus. That is, *the dual value of the contract is highest after a payment was made.*

However, if costs are private information, "selling the firm" creates the additional incentive to shirk. The contract after payment does not depend on *future* outcomes, but does depend on *previous* outcomes.

PROPOSITION 4 *If the agent is ever paid in a period, the contract in all periods after the payment asks the agent to work until period $N(\delta)$ and pays the agent $\frac{c_{N(\delta)}}{p-p_0}$ for all remaining successes. δ is the information rent cost in the first period after payment and $N(\delta)$ is given by*

$$N(\delta) = \max n : \quad v(p-p_0) - c_n - \delta \cdot d_n \geq 0 \quad .$$

The dual value to the principal is the remaining surplus less the information rent cost $\delta \cdot \left(c_{N(\delta)} - c_{n-1}\right)$

PROOF: The stopping period $N(\delta)$ is derived as the last period in which μ is still positive. To determine the wage, it is easiest to use the LDIC in problem 5.1. As the continuation utilities and work plans are fixed $U^1 = U^0$ and $D^1 = D^0 = q\left(c_{N(\delta)} - c_n\right)$.

The LDIC then simplifies to

$$w(p-p_0) - q \cdot c_n - q\left(c_{N(\delta)} - c_n\right) \geq 0$$

13

$$p\frac{v(p-p_0)}{p-p_0} - v(p-p_0) = vp_0$$

Or simply

$$\frac{w}{q} = \frac{c_{N(\delta)}}{p - p_0} \ .$$

The agent's wage for success in a period in which he works is exactly $\frac{w}{q}$. *Q.E.D.*

Proposition 4 describes the optimal contract as a "dynamic quota". Once the agent meets a goal (his "quota"), he gets a fixed linear rate on all remaining sales. The quota here is dynamic because the goal and the expected linear rate potentially change until the agent does meet his quota.

The optimal contract therefore "ratchets" incentives. An agent that 'proves' his capabilities is asked to work more. Ratcheting has often been considered a misguided approach to incentives (see e.g. Weitzman (1980)). However, while sub-optimal ratcheting may well be worse than a fixed contract, the contract identified here is an optimal ratcheting mechanism that outperforms the best fixed contract.

If costs are public information, the optimal contract will, in some realizations be ex-post efficient – the agent will work the first best number of periods. However, if costs are private the contract accumulates "private information costs" even if the agent succeeds in all periods. As a result, the contract must terminate before the first-best. The next proposition identifies the most efficient ex-post realization:

PROPOSITION 5 *In the optimal contract, the agent never works more for more than $\overline{N}$ periods.*

$$\overline{N} \quad = \quad \max n : \quad v\left(p - p_0\right) \geq c_n + d_n \frac{p_0}{p - p_0}$$

The agent works for exactly $\overline{N}$ periods if he never fails before the first payment. In particular, if $v\left(p - p_0\right) < c_{N^{FB}} + d_{N^{FB}} \frac{p_0}{p - p_0}$ then the optimal contract is never ex-post efficient.

PROOF: The linear contract with the highest $N(\delta)$ is the linear contract with the lowest possible information rent δ. As $\frac{p_0}{p} < \frac{1-p_0}{1-p}$, the lowest δ for any specific sequence of $\lambda' s$ is obtained if the agent constantly succeeds:

$$\delta^h \geq \frac{p_0}{p} \sum_{h \succeq \tilde{h}} \lambda^{\tilde{h}}$$

The inequality is an equality if the agent never failed in the past.

At the start of the contract $\gamma = 0$. For the contract to move to the linear rate it must be that $\gamma^h = -1$. By the law of motion for γ, this implies that the sequence of past λ must satisfy at least

$$\frac{p - p_0}{p} \sum_{h \succeq \tilde{h}} \lambda^{\tilde{h}} \geq 1 \ .$$

The inequality is an equality if the agent never failed in the past. In this case

$$\sum_{h \succeq \tilde{h}} \lambda^{\tilde{h}} = \frac{p}{p - p_0} \ ,$$

and the lowest possible value for δ^h is:

$$\delta^h = \frac{p_0}{p} \frac{p}{p - p_0} = \frac{p_0}{p - p_0} \ .$$

Placing this in $N(\delta)$ yields the desired result. *Q.E.D.*

Another interesting result is that "second chances" are always worse for the agent. Suppose that the agent can make his quota by succeeding in history h. That is, he would be paid for success in history h, but was not paid in any history that precedes h. If the agent fails, he may still meet his quota in a future history. However, the next result shows that the agent can never be ex-post better off from succeeding in such second chances, compared to succeeding at first. That is:

LEMMA 7 *If the agent is paid for a success in history h, then any linear rate in any history that follows a failure in h is lower than the linear rate that starts in $\langle h, 1\rangle$. In particular, if the agent is paid for success in history h and works in any history $\left\langle h, 0, \tilde{h} \right\rangle$ then he works in the history $\left\langle h, 1, \tilde{h} \right\rangle$.*

PROOF: By construction, $\gamma = -1$ in all histories of the form $\left\langle h, 1, \tilde{h} \right\rangle$, and this is the lowest possible value for γ. By the law of motion for δ,

$$\delta^{\left\langle h,1,\tilde{h} \right\rangle} = \delta^{\langle h,1\rangle} < \delta^{\langle h,0\rangle} \leq \delta^{\left\langle h,0,\tilde{h} \right\rangle}$$

As μ decreases in both state variables, for any history $\left\langle h, 0, \tilde{h} \right\rangle$, $\mu^{\left\langle h,1,\tilde{h} \right\rangle} > \mu^{\left\langle h,0,\tilde{h} \right\rangle}$. Therefore, if the agent works $\left\langle h, 0, \tilde{h} \right\rangle$ and would have been paid for success in history

h, he must also work in history $\left\langle h, 1, \tilde{h} \right\rangle$. As the linear rate is set by the latest period of work, it cannot be larger after a failure than after success. *Q.E.D.*

Lemma 8 identifies another simple result that follows from proposition 3:

LEMMA 8 *The IC always binds in the optimal contract.*

PROOF: In all periods in which the agent is paid, reducing the wage will increase profits and therefore the IC must bind. If the agent is not paid, it must be that the wage constraint does not bind. By complementary slackness the IC does not bind only if $\lambda = 0$. If the optimal contract sets $\lambda = 0$ the state variables in the next period are the same in all outcomes. Thus, the continuation contract and the agent's expected continuation utility from the next period is the same regardless of outcomes. As w_h is zero, the agent is not paid for success nor rewarded for his success in the future in any way. Therefore, the agent's optimal plan must be to shirk in this period, contradicting incentive compatibility. *Q.E.D.*

5.3. *Implementation Examples*

I am not aware of explicit implementations of this theoretic mechanism. However, as the quote from Prendergast (1999) in the introduction suggests, current theory is challenged to explain frequently used mechanisms such as quotas and convex reward schemes. The following examples show that the analysis so far can provide a microeconomic foundation for these.[14]

The first example shows that a twist on a simple quota contract implements the optimal mechanism and outperforms the optimal linear contract by about 40%. Quota contracts with different variants are popular in sales organizations (see e.g. Misra and Nair (2009) and Joseph and Kalwani (1998)).

EXAMPLE 1 Suppose costs are $c = \{8, 10, 12, 14, 16, 18, 20, 40\}$, $v = 100$, $p = \frac{1}{2}$ and $p_0 = \frac{1}{10}$. The optimal contract pays the agent 50 $\left(= \frac{20}{p - p_0}\right)$ for any success *except* the first. The agent's optimal plan is to stop if and only if he fails in all the first three periods. The principal's expected profit is 141.25 and the agent's expected utility is 29.25. The only difference between this contract and a regular quota contract is for the case that the agent decides to stop (i.e. after failing in the first three periods). The

[14]The code for the simulations is available online at the author's website: www.guyarie.com

optimal contract must terminate. In a regular quota contract, if the agent happens to succeed without effort, say in period four, he will resume working.

There are various ways to implement this modified quota. A simple method is to require that the quota is met within some deadline. A more complicated method occurs in sales situations in which the agent's success depends in part on getting "leads" from the principal (as in the popular David Mamet's play Glengarry Glen Ross). In these settings the costs increase over the selling cycle (month or quarter) in which the leads are used. At the end of the selling cycle the process restarts. An easy implementation is to set the agent the quota as here. However, if the agent does not accumulate enough successes early on, he does not get any new leads (is "dried up") for the remainder of the selling cycle. In the next selling cycle, everyone starts again with the same contract.

In comparison, if the contract is limited to a linear, fixed "no ratcheting" type, the agent is paid 50 for all successes. This increases ex-ante efficiency (the agent works more in expectation) and the agent's utility from 29.25 to 77. However, the principal's profit reduces from 141.25 to 105. In the fixed contract, the principal's expected profit is 15 per period $\left(100 \cdot \frac{2}{5} - 50 \cdot \frac{1}{2}\right)$. Using a quota mechanism saves the principal the first payment of 50, a net increase in profits as long as the first success happens fast enough. Only if the agent fails three times in a row (probability of $\frac{1}{8}$), the optimal contract does worse than the regular contract, losing the possible gains in the remaining four periods.[15] Quota contracts allow the principal to pay less, risking that in some bad realizations, the agent may be discouraged and quit early on.

The agent's optimal plan is derived from the dual solution – the agent works only in periods in which $\mu \geq 0$. It is also possible to reconstruct the optimal plan using backward induction.

The next example implements the optimal contract using a non-linear (convex) reward scheme. In non-linear schemes, the agent's reward per sale increases (often drastically) with the volume of sales. Joseph and Kalwani (1998) document the use of such schemes, and Larkin (2007) provides a recent example.

EXAMPLE 2 Suppose that the costs are $c = \{4, 6, 8, 10, 20\}$ and that $v = 60$, $p = \frac{1}{2}$ and $p_0 = \frac{1}{4}$. The optimal contract pays the agent nothing for the first success, 24 for the second success and 40 for any later success, terminating the contract after

[15] Indeed, the difference between the two profits is $\frac{7}{8} \cdot 50 - \frac{1}{8} \cdot 15 \cdot 4 = 36.25$.

four periods. This is a standard convex reward scheme. Note that in a single period problem 40 is the required reward if cost is 10 (c_4) and 24 is the required reward if cost is 6 (c_2).

If the agent fails in the first period, he stops. However, if the agent succeeds in the first period, he works in all following histories unless he fails in the following two periods, which causes him to stop after three periods.

The principal's expected profit is 14.125 and the agent's expected rent is 6.75. The total surplus is 20.875. The agent is expected to work for a total of 2.375 periods.

In comparison, if the contract is limited to a linear, fixed "no ratcheting" type, the optimal contract has the agent work *only* in the first period. The agent is paid 16 ($= \frac{4}{.25}$) for his first success. The principal's expected profit from work is seven ($60 \times (p - p_0) - 16 \times p$) and the agent's expected rent is four.

This second example illustrates two aspects of the optimal contract. First, the principal's "continuation profit" is negative in most histories in which the agent works. It is positive only in the first period and the third period following a failure in the second period. Thus, commitment plays an important role. Second, the optimal contract more than doubles the principal's profit and the expected work, nearly doubles surplus and *increases* the agent's rent by over 50%.

Implementing a "ratcheting" contract may therefore be a Pareto improvement. Both the agent and the principal are better off. Without ratcheting, the principal has very limited tools to mitigate the private information problem and as a result, production is very limited. By stopping work only after a failure, the ratcheting contract allows the principal to stop production only when the private information problem is most severe.

A more general implementation of the contract can be a pre-specified quota with "adjustments". In the sales example, at the start of the quarter, the agent is offered a quota contract. At any period, the agent can come to the firm and "complain" that the quota contract is too aggressive – requires too many successes. Based on the agent's performance, the firm then reduces the threshold to make the quota, but also reduces the linear rate the agent receives when making the quota. As long as the firm can commit in advance to the adjustments (e.g. have a policy in place), a contract arbitrarily close to the optimal contract may be implemented.

5.4. *Higher Costs vs. Increasing Costs*

A general presumption is that when it comes to costs, less is always better. However, lower costs in early periods generate a private information problem that may be more costly than the efficiency gain. A two period setting is sufficient to see this. If the optimal contract asks the agent to work in the second period *only* after success, the reduction in the difference between c_1 and c_2 increases the overall expected costs of effort.The increase in c_1 is paid in all possible realizations, but the second period cost decrease happens ex-post only with probability p. However, this also reduces the information problem. Proposition 6 shows that, under certain conditions, this increase in expected costs increases the contract's expected profit.

PROPOSITION 6 *In a two period problem, for any* $\varepsilon \in \left(0, \frac{c_2 - c_1}{2}\right)$ *consider increasing* c_1 *by* ε *and decreasing* c_2 *by* ε*. The change strictly increases expected profits if and only if:*

- *The optimal contract asks the agent to work after failing in the first period; or*
- *The optimal contract asks the agent to work only after success in the first period and* $p + p_0 > 1$

The key to the proof is the following lemma, which provides a closed form solution to *any* two period problem:

LEMMA 9 *If the contract must terminate within at most two periods, the wage constraint binds in both remaining periods. In a two period problem, the dual values* $\mu^{\emptyset}, \mu^0$ *and* μ^1 *are given by:*

$$\begin{aligned}
\mu^1 &= \max\left[0, v\left(p - p_0\right) - c_2 - \frac{p_0}{p - p_0}\left(c_2 - c_1\right)\right] \\
\mu^0 &= \max\left[0, v\left(p - p_0\right) + c_1 \frac{1 - p_0}{p - p_0} \frac{p}{1 - p} - c_2 \frac{p}{1 - p} \frac{2 - p}{p - p_0}\right] \\
\mu &= v\left(p - p_0\right) - \frac{p}{p - p_0} c_1 + p\mu^1 + \left(1 - p\right)\mu^0
\end{aligned}$$

The only challenging part of lemma 9 is to determine that the wage constraint does indeed bind at all cases. Once that is obtained, the rest is simple algebra. Note that the first part of the lemma applies to *any* last two periods in a multi-period optimal contract.

If the optimal contract requires work in both periods, then by proposition 4, the agent's wage is $\frac{c_2}{p - p_0}$ in both periods and c_1 does not matter for the principal's expected profit. The closed form for μ^0 shows that increasing c_1 and decreasing c_2 only increases μ^0 and so the change proposed in the proposition will not change the optimal work plan.

If the optimal contract requires work in the first period and in the second period only after success

$$\frac{\partial\mu}{\partial c_1} = -\frac{p}{p-p_0} + \frac{pp_0}{p-p_0} = -\frac{p\,(1-p_0)}{p-p_0} \quad \text{and} \quad \frac{\partial\mu}{\partial c_2} = -\frac{p^2}{p-p_0}$$

By the condition $p + p_0 > 1$, $\frac{\partial\mu}{\partial c_2} < \frac{\partial\mu}{\partial c_1}$, which proves proposition 6.

Even though second period utility is used to provide first period incentives *and* work in the second period will only be required ex-post with probability $p < 1$, the firm prefers a reduction in the second period cost to a reduction in the first period costs. As for the previous case, the proposed change only increases μ^1 and so the total work required by the optimal contract as a result of the change cannot decrease.

To see that the effect is only a result of the private information problem, suppose that costs are known. If the optimal contract requires work in both periods regardless of outcomes, the wage in each period is $\frac{c_n}{p-p_0}$ and the profit is not affected by the change considered in the proposition. For the general case, the closed form solution depends on the sign of $c_1 - p_0 c_2$. Here, we assume that $c_1 > p_0 c_2$, in which case the wage constraint always binds in the first period $\lambda^{\emptyset} = \frac{p}{p-p_0}$. The closed form solutions are[16]

$$\begin{aligned} \hat{\mu}^1 &= v\,(p-p_0) - c_2 \geq 0 \\ \hat{\mu}^0 &= \left[v\,(p-p_0) - p\frac{c_2}{p-p_0} - \frac{p_0}{1-p}\frac{p-p_0}{p}c_2 \right]^+ \end{aligned}$$

Focusing on the case that the optimal contract does not require work after failure,

$$\hat{\mu} = \left[v\,(p-p_0)\,(1+p) - p\frac{c_1}{p-p_0} - pc_2 \right]^+$$

Then

$$\frac{\partial\hat{\mu}}{\partial c_1} = -\frac{1}{p-p_0} < -1 < -p = \frac{\partial\hat{\mu}}{\partial c_2}$$

The principal strictly gains by making the two costs more similar, exactly in contrast

[16] The 'hat' is used to identify these as the values with publicly known costs.

to the case of private costs.

5.5. *Discounting*

The standard procedure to introducing a common discount factor β to problem 3.3 is by multiplying all the continuation terms (in the objective, the FDIC and the regeneration constraints). However, since no payments are deferred across periods, an equivalent, and much simpler way is to change the probability constraint to

$$\beta \cdot \overline{q} \leq q$$

Then, q_h is the "discounted probability" that the contract was not terminated. It is then immediate that the only part of the analysis that can possibly change (other than the specific examples) is the required payment identified in proposition 4. Discounting "dilutes" the agent's information rent and so the agent can be paid less to forgo it. The next lemma identifies the solution.

LEMMA 10 *If the principal and agent have a common discount factor β, then after the first payment was made, the contract is fixed to $N(\delta)$ as in proposition 4 and the payment in period n is defined by*

$$\begin{aligned} W_{N(\delta)} &= \frac{c_{N(\delta)}}{p-p_0} \\ W_n &= \beta W_{n+1} + (1-\beta)\frac{c_n}{p-p_0}\ . \end{aligned}$$

PROOF: As the contract after payment still does not depend on new outcomes, the relevant IC is still as in the proof of proposition 4:

$$(p-p_0)\, w_h - q_h c_h \geq D^1$$

Discounting means that $q_{\langle h,y\rangle} = \beta \cdot q_h$. Thus, at any period n, for a contract that will end in period $N(\delta)$

$$D^1 = q_h \sum_{m=n+1}^{N(\delta)} \beta^{(m-n)} d_m = q_h \sum_{m=n+1}^{N(\delta)} \beta^{(m-n)} \left(c_m - c_{m-1}\right)$$

So,

$$w_h = \frac{c_h q_h + q_h \sum_{m=n+1}^{N(\delta)} \beta^{(m-n)} (c_m - c_{m-1})}{p - p_0}$$

As $w_h = W_h \cdot q_h$ where W_h is the actual payment for success in the history h, divide both sides by q_h:

$$W_h = \frac{c_h + \sum_{m=n_h+1}^{N(\delta)} \beta^{m-n_h} (c_m - c_{m-1})}{p - p_0} .$$

The result for $W_{N(\delta)}$ is immediate. Now observe that

$$\beta W_{n+1} = \frac{\beta c_{n+1} + \sum_{m=n+2}^{N(\delta)} \beta^{m-n_h} (c_m - c_{m-1})}{p - p_0} .$$

So

$$\begin{aligned} W_n - \beta W_{n+1} &= \frac{c_n + \sum_{m=n_h+1}^{N(\delta)} \beta^{m-n_h} (c_m - c_{m-1}) - \beta c_{n+1} - \sum_{m=n+2}^{N(\delta)} \beta^{m-n_h} (c_m - c_{m-1})}{p - p_0} \\ &= \frac{c_n + \beta c_{n+1} - \beta c_n - \beta c_{n+1}}{p - p_0} \\ &= (1 - \beta) \frac{c_n}{p - p_0} \end{aligned}$$

Q.E.D.

6. CONCLUSION

This paper restored the sufficiency of local deviations in a dynamic moral hazard setting with persistent private information. The optimal contract problem was reformulated based on the two agency frictions – the cost to the principal of providing future utility (γ) and future private information (δ) to the agent.

The resulting optimal contract was characterized as a *dynamic quota.* At the start of the contract the agent is not paid for successes. Once the agent is paid, he is paid a fixed linear piece-rate that depends only on his outcomes prior to the first payment.

The optimal contract explains features of real world contracts that puzzled economic observers. The variance in the expected total effort is larger with private cost information than without. Such large variation in ex-post incentives and effort across agents is inefficient and led several authors (see e.g. Oyer (1998); Larkin (2007); Misra and Nair (2009)) to suggest that there is significant room for improvement in either the design of real world incentives or models of the moral hazard setting. The model shows that this variance allows the firm to provide sufficient incentives for effort when it is relatively cheap and to provide high powered incentives when those are required *without fear that agents misrepresent their effort* (delay "easy sales" to the end of the period). The optimal contract must balance between efficiency (having the agents work longer) and profitability. While a high linear commission would guarantee all agents make the efficient level of effort, the firm's profits would all be provided as rents. Consistent with the model, in the firm documented by Larkin (2007) the top end of the reward scale provides the salesperson a 25% commission on *revenues,* a figure very close to the industry's accounting profit margins.

The analysis used arguments based on the duality of linear programs to design a dynamic program. The duality based analysis allows applying standard mechanism design techniques to the dynamic private information problem. The dynamic dual problem has desirable features – namely monotonicity and single-crossing in the state space. These properties have so far been absent from dynamic moral hazard problems but are generally instrumental in the characterization of economic outcomes and comparative statics.

The use of duality in dynamic moral hazard problem has been advanced recently by Marcet and Marimon (2011) and Mele (2011) as well. Compared to these studies, the work here provides stronger and more direct results for a simpler and more specific framework. This allowed the dual formulation to provide new and important results that cannot be obtained using standard methods. In addition, the dual value of a period and the separation of the utility cost and the private information cost should prove useful for similar problems in future research.

REFERENCES

Ábrahám, Á., S. Koehne, and N. Pavoni (2011): "On the first-order approach in principal–agent models with hidden borrowing and lending," *Journal of Economic Theory*, 146(4), 1331–1361.

Abraham, a., and N. Pavoni (2008): "Efficient allocations with moral hazard and hidden borrowing and lending: A recursive formulation," *Review of Economic Dynamics*, 11(4), 781–803.

Bergemann, D., and U. Hege (2005): "The Financing of Innovation: Learning and Stopping," *The RAND Journal of Economics*, 36(4), 719–752.

Bonatti, A., and J. Hörner (2011): "Collaborating," *American Economic Review*, 101(2), 632–663.

Clementi, G. L., and H. a. Hopenhayn (2006): "A Theory of Financing Constraints and Firm Dynamics," *The Quarterly Journal of Economics*, 121(1), 229–265.

Dantzig, G. B. (1963): *Linear programming and extensions*. Princeton University Press.

DeMarzo, P., and Y. Sannikov (2008): "Learning in dynamic incentive contracts," *Unpublished Manuscript*.

Fernandes, A., and C. Phelan (2000): "A Recursive Formulation for Repeated Agency with History Dependence," *Journal of Economic Theory*, 91(2), 223–247.

Halac, M., N. Kartik, and Q. Liu (2012): "Optimal Contracts for Experimentation," .

Holmstrom, B., and P. Milgrom (1987): "Aggregation and linearity in the provision of intertemporal incentives," *Econometrica: Journal of the Econometric Society*, 55(2), 303–328.

Joseph, K., and M. U. Kalwani (1998): "The Role of Bonus Pay in Salesforce Compensation Plans," *Industrial Marketing Management*, 27(2), 147–159.

Larkin, I. (2007): "The cost of high-powered incentives: Employee gaming in enterprise software sales," *Harvard Business School*.

Marcet, A., and R. Marimon (2011): "Recursive contracts," .

Mele, A. (2011): "Repeated moral hazard and recursive Lagrangeans," .

Misra, S., and H. Nair (2009): "A structural model of sales-force compensation dynamics: Estimation and field implementation," *Quantitative Marketing and Economics*, p. 1–47.

Mukoyama, T., and A. Şahin (2005): "Repeated moral hazard with persistence," *Economic Theory*, 25(4), 831–854.

Myerson, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6(1), 58–73, ArticleType: research-article / Full publication date: Feb., 1981 / Copyright © 1981 INFORMS.

Oyer, P. (1998): "Fiscal Year Ends and Nonlinear Incentive Contracts: The Effect on Business Seasonality," *The Quarterly Journal of Economics*, 113(1), 149–185.

Prendergast, C. (1999): "The provision of incentives in firms," *Journal of economic literature*, 37(1), 7–63.

Rockafellar, R. T. (1970): *Convex analysis*, Princeton Landmarks in Mathematics. Princeton University Press.

Spear, S. E., and S. Srivastava (1987): "On Repeated Moral Hazard with Discounting," *The Review of Economic Studies*, 54(4), 599–617.

Tchistyi, A. (2006): "Security design with correlated hidden cash flows: The optimality of performance pricing," *Unpublished manuscript, New York University*.

Topkis, D. M. (1998): *Supermodularity and complementarity*. Princeton University Press.

Vohra, R. V. (2005): *Advanced Mathematical Economics*. Routlege.

——— (2011): *Mechanism Design: A Linear Programming Approach (Econometric Society Monographs)*. Cambridge University Press.

Weitzman, M. L. (1980): "The "Ratchet Principle" and Performance Incentives," *The Bell Journal of Economics*, 11(1), 302.

Williams, N. (2011): "Persistent private information," *Econometrica*, 79(4), 1233–1275.

APPENDIX A: DETAILED PROOFS AND DERIVATIONS

A.1. *Lemma 1: If a contract is FDIC it is IC*

Proof: Suppose the contract q, w is not IC. Then it violates FDIC:

1. As the set of possible work plans for the agent is finite and the agent's expected profit is well defined and bounded for each work plan given q, w, there is a set $\hat{E}(q, w)$ of most profitable work plans given q, w.
2. Suppose $e^c \notin \hat{E}$ and let $\hat{e} \in \hat{E}$, be a most profitable deviating work plan.
3. Consider the set of histories $\hat{H}$ in which the agent makes a "final deviation" according to $\hat{e}$. That is, $\hat{h} \in \hat{H}$ if $q_{\hat{h}} > 0$, $\hat{e}_{\hat{h}} = 0$ and for every $h \succeq \hat{h}$, $h \neq \hat{h}$, either $q_h = 0$ or $\hat{e}_h = 1$. Let $\hat{s}$ be the number of past deviations at $\hat{s}$ according to $\hat{e}$. Clearly, if the agent profits from making this final deviation, the FDIC for $\hat{h}, \hat{s}$ is violated and the proof is complete.
4. If the agent does not profit from making this final deviation then the effort plan that complies in this last period provides at least the same expected profit to the agent. Thus, the effort plan with $\hat{e}_{\hat{h}} = 1$ provides at least the same expected profit for the agent. We can now repeat the process of searching for a profitable final deviation after setting $\hat{e}_{\hat{h}} = 1$. As H is a finite set, the process ends either in finding a history in which FDIC is violated or if we change all periods in which $\hat{e}_h = 0$ to $\hat{e}_h = 1$ while weakly increasing the agent's expected profit, implying that $\hat{e}$ was not more profitable than e^c.

Q.E.D.

A.2. *Lemma 2: LDIC does not imply IC*

Proof: Suppose $c_n = n$, $p = \frac{1}{2}$ and $p_0 = 0$. We will show that the following contract violates IC but not LDIC:

- If the agent succeeds in any of the first two periods, he is paid 48 and is asked to stop working.
- If the agent fails in both first two periods, he is asked to work for eight more periods regardless of new outcomes and is paid 20 for each success.

First, we calculate U_s^h for any period after the second. The agent is paid 20 for each remaining success and so:

$$U_s^h = \sum_{m=n_h}^{10} \frac{1}{2} \cdot 20 - \sum_{m=n_h}^{10} c_{m-s} = (11 - n_h) \cdot 10 - \sum_{m=n_h}^{10} c_{m-s}$$

In particular, if the agent failed in the first two periods, his expected utility from following the contract is:

$$U_s^{\langle 0,0 \rangle} = 80 - 52 + 8s = 28 + 8s \quad \text{(A.1)}$$

Therefore, if the agent shirks in the first two periods, he is guaranteed to fail in those periods and his expected utility is $U_2^{\langle 0,0\rangle} = 44$.

By complying, the agent's expected utility is

$$U^{\emptyset} = \frac{1}{2} \cdot 48 - 1 + \frac{1}{2}\left(\frac{1}{2} \cdot 48 + \frac{1}{2}U_0^{\langle 0,0\rangle} - 2\right) = \frac{1}{2}48 - 1 + \frac{1}{2}36 = 41$$

Therefore

$$U^{\emptyset} < U_2^{\langle 0,0\rangle}$$

and the contract is not incentive compatible.

It remains to show that the contract does satisfy all LDIC: *Q.E.D.*

- If the agent makes a first and last shirk in history h with $n_h > 2$ he surely fails in that period and is expected utility is $U_1^{\langle h,0\rangle}$. Thus, the LDIC for any h with $n_h > 2$ is

  $$U_0^h \geq U_1^{\langle h,0\rangle} \quad .$$

  Which is:

  $$(11-n) \cdot 10 - \sum_{m=n}^{10} c_m \quad \geq \quad (11-(n+1)) \cdot 10 - \sum_{m=n+1}^{10} c_{m-1}$$

  Simplifying, the LDIC when $n_h > 2$ is

  $$10 \geq \sum_{m=n}^{10} c_m - \sum_{m=n}^{9} c_m = c_{10}$$

  As $c_n = n$, the LDIC binds in all periods starting from $n > 2$.

- Next consider period 2. To work in period 2, the agent must have failed in the first period. The LDIC is

  $$\frac{1}{2} \cdot 48 + \frac{1}{2} \cdot U_0^{\langle 0,0\rangle} - 2 \geq U_1^{\langle 0,0\rangle}$$

  Using equation A.1, $U_0^{\langle 0,0\rangle} = 28$ and $U_1^{\langle 0,0\rangle} = 36$ so the LDIC strictly binds:

  $$U_0^0 = 24 + 14 - 2 = 36$$

- Finally, for period 1, the LDIC is

  $$\frac{1}{2} \cdot 48 + \frac{1}{2} \cdot U_0^0 - 1 \quad \geq \quad U_1^0$$

  $U_0^0 = 36$ was obtained when the second period was considered. Therefore, the LHS of the

LDIC is

$$U^{\emptyset} = \frac{1}{2} \cdot 48 + \frac{1}{2} \cdot 36 - 1 = 41$$

For U_1^0, if the agent shirked in the first period and complies starting from the second period then with probability $\frac{1}{2}$, a complying agent works only in period 2 and so gains just one more util from shirking before and with probability $\frac{1}{2}$, a complying agent works in all remaining periods and so gains nine more utils (one per period) from shirking before. Thus:

$$U_1^0 = U_0^0 + \frac{1}{2} + \frac{1}{2} \cdot 9 = 36 + 5 = 41$$

The LDIC in the first period binds as well.

A.3. *Duality - Main Theorems*

The classic reference is Dantzig (1963). The results are given in current textbooks on static optimization (see e.g. Vohra (2005)). Any linear problem may be written as

$$\max_{x\geq 0} c\cdot x \qquad s.t. \quad Ax \leq b \quad . \tag{A.2}$$

With c a vector of coefficients and A a matrix that holds in each row the coefficients on a constraint. The dual of the problem is

$$\min_{y\geq 0} y\cdot b \qquad s.t. \quad yA \geq c \tag{A.3}$$

The main results of interest are:

1. Each primal variable (x) translates to a constraint in the dual problem. Each primal constraint translates to a dual variable (y)
2. The Duality Theorem: If x^* and y^* are optimal, $y^*\cdot b = c\cdot x^*$ whenever both exist and are finite; and
3. Complementary Slackness: y_i^* is the Lagrange multiplier in the primal solution for the constraint associated with the i-th row in A . If $y_i^* = 0$ then the constraint associated with the i-th row in A does not bind when solving the primal.
4. The Dual of the Dual is the primal. Therefore, the primal variables x_i^* are the Lagrange multipliers in the dual's solution.

The linearity of the objective implies:

1. If $y_i^* = 0$ then the solution to problem A.2 is not changed if the constraint associated with the i-th row in A is removed.

This last result is a combination of the Complementary Slackness result and the Fundamental Theorem of Linear Programming. See e.g. the discussion in Vohra (2005) preceding theorem 4.10 (Complementary Slackness).

A.4. *Proposition 1 – Problem 3.11 is the dual of problem 3.3.*

The proof first reconstructs problem 3.3 as a linear problem. Then derives the dual of the linear problem and finally reconstructs the dual as a recursive problem.

LEMMA 11 *Problem A.4 identifies the optimal contract:*

$$V^{FD} = \max_{q\geq 0, w\geq 0} \sum_{h\in H} P_h \left[q_h \left(p-p_0\right) v - w_h p\right] \tag{A.4}$$

$s.t.$

$$\begin{array}{lll} & q_\emptyset \leq 1 & \mu^\emptyset \\ \forall h, y \in \{0,1\} & P_{\langle h,y\rangle} q_{\langle h,y\rangle} - P_{\langle h,y\rangle} q_h \leq 0 & \mu^{\langle h,y\rangle} \\ \forall h, s & FDIC\ (h,s) & \lambda_s^h \end{array}$$

With the FDIC:

$$\begin{aligned} FDIC\ (h,s) : \quad & -P_h \left(p-p_0\right) w_h + P_h q_h c_{h-s} \\ & - \left(p-p_0\right) \left(\sum_{\tilde{h}\succeq\langle h,1\rangle} \tfrac{P_{\tilde{h}}}{p} p w_{\tilde{h}} - \sum_{\tilde{h}\succeq\langle h,1\rangle} \tfrac{P_{\tilde{h}}}{p} q_{\tilde{h}} \cdot c_{\tilde{h}-s}\right) \\ & + \left(p-p_0\right) \left(\sum_{\tilde{h}\succeq\langle h,0\rangle} \tfrac{P_{\tilde{h}}}{1-p} p w_{\tilde{h}} - \sum_{\tilde{h}\succeq\langle h,0\rangle} \tfrac{P_{\tilde{h}}}{1-p} q_{\tilde{h}} c_{\tilde{h}-s}\right) \\ & + p_0 \sum_{\tilde{h}\succeq\langle h,1\rangle} \tfrac{P_{\tilde{h}}}{p} q_{\tilde{h}} \cdot d_{\tilde{h}-s} \\ & + \left(1-p_0\right) \sum_{\tilde{h}\succeq\langle h,0\rangle} \tfrac{P_{\tilde{h}}}{1-p} q_{\tilde{h}} \cdot d_{\tilde{h}-s} \leq 0 \quad . \end{aligned} \tag{A.5}$$

PROOF: The objective is the same as 2.3. The probability constraint is the same as in problem 3.3, with both sides multiplied by P_h which is strictly positive. For the FDIC, start from the FDIC (3.2):

$$w\left(p-p_0\right) - q \cdot c_{n-s} + \left(p-p_0\right)\left(U_s^1 - U_s^0\right) - p_0 D_s^1 - \left(1-p_0\right) D_s^0 \geq 0 \tag{A.6}$$

By construction, U_s^h is the agent's continuation utility starting at the private history (h,s) is the agent complies in all remaining histories. This yields:

$$\begin{aligned} U_s^h &\equiv \sum_{\tilde{h}\succeq h} \frac{P_{\tilde{h}}}{P_h} \left[p w_{\tilde{h}} - q_{\tilde{h}} c_{\tilde{h}-s}\right] \\ D_s^h &\equiv U_{s+1}^h - U_s^h = \sum_{\tilde{h}\in H} \frac{P_{\tilde{h}}}{P_h} \left[q_{\tilde{h}} d_{\tilde{h}-s}\right] \end{aligned}$$

Then, for example

$$\begin{aligned} U_s^{\langle h,1\rangle} &= \sum_{\tilde{h}\succeq\langle h,1\rangle} \frac{P_{\tilde{h}}}{p\cdot P_h} \left[p w_{\tilde{h}} - q_{\tilde{h}} c_{\tilde{h}-s}\right] \quad \text{and} \\ D_s^{\langle h,1\rangle} &= \sum_{\tilde{h}\succeq\langle h,1\rangle} \frac{P_{\tilde{h}}}{p\cdot P_h} q_{\tilde{h}} d_{\tilde{h}-s} \end{aligned}$$

Apply these to all continuation values in A.6 and multiplying both sides by P_h gets the FDIC A.5. *Q.E.D.*

LEMMA 12 *Problem A.7 is the dual of problem A.4.*

$$\text{(A.7)} \quad \begin{aligned} &\min_{(\mu,\lambda)\geq 0} && \mu^{\emptyset} \\ &s.t. \quad \forall h: && \\ \text{Wage } (w_h) \;\; &\text{constraint:} && (p-p_0)\sum_{s=0}^{n_h-1}\lambda_s^h \leq p + p\cdot\sum_{s=0}^{n_h-1}\gamma_s^h \\ q_h &\text{constraint:} && \mu^h \geq v\,(p-p_0) + (1-p)\,\mu^{\langle h,0\rangle} + p\mu^{\langle h,1\rangle} \\ & && +\sum_{s=0}^{n_h-1}\left(\gamma_s^h\cdot c_{h-s} - \lambda_s^h\cdot c_{h-s} - \delta_s^h\cdot d_{h-s}\right) \\ \text{stopping } &\text{condition:} && n_h > N^{FB} \implies \mu^h = 0 \end{aligned}$$

PROOF: The objective $\mu^{\emptyset}$ is straightforward as it is the multiplier on the only constraint with non-zero RHS. The stopping condition follows as $q_h = 0$ is the optimal solution whenever $n_h > N^{FB}$ and so the probability constraint cannot bind. It remains to derive the constraints for each w_h and q_h.

- For $\underline{w_h}$, we show that the constraint is

$$-(p-p_0)\sum_{s=0}^{n_h-1}\lambda_s^h + p\cdot\sum_{s=0}^{n_h-1}\gamma_s^h \geq -p$$

  - The right hand side of the constraint is the coefficient on w_h in the objective of V^{FD}: $-P_h p$.
  - w_h appears with a coefficient $-P_h\,(p-p_0)$ in all the FDIC (A.5) for h. Each has a shadow variable λ_s^h. Summing obtains the first element of the left hand side

$$-P_h\,(p-p_0)\sum_{s=0}^{n_h-1}\lambda_s^h$$

  - The variable w_h also appears with a coefficient $-P_h\frac{p}{p}\,(p-p_0)$ in all FDIC (A.5) for $\hat{h}$ such that $h \succeq \left\langle \hat{h},1\right\rangle$ and with a coefficient $P_h\frac{p}{1-p}\,(p-p_0)$ in all FDIC for $\hat{h}$ such that $h \succeq \left\langle \hat{h},0\right\rangle$.[17] Taking p out and summing up, this yields:

$$P_h p\cdot\sum_{s=0}^{n_h}(p-p_0)\left[-\sum_{\tilde{h}:h\succeq\left\langle\tilde{h},1\right\rangle}\frac{\lambda_s^{\tilde{h}}}{p} + \sum_{\tilde{h}:h\succeq\left\langle\tilde{h},0\right\rangle}\frac{\lambda_s^{\tilde{h}}}{1-p}\right] \quad = P_h p\cdot\sum_{s=0}^{n_h-1}\gamma_s^h$$

- For q_h, we show that the constraint is

$$P_h\cdot\left[\mu^h - (1-p)\,\mu^{\langle h,0\rangle} - p\mu^{\langle h,1\rangle} + \sum_{s=0}^{n_h-1}\left(\lambda_s^h\cdot c_{h-s} - \gamma_s^h\cdot c_{h-s} + \delta_s^h\cdot d_{h-s}\right)\right] \geq P_h v\,(p-p_0)$$

[17] Note that the history h here is $\tilde{h}$ in those relevant FDIC (A.5) and the history h in the FDICs is the history $\hat{h}$ here.

Then , dividing both sides by P_h and isolating μ^h obtains the result.

- The variable q_h appears in the objective with coefficient $P_h v(p - p_0)$, obtaining the right hand side.
- q_h appears in the three probability constraints . These generate the first three terms in the left hand side

$$P_h \mu^h - P_h (1-p) \mu^{\langle h,0\rangle} - P_h p \mu^{\langle h,1\rangle} \tag{A.8}$$

- In all the FDIC for history h (i.e., for each s), q_h appears with a coefficient c_{h-s} . This generates the term

$$\sum_{s=0}^{n_h} P_h c_{h-s} \lambda_s^h \ .$$

- The variable q_h also appears twice in each of the FDIC for $\tilde{h}, s$ such that $h \succeq \tilde{h}$, once as part of the continuation utility term (in either the second or third row of A.5) and once as part of the future gains from shirking term (in either the fourth or fifth row of A.5). The continuation utility term will determine the coefficients on γ. The shirking gains term will determine the coefficients on δ.

  1. In the continuation utility term in the FDIC for all histories that h follows, the coefficient for $\lambda_s^{\tilde{h}}$ is multiplied by the current cost (c_{h-s}) and $\frac{p-p_0}{p}$ if $h \succeq \left\langle \tilde{h}, 1 \right\rangle$ or $-\frac{p-p_0}{1-p}$ if $h \succeq \left\langle \tilde{h}, 0 \right\rangle$. Summarizing these terms obtains the sum:

     $$P_h \sum_{s=0}^{n_h} c_{h-s} (p - p_0) \left[\sum_{\tilde{h}: h \succeq \left\langle \tilde{h}, 1 \right\rangle} \frac{\lambda_s^{\tilde{h}}}{p} - \sum_{\tilde{h}: h \succeq \left\langle \tilde{h}, 0 \right\rangle} \frac{\lambda_s^{\tilde{h}}}{1-p} \right] \ . \tag{A.9}$$

     Simple algebra yields the term

     $$-P_h \sum_{s=0}^{n_h - 1} \gamma_s^h c_{h-s}$$

  2. In the information rents terms in each FDIC constraint (the last two rows in A.5), the coefficient for $\lambda_s^{\tilde{h}}$ is $P_h \frac{p_0}{p} d_{h-s}$ if $h \succeq \left\langle \tilde{h}, 1 \right\rangle$ and $\frac{1-p_0}{1-p} P_h d_{h-s}$ if $h \succeq \left\langle \tilde{h}, 0 \right\rangle$. Again the coefficients depend only on whether h follows a success or failure in $\tilde{h}$. Adding up all the relevant terms obtains:

     $$P_h \sum_{s=0}^{n_h - 1} d_{h-s} \left[\frac{p_0}{p} \sum_{\tilde{h}: h \succeq \left\langle \tilde{h}, 1 \right\rangle} \lambda_s^{\tilde{h}} + \frac{1-p_0}{1-p} \sum_{\tilde{h}: h \succeq \left\langle \tilde{h}, 0 \right\rangle} \lambda_s^{\tilde{h}} \right] \ . \tag{A.10}$$

     Which yields

     $$P_h \sum_{s=0}^{n_h - 1} \delta_s^h d_{h-s}$$

Q.E.D.

LEMMA 13 *In the optimal solution to the dual problem A.7*

$$\mu^h = \max\left[0, v\left(p - p_0\right) + \left(1 - p\right)\mu^{\langle h,0\rangle} + p\mu^{\langle h,1\rangle} - \sum_{s=0}^{n_h - 1}\left(c_{h-s}\lambda_s^h - c_{h-s}\gamma_s^h + d_{h-s}\delta_s^h\right)\right]$$

PROOF: By construction, $\mu^h \geq 0$. The lemma states that if $\mu^h > 0$ then

$$\mu^h = v\left(p - p_0\right) + \left(1 - p\right)\mu^{\langle h,0\rangle} + p\mu^{\langle h,1\rangle} - \sum_{s=0}^{n_h - 1}\left(c_{h-s}\lambda_s^h - c_{h-s}\gamma_s^h + d_{h-s}\delta_s^h\right)$$

Suppose the statement is false.

1. If $\mu^\emptyset$ violates the condition, decrease $\mu^\emptyset$. This is feasible and decreases the objective. Therefore, $\mu^\emptyset$ was not optimal.
2. If any other history violates the condition, there must be an h that violates the condition and that in all histories that precede h the condition holds. Decrease μ^h by ε. As the constraint in the previous period binds, this allows decreasing the previous history's μ by either $\varepsilon \cdot p$ or $\varepsilon\left(1 - p\right)$. Continuing backwards, this will decrease $\mu^\emptyset$. This is a feasible decreases of the objective.

Q.E.D.

LEMMA *For every h let μ^h, λ^h be the solution to the dual problem A.7. Then $\mu^h = \mu\left(n_h, \gamma^h, \delta^h\right)$ and $\lambda^h \in \lambda^*\left(n_h, \gamma^h, \delta^h\right)$. In particular, $\mu\left(1, 0, 0\right) = \mu^\emptyset$. Where $\mu\left(n, \gamma, \delta\right)$ is defined recursively in 3.11.*

PROOF: For any history h, let σ^h be the solution to the problem of minimizing μ^h subject to the constraints in problem A.7 for all histories that follow h and given the δ^h and γ^h that correspond to the problem's solution starting at $\mu^\emptyset$. The solutionσ^h must be identical on all common histories to the solution for the original problem. In addition, the solution σ^h only specifies variables for histories that follow h.

Applying the Principle of Optimality, any problem starting at history h can be broken down to choosing the λ_s^h in history h and optimizing the continuation problem. It can be verified that the constraints for the recursive formulation are identical to the constraints in the linear formulation. Lemma 13 proves that the value of the solution in all continuations is the one defined in the objective. *Q.E.D.*

A.5. *Basic Properties of the Dual*

THEOREM *The following hold for the dual problem 4.10:*

- *$F\left(n, \Gamma, \Delta\right)$ is convex in $\left(\Gamma, \Delta\right)$. $\hat{F}\left(n, \Gamma, \Delta, \Lambda\right)$ is convex and continuous in Λ for every Γ, Δ.*

- *$F(n,\Gamma,\Delta)$ decreases in Γ_0 and Δ_0 for any $m>0$*
- *$F(n,\Gamma,\Delta)$ increases in Γ_m and Δ_m for any $m>0$*

Proof: The proof uses F_y to denote the derivative of F with respect to y, or if the derivative does not exit, the supergradient.

1. $F(n,\Gamma,\Delta)$ is convex in (Γ,Δ)

   (a) In any last period, $F(n,\Gamma,\Delta)$ is linear and thus continuous and convex.

   (b) Assume that $F(n+1,\Gamma,\Delta)$is continuous and convex. As the positive sum of three continuous and convex functions is continuous and convex, for every Λ, $\hat{F}(n,\Gamma,\Delta,\Lambda)$ is convex in (Γ,Δ). As the feasible set is convex and the objective is to minimize a convex function, $F(n,\Gamma,\Delta)$ is continuous and convex.

2. $\hat{F}(n,\Gamma,\Delta,\Lambda)$ is convex in Λ for every Γ,Δ.

   (a) For $\hat{F}(n,\Gamma,\Delta,\Lambda)$, the period return is linear in λ and so it is sufficient to show that the continuation is convex in Λ. I show this for the continuation after success – $F\left(n+1,\Gamma-\frac{p-p_0}{p}\Lambda,\Delta+\frac{p_0}{p}\Lambda\right)$. The same proof applies for the continuation after failure. As the sum of convex functions is convex, this completes the proof. Let Λ^1 and Λ^2 be feasible solutions. Then by convexity of $F(n+1,\Gamma,\Delta)$ for any $\alpha\in(0,1)$:

$$\alpha F\left(n+1,\Gamma-\tfrac{p-p_0}{p}\Lambda^1,\Delta+\tfrac{p_0}{p}\Lambda^1\right)+(1-\alpha)F\left(n+1,\Gamma-\tfrac{p-p_0}{p}\Lambda^2,\Delta+\tfrac{p_0}{p}\Lambda^2\right)\leq$$
$$F\left(n+1,\alpha\left(\Gamma-\tfrac{p-p_0}{p}\Lambda^1\right)+(1-\alpha)\left(\Gamma-\tfrac{p-p_0}{p}\Lambda^2\right),\alpha\left(\Delta+\tfrac{p_0}{p}\Lambda^1\right)+(1-\alpha)\left(\Delta+\tfrac{p_0}{p}\Lambda^2\right)\right)=$$
$$F\left(n+1,\Delta-\tfrac{p-p_0}{p}\left(\alpha\Lambda^1+(1-\alpha)\Lambda^2\right),\Delta+\tfrac{p_0}{p}\left(\alpha\Lambda^1+(1-\alpha)\Lambda^2\right)\right)$$

3. Given the previous result, the optimal Λ is either unique or an interval. Continuity is a standard result.

4. $F(n,\Gamma,\Delta)$ decreases in Δ_0 and increases in Δ_m for any $m>0$. By backward induction:

   - In any last period if $F(n,\Gamma,\Delta)>0$:
     - $F_{\Delta_0}=-d_n<0$ and
     - $F_{\Delta_m}=d_{n-m+1}-d_{n-m}>0$ (by c_n convex $\implies d_n$ increasing).
   - Suppose $F(n+1,\Gamma,\Delta)$ decreases in Δ_0 and increases in Δ_m for any $m>0$. Let Λ^* be optimal at the state (n,Γ,Δ).
   - Let e_0 be a vector of size $n-1$ with the first element some $\varepsilon>0$ and all other elements zero. As the constraint is not affected by Δ, Λ^* is feasible for any $(n,\Gamma,\Delta+e_0)$.

Therefore

$$\begin{aligned} F(n,\Gamma,\Delta+e_0) &\leq p\mu\left(n+1,\Gamma-\frac{p-p_0}{p}\Lambda^*,\Delta+e_0+\frac{p_0}{p}\Lambda^*\right)\\ &\quad+(1-p)\,\mu\left(n+1,\Gamma+\frac{p-p_0}{1-p}\Lambda^*,\Delta+e_0+\frac{1-p_0}{1-p}\Lambda^*\right)\\ &\quad+f(n,\Gamma,\Delta,\Lambda^*)-e_0\cdot d_n\\ &< F(n,\Gamma,\Delta)\end{aligned}$$

- Let e_m be a vector of size $n-1$ with element m some $\varepsilon>0$ and all other elements zero. As the constraint in not affected by Δ, Λ^* is feasible for any $(n,\Gamma,\Delta-e_m)$. Therefore

$$\begin{aligned} F(n,\Gamma,\Delta-e_m) &\leq p\mu\left(n+1,\Gamma-\frac{p-p_0}{p}\Lambda^*,\Delta-e_m+\frac{p_0}{p}\Lambda^*\right)\\ &\quad+(1-p)\,\mu\left(n+1,\Gamma+\frac{p-p_0}{1-p}\Lambda^*,\Delta-e_m+\frac{1-p_0}{1-p}\Lambda^*\right)\\ &\quad+f(n,\Gamma,\Delta,\Lambda^*)-e_m\cdot(d_{n-m+1}-d_{n-m})\\ &< F(n,\Gamma,\Delta)\end{aligned}$$

5. $F(n,\Gamma,\Delta)$ increases in Γ_0 and decreases in Γ_m for any $m>0$

- In any last period if $F(n,\Gamma,\Delta)>0$: $F_{\Gamma_m}=-d_{n-m+1}<0$. As Γ_m does not affect the constraint, the proof is the same as for Δ_0.
- In any last period, f decreases in Λ_0 and increases in all $\Lambda_{m>0}$, thus, the constraint must bind and

$$\Lambda_0=\frac{p}{p-p_0}(1+\Gamma_0)$$

- In any last period in which $F(n,\Gamma,\Delta)>0$

$$\frac{\partial F(n,\Gamma,\Delta)}{\Gamma_0}=-c_n\frac{p}{p-p_0}+c_n=-\frac{p_0}{p-p_0}c_n<0$$

- For any period (n,Γ,Δ), let Λ^* be an optimal solution. An increase ε in Γ makes it feasible to increase Λ_0 by $\varepsilon\frac{p}{p-p_0}$. Let e_0 be a vector with the first element $\varepsilon>0$ and the rest zeros. Increasing Λ_0 based on the increase in ε yields :

$$\begin{aligned}\Gamma^1 &= \Gamma+e_0-\frac{p-p_0}{p}\Lambda^*-e_0\frac{p}{p-p_0}\frac{p-p_0}{p}=\Gamma-\frac{p-p_0}{p}\Lambda^*\\ \Gamma^0 &= \Gamma+e_0+\frac{p-p_0}{1-p}\Lambda^*+e_0\frac{p}{p-p_0}\frac{p-p_0}{1-p}=\Gamma+\frac{p-p_0}{1-p}\Lambda^*+\frac{e_0}{1-p}\end{aligned}$$

In addition, the first elements of Δ^1 and Δ^0 increase and from the previous step this decreases the continuation values even more. Thus

$$\begin{aligned} F\left(n,\Gamma+e_0,\Delta\right) &\leq p\mu\left(n+1,\Gamma-\frac{p-p_0}{p}\Lambda^*,\Delta+\frac{p_0}{p}\Lambda^*\right) \\ &\quad +\left(1-p\right)\mu\left(n+1,\Gamma+\frac{p-p_0}{1-p}\Lambda^*,\Delta+\frac{1-p_0}{1-p}\Lambda^*\right) \\ &\quad +f\left(n,\Gamma,\Delta,\Lambda^*\right)-c_n\cdot\varepsilon\cdot\frac{p}{p-p_0} \\ &< F\left(n,\Gamma,\Delta\right) \end{aligned}$$

Q.E.D.

A.6. *Sufficiency of Local Deviations*

THEOREM *Any optimal contract subject to LDIC is an optimal contract.*

PROOF: By corollary 2 , it is sufficient to show that in the solution of the dual problem 4.10 starting from $F(1,0,0)$, $\Lambda_m=0$ for every $m>0$ at every state.

The objective in each state is to minimize $F(\cdot)$. Each Λ_m has three effects on $F(\cdot)$ – the period return $f(\cdot)$, the law of motion for Γ and the law of motion for Δ. The proof considers each of these separately and shows that for any $m>0$, $F(\cdot)$ increases through each of these effects and thus the optimal Λ_m is the lowest possible: $\Lambda_m=0$.

For this, consider a relaxed problem that allows, for every $m>0$ to choose separately $\Lambda_m^p,\Lambda_m^\Gamma,\Lambda_m^\Delta$ (all non-negative) such that Λ_m^p affects the period return term $f(\cdot)$, Λ_m^Γ affects the law of motion for Γ , and Λ_m^Δ affects the law of motion for Δ. Moreover, the constraints $\Lambda_m\geq\Lambda_{m+1}$ are ignored. The proof will show that the optimal solution sets $\Lambda_m^p=\Lambda_m^\Gamma=\Lambda_m^\Delta=0$. Thus, in the original problem $F(\cdot)$, it must be that the optimal solution is $\Lambda_m=0$.

We first establish that for any $m>0$, it is optimal to set $\Lambda_m^p=\Lambda_m^\Delta=0$. Recall that all the Λ variables are chosen to minimize $F(\cdot)$. Then:

- Λ_m^p only appears in $f(\cdot)$ with a positive coefficient as for all n, $d_n\geq0$. Therefore, the only effect of a reduction in Λ_m^p for $m>0$ is a decrease in $f(\cdot)$. As setting $\Lambda_m^p=0$ is feasible it must be optimal.

- Λ_m^Δ only appears in the law of motion for Δ_m. Moreover, both Δ_m^s and Δ_m^f increase with Λ_m^Δ. By theorem 2 above, $F(n+1,\cdot)$ is increasing in Δ_m for $m>0$. Thus, $\Lambda_m^\Delta=0$ is feasible and optimal.

It remains to consider Λ_m^Γ. As the purpose of the analysis is to show that $\Lambda_m^\Gamma=0$ for all $m>0$, let $\overline{\Lambda}$ be the vector $\left(\Lambda_1^\Gamma,\Lambda_2^\Gamma,...,\Lambda_{n+1}^\Gamma\right)$ and $\overline{\Gamma}$ be the vector $(\Gamma_1,...,\Gamma_n)$ so that $\overline{\Lambda}_m=\Lambda_m^\Gamma$ and similarly $\overline{\Gamma}_m=\Gamma_m$.

By the law of motion for Δ, if $\Lambda_m^\Delta=0$ at all states then in all states along the optimal solution $\Delta_m=0$. Define the implied problem $G\left(n,\Gamma_0,\Delta_0,\overline{\Gamma}\right)$ derived by removing from $F(\cdot)$ all elements that are known to be zero (Δ_m,Λ_m^p and Λ_m^Δ for $m>0$) and using the definition of $\overline{\Gamma}$ and $\overline{\Lambda}$ as above. As we will use monotone comparative static results, the existing results will be more familiar for maximizing $-F(\cdot)$ instead of minimizing $F(\cdot)$. Thus, the problem is:

$$G\left(n,\Gamma_0,\Delta_0,\overline{\Gamma}\right) \;=\; \min\left[0,\max_{\Lambda_0\geq 0,\overline{\Lambda}\geq 0}\hat{G}\left(n,\Gamma_0,\Delta_0,\overline{\Gamma},\Lambda_0,\overline{\Lambda}\right)\right]$$

With

$$\begin{aligned}\hat{G}\left(n,\Gamma_0,\Delta_0,\overline{\Gamma},\Lambda_0,\overline{\Lambda}\right) \;=\;& pG\left(n+1,\Gamma_0^s,\Delta_0^s,\overline{\Gamma}^s\right)\\ &+(1-p)\,G\left(n+1,\Gamma_0^f,\Delta_0^f,\overline{\Gamma}^f\right)\\ &-v\,(p-p_0)+c_n\,(\Lambda_0-\Gamma_0)+d_n\Delta_0+\sum_{m=1}^{n_h}d_{n-m}\overline{\Gamma}_m\end{aligned}$$

$s.t.$

$$\begin{aligned}&\Lambda_0\leq\frac{p}{p-p_0}\,(1+\Gamma_0)\\ &\Gamma_0^s=\Gamma_0-\frac{p-p_0}{p}\Lambda_0\quad;\quad \Gamma_0^f=\Gamma_0+\frac{p-p_0}{1-p}\Lambda_0\\ &\overline{\Gamma}^s=\Gamma-\frac{p-p_0}{p}\overline{\Lambda}\quad;\quad \overline{\Gamma}^f=\overline{\Gamma}+\frac{p-p_0}{1-p}\overline{\Lambda}\\ &\Delta_0^s=\Delta_0+\frac{p_0}{p}\Lambda_0\quad;\quad \Delta_0^f=\Delta_0+\frac{1-p_0}{1-p}\Lambda_0\end{aligned}$$

It remains to show that in the solution to $G(1,0,0,0)$, $\overline{\Lambda}=0$ is optimal for all feasible states. Observe that in any last period N^{FB}, the optimal solution sets $\Lambda_0=\frac{p}{p-p_0}\,(1+\Gamma_0)$. Therefore

$$G\left(N^{FB},L_0,\Delta_0,\overline{L}\right)= \min\left[0,-v\,(p-p_0)+\tfrac{p_0}{p-p_0}c_{N^{FB}}+\tfrac{p_0}{p-p_0}c_{N^{FB}}L_0+\textstyle\sum_{m=1}^{N^{FB}-1}d_{N^{FB}-m+1}\overline{L}_m+d_n\Delta_0\right] \tag{A.11}$$

The following preliminary results will be used in the proof.

Lemma 14 *$G\left(n,\Gamma_0,\Delta_0,\overline{\Gamma}\right)$ and $\hat{G}\left(n,\Gamma_0,\Delta_0,\overline{\Gamma},\Lambda_0,\overline{\Lambda}\right)$ are increasing continuous and concave in Γ_0,Δ_0 and $\overline{\Gamma}_m$ for any m. $\hat{G}\left(n,\Gamma_0,\Delta_0,\overline{\Gamma},\Lambda_0,\overline{\Lambda}\right)$ is concave in $\Lambda_0,\overline{\Lambda}$.*

Proof: As $F(n,\cdot)$ and $\hat{F}(n,\cdot)$ are convex and continuous in all their arguments for every n, $G(n,\cdot)$ and $\hat{G}(n,\cdot)$ are concave and continuous in all their arguments.

To prove the remaining claims it is sufficient to prove that in any last period $G\left(N^{FB},\Gamma_0,\Delta_0,\overline{\Gamma}\right)$ is increasing inΓ_0,Δ_0,$\overline{\Gamma}$ and that backward induction implies $\hat{G}\left(n,\Gamma_0,\Delta_0,\overline{\Gamma},\Lambda_0,\overline{\Lambda}\right)$ is increasing in $\Gamma_0,\Delta_0,\overline{\Gamma}$. The last period result is observable in equation A.11. The backward induction step:

1. For Δ_0 (and $\overline{\Gamma}_m$): For any n, suppose that $G(n+1,\cdot)$ increases in Δ_0 (and $\overline{\Gamma}_m$). Then starting with a higher Δ_0 (and $\overline{\Gamma}_m$) does not affect any constraint but for any $\left(\Lambda_0,\overline{\Lambda}\right)$ this increases both the period return and the continuation values (Δ^s,Δ^f,$\overline{\Gamma}^s$,$\overline{\Gamma}^f$). Thus, the proof is complete for Δ_0 and $\overline{\Gamma}$.

2. For Γ_0 : For any n, suppose that $G(n+1,\cdot)$ increases in Γ_0 and Δ_0. Let $\Lambda_0^*,\overline{\Lambda}^*$ be optimal for Γ_0. For any $\varepsilon>0$ increase in Γ_0 , it is feasible to increase Λ_0^* by $\varepsilon\frac{p}{p-p_0}>\varepsilon$. The period return increases by $\varepsilon\frac{p_0}{p-p_0}c_n$. Thus, it remains to show that all the state variables weakly increase. The continuation states $\overline{\Gamma}^s$,$\overline{\Gamma}^f$ are unaffected as $\overline{\Lambda}^*$ did not change. The continuation Δ^s,Δ^f

increase as they both increase with Λ_0. Γ_0^f increases with Γ_0 and with Λ_0. Finally, the effect on Γ_0^s is exactly zero for any ε:

$$\begin{aligned}\Gamma_0^s &= \Gamma_0 + \varepsilon - \Lambda_0^* \frac{p-p_0}{p} - \varepsilon \frac{p}{p-p_0} \frac{p-p_0}{p} \\ &= \Gamma_0 - \Lambda_0^* \frac{p-p_0}{p}\end{aligned}$$

Q.E.D.

Given the concavity result in the previous lemma, it is sufficient to evaluate the first order effect on $\hat{G}(n,\cdot)$ of any possible marginal increase in $\overline{\Lambda}$ starting from $\overline{\Lambda}=0$ to determine whether $\overline{\Lambda}=0$ is the optimal solution for $\hat{G}(n,\cdot)$.

For this, define $G_{\overline{\Lambda}^+}(n,\cdot)$ as the positive gradient of $G\left(n,\Gamma_0,\Delta_0,\overline{\Gamma}\right)$ along any direction $\overline{\Lambda}$. That is

$$G_{\overline{\Lambda}^+}\left(n,\Gamma_0,\Delta_0,\overline{\Gamma}\right) = \lim_{\alpha\to 0,\alpha>0} \frac{G\left(n+1,\Gamma_0,\Delta_0,\overline{\Gamma}+\alpha\overline{\Lambda}\right) - G\left(n+1,\Gamma_0,\Delta_0,\overline{\Gamma}\right)}{\alpha}\ .$$

Define $G_{\overline{\Lambda}^-}(n,\cdot)$ as the negative gradient along direction $\overline{\Lambda}$ (the standard definition given $G_{\overline{\Lambda}^+}(n,\cdot)$).

For any Λ_0, let

$$\hat{G}_{\overline{\Lambda}^+}\left(n,\Gamma_0,\Delta_0,\overline{\Gamma},\Lambda_0,0\right) = \lim_{\alpha\to 0,\alpha>0} \frac{\hat{G}\left(n+1,\Gamma_0,\Delta_0,\overline{\Gamma},\Lambda_0,\alpha\overline{\Lambda}\right) - \hat{G}\left(n+1,\Gamma_0,\Delta_0,\overline{\Gamma},\Lambda_0,0\right)}{\alpha}\ .$$

A sufficient condition for $\overline{\Lambda}=0$ to be optimal is that, for any $\Lambda_0 \geq 0$

$$\hat{G}_{\overline{\Lambda}^+}\left(n,\Gamma_0,\Delta_0,\overline{\Gamma},\Lambda_0,0\right) \leq 0\ .$$

Given the definitions of $G_{\overline{\Lambda}^+}(n,\cdot)$ and $G_{\overline{\Lambda}^-}(n,\cdot)$:

$$\hat{G}_{\overline{\Lambda}^+}\left(n,\Gamma_0,\Delta_0,\overline{\Gamma},\Lambda_0,0\right) = (p-p_0)\left(G_{\overline{\Lambda}^+}\left(n+1,\Gamma_0^f,\Delta_0^f,\overline{\Gamma}\right) - G_{\overline{\Lambda}^-}\left(n+1,\Gamma_0^s,\Delta_0^s,\overline{\Gamma}\right)\right)\ .$$

As $p > p_0$, it remains to show that for any n and $\left(\Lambda_0,\overline{\Lambda}\right)$ non-negative:

$$G_{\overline{\Lambda}^+}\left(n,\Gamma_0^f,\Delta_0^f,\overline{\Gamma}\right) \leq G_{\overline{\Lambda}^-}\left(n,\Gamma_0^s,\Delta_0^s,\overline{\Gamma}\right) \tag{A.12}$$

As $G(n,\cdot)$ is concave in $\left(\Gamma_0,\Delta_0,\overline{\Gamma}\right)$, $G_{\overline{\Lambda}^+}(n,\cdot)$ and $G_{\overline{\Lambda}^-}(n,\cdot)$ exist everywhere and for any Γ_0,Δ_0:

$$G_{\overline{\Lambda}^+}\left(n,\Gamma_0,\Delta_0,\overline{\Gamma}\right) \leq G_{\overline{\Lambda}^-}\left(n,\Gamma_0,\Delta_0,\overline{\Gamma}\right). \tag{A.13}$$

Thus, a sufficient condition for A.12 is

$$G_{\overline{\Lambda}^+}\left(n,\Gamma_0^f,\Delta_0^f,\overline{\Gamma}\right) \leq G_{\overline{\Lambda}^+}\left(n,\Gamma_0^s,\Delta_0^s,\overline{\Gamma}\right) \tag{A.14}$$

at $\Lambda_0 = 0$ $\gamma^f = \gamma^s$ and $\Delta^s = \Delta^f$ and thus condition A.14 trivially holds.

For any $\Lambda_0 > 0$, recall that $\Gamma_0^f > \Gamma_0^s$ and $\Delta_0^f > \Delta_0^s$. Thus it is sufficient to show that as Γ_0 and Δ_0 increase $G_{\overline{\Lambda}^+}\left(n, \Gamma_0, \Delta_0, \overline{\Gamma}\right)$ decrease. This is equivalent to the requirement that $G\left(n, \Gamma_0, \Delta_0, \overline{\Gamma}\right)$ is supermodular in $\left(-\Gamma_0, \overline{\Gamma}_m\right)$ and in $\left(-\Delta_0, \overline{\Gamma}_m\right)$ for every m. The next result therefore concludes the proof.

Lemma 15 *$G\left(n, \Gamma_0, \Delta_0, \overline{\Gamma}\right)$ is supermodular in $\left(-\Gamma_0, \overline{\Gamma}_m\right)$ and in $\left(-\Delta_0, \overline{\Gamma}_m\right)$ for every m.*

Proof: By backward induction.

- In any last period N^{FB}, the optimal solution sets $\Lambda_0 = \frac{p}{p-p_0}\left(1+\Gamma_0\right)$. Therefore

$$G\left(N^{FB}, L_0, \Delta_0, \overline{L}\right) = \min\left[0, -v\left(p-p_0\right) + \frac{p_0}{p-p_0} c_{N^{FB}} + \frac{p_0}{p-p_0} c_{N^{FB}} L_0 + \sum_{m=1}^{N^{FB}-1} d_{N^{FB}-m+1} \overline{L}_m + d_n \Delta_0\right] \tag{A.15}$$

1. If $G\left(N^{FB}, \cdot\right) < 0$ then $\frac{\partial G\left(N^{FB}, \cdot\right)}{\partial \overline{\Gamma}_m} = d_{N^{FB}-m+1} > 0$.
2. As Γ_0 or Δ_0 increases, eventually $G\left(N^{FB}, \cdot\right) = 0$ and $\frac{\partial G\left(N^{FB}, \cdot\right)}{\partial \Gamma_m}$ decreases to $\frac{\partial G\left(N^{FB}, \cdot\right)}{\partial \Gamma_m} = 0$.
3. Therefore, for every Γ_m, $G\left(N^{FB}, \Gamma_0, \Delta_0, \overline{\Gamma}\right)$ is supermodular in $\left(-\Gamma_0, \overline{\Gamma}_m\right)$ and in $\left(-\Delta_0, \overline{\Gamma}_m\right)$ for every m.

Suppose that $G\left(n+1, \Gamma_0, \Delta_0, \overline{\Gamma}\right)$ is supermodular in $\left(-\Gamma_0, \overline{\Gamma}_m\right)$ and in $\left(-\Delta_0, \overline{\Gamma}_m\right)$ for every m.

For any $\Lambda_0, \overline{\Lambda}$, $G\left(n+1, \Gamma_0^s, \Delta_0^s, \overline{\Gamma}^s\right)$ and $G\left(n+1, \Gamma_0^f, \Delta_0^f, \overline{\Gamma}^f\right)$ are also supermodular as required.

All remaining parts of the objective of $\hat{G}\left(n, \Gamma_0, \Delta_0, \overline{\Gamma}, \Lambda_0, \overline{\Lambda}\right)$ are linear in all variables.

Therefore, for any $\left(\Lambda_0, \overline{\Lambda}\right)$, the objective of $\hat{G}\left(n, \Gamma_0, \Delta_0, \overline{\Gamma}, \Lambda_0, \overline{\Lambda}\right)$ is supermodular as required.

As the feasible set is a sub-lattice, supermodularity is preserved under maximization (see section A.7 below for details). Therefore, $\max_{\left(\Lambda_0, \overline{\Lambda}\right)} \hat{G}\left(n, \Gamma_0, \Delta_0, \overline{\Gamma}, \Lambda_0, \overline{\Lambda}\right)$ is supermodular in $\left(-\Gamma_0, \overline{\Gamma}_m\right)$ and in $\left(-\Delta_0, \overline{\Gamma}_m\right)$ for every m.

It remains to show that $G\left(n, \Gamma_0, \Delta_0, \overline{\Gamma}\right) = \min\left[0, \max_{\Lambda_0, \overline{\Lambda}} \hat{G}\left(n, \Gamma_0, \Delta_0, \overline{\Gamma}, \Lambda_0, \overline{\Lambda}\right)\right]$ is supermodular in $\left(-\Gamma_0, \overline{\Gamma}_m\right)$ and in $\left(-\Delta_0, \overline{\Gamma}_m\right)$ for every m. Lemma 17 shows that, given the supermodularity result for $\hat{G}(\cdot)$, it is sufficient to show that $\max_{\Lambda_0, \overline{\Lambda}} \hat{G}\left(n, \Gamma_0, \Delta_0, \overline{\Gamma}, \Lambda_0, \overline{\Lambda}\right)$ is increasing in Γ_0, Δ_0 and $\overline{\Gamma}_m$ for any m. This was proved in lemma 14. *Q.E.D.*

Q.E.D.

A.7. *Monotone Comparative Static Results*

The result used in the proof is an application of Theorem 2.7.6 in Topkis (1998). I repeat the relevant construction here.

- For any $y \in Y \subset R^n$ and $X \subset R^m$, let $X(y) \subset X$ denote a subset of X for each y
- For any $x, x' \in R^n$, let $x \wedge x'$ denote the meet (pairwise minimum) of x and x' and $\vee$ denote the join (pairwise maximum).

- The space $X(y)\times Y$ is a sub-lattice iff for any $x \in X(y)$ and $x^{'} \in X\left(y^{'}\right)$, $x \wedge x^{'} \in X\left(y \wedge y^{'}\right)$ and $x \vee x^{'} \in X\left(y \vee y^{'}\right)$
- For $f(x,y) : X\times Y \to R$ suppose that $X(y)\times Y$ is a sub-lattice and let $h(y) = \max_{x\in X(y)} f(x,y)$. Then if $f(x,y)$ is supermodular in y_i, y_j , so is $h(y)$.

Note that the commonly used result is simpler as it assumes the feasible set does not depend on y.

LEMMA 16 *The set of Λ and Γ such that Λ is feasible for Γ in $\hat{F}(\cdot)$ or $\hat{G}(\cdot)$ defined above is a sub-lattice.*

PROOF: The statement is equivalent to the following: For every two pairs Γ, Λ and $\Gamma^{'}, \Lambda^{'}$ such that in problem $\hat{F}(n,\cdot)$, Λ is feasible for Γ and $\Lambda^{'}$ is feasible for $\Gamma^{'}$, it must be that $\Lambda \wedge \Lambda^{'}$ is feasible for $\Gamma \wedge \Gamma^{'}$ and $\Lambda \vee \Lambda^{'}$ is feasible for $\Gamma \vee \Gamma^{'}$.

This may be verified directly, and is also worked out in part (d) of example 2.6.2 in Topkis (1998) as all constraints are of the form $x_i - x_j \leq b$. *Q.E.D.*

LEMMA 17 *Suppose $z(x) : R^n \to R$ is increasing and continuous in x_i, x_j and supermodular in $(-x_i, x_j)$. Then $\min[0, z(x)]$ is supermodular in $(-x_i, x_j)$*

PROOF: By standard monotone comparative static results, it is sufficient to consider the two variable function $z(x_i, x_j)$. For every $x_i^{'} \geq x_i$ and $x_j^{'} \geq x_j$, the lemma's assumption is that

$$z\left(x_i, x_j^{'}\right) - z(x_i, x_j) \geq z\left(x_i^{'}, x_j^{'}\right) - z\left(x_i^{'}, x_j\right)$$

We need to show that the inequality is preserved under application of the min operator for each element:

$$\min\left[0, z\left(x_i, x_j^{'}\right)\right] - \min\left[0, z(x_i, x_j)\right] \geq \min\left[0, z\left(x_i^{'}, x_j^{'}\right)\right] - \min\left[0, z\left(x_i^{'}, x_j\right)\right] \tag{A.16}$$

Consider each possible case separately:

1. If $z(x_i, x_j) \geq 0$ then by $z()$ increasing, all min operators bind and both sides of A.16 are zero.
2. If $z\left(x_i^{'}, x_j^{'}\right) \leq 0$ then by $z()$ increasing all min operators are redundant and the inequality holds by assumption.
3. If $z\left(x_i^{'}, x_j^{'}\right)$ is the only non-negative number, by $z()$ increasing there is some $\tilde{x}_j \in \left(x_j, x_j^{'}\right)$ such that $z\left(x_i^{'}, \tilde{x}_j\right) = 0 = \min\left[0, z\left(x_i^{'}, x_j^{'}\right)\right]$ and so

$$z\left(x_i^{'}, \tilde{x}_j\right) - z\left(x_i^{'}, x_j\right) = \min\left[0, z\left(x_i^{'}, x_j^{'}\right)\right] - z\left(x_i^{'}, x_j\right) \ .$$

By $x_j^{'} \geq \tilde{x}_j$,

$$z\left(x_i^{'}, x_j^{'}\right) - z\left(x_i^{'}, x_j\right) \geq z\left(x_i^{'}, \tilde{x}_j\right) - z\left(x_i^{'}, x_j\right)$$

Therefore, by assumption

$$z\left(x_i,x_j^{'}\right)-z\left(x_i,x_j\right)\geq z\left(\tilde{x}_i,x_j^{'}\right)-z\left(x_i^{'},x_j\right)=\min\left[0,z\left(x_i^{'},x_j^{'}\right)\right]-z\left(x_i^{'},x_j\right) .$$

4. If only $z\left(x_i^{'},x_j^{'}\right)$ and $z\left(x_i,x_j^{'}\right)$ are positive then A.16 is the same as

$$z\left(x_i^{'},x_j\right)\geq z\left(x_i,x_j\right)$$

which holds as $z\left(x_i,x_j\right)$ is increasing in x_i.

5. If only $z\left(x_i^{'},x_j^{'}\right)$ and $z\left(x_i^{'},x_j\right)$ are positive the A.16 is the same as

$$z\left(x_i,x_j^{'}\right)\geq z\left(x_i,x_j\right)$$

which holds as $z\left(x_i,x_j\right)$ is increasing in x_j.

6. If $z\left(x_i^{'},x_j^{'}\right)$, $z\left(x_i^{'},x_j\right)$ and $z\left(x_i,x_j^{'}\right)$ are positive than A.16 is the same as $z\left(x_i,x_j\right)\leq 0$ which holds by assumption.

Q.E.D.

A.8. *Lemma 5*

LEMMA *Problem 5.2 is convex in (γ,δ). The optimal solution decreases in the costs (γ and δ), γ,δ are substitutes. If $\mu\left(n,\gamma,\delta,\lambda\right)<0$, the contract is terminated.*

PROOF: The proof for the first two statements (convex and decreasing) is identical to the proof of proposition 2. That γ,δ are substitutes is proved as part of the proof of theorem 1. If $\mu\left(n,\gamma,\delta,\lambda\right)<0$ then in the optimal contract problem the constraint $q\geq 0$ for the history binds. Thus, $q=0$ and the contract is terminated. *Q.E.D.*

A.9. *Two Period Solution*

First prove the lemma:

LEMMA *If the contract must terminate within at most two periods, the wage constraint binds in both remaining periods. In a two period problem, the dual values $\mu^{\emptyset},\mu^0$ and μ^1 are given by:*

$$\begin{aligned}
\mu^1 &= \max\left[0,v\left(p-p_0\right)-c_2-\frac{p_0}{p-p_0}\left(c_2-c_1\right)\right]\\
\mu^0 &= \max\left[0,v\left(p-p_0\right)+c_1\frac{1-p_0}{p-p_0}\frac{p}{1-p}-c_2\frac{p}{1-p}\frac{2-p}{p-p_0}\right]\\
\mu &= v\left(p-p_0\right)-\frac{p}{p-p_0}c_1+p\mu^1+\left(1-p\right)\mu^0
\end{aligned}$$

PROOF: In any last period, the dual problem is:

$$\begin{aligned}\mu(n,\gamma,\delta,\lambda) &= \min_{\lambda\geq 0} v(p-p_0) - c_n\lambda - \delta d_n + \gamma\cdot c_n \\ s.t. &\quad (p-p_0)\lambda \leq p(1+\gamma)\end{aligned}$$

The objective decreases with λ (by $-c_n$) and therefore must bind in the optimal solution. Placing λ in the objective, simplifying and binding by zero

$$\mu(n,\gamma,\delta) = \max\left[0, v(p-p_0) - c_n\frac{p}{p-p_0} - \delta d_n - \gamma\cdot c_n\frac{p_0}{p-p_0}\right]$$

The first two terms are the 'static return' the next term is the private information cost, and the last term is the utility cost ($c_n\frac{p_0}{p-p_0}$ is the agent's expected utility).

Directly using the continuation values for the state variables, the problem in the previous period is:

$$\begin{aligned}\mu(n-1,\gamma,\delta,\lambda) &= \min_{\lambda\geq 0} v(p-p_0) - c_{n-1}\lambda - \delta d_{n-1} + \gamma\cdot c_{n-1} \\ &\quad + p\max\left[0, v(p-p_0) - c_n\frac{p}{p-p_0} - \delta d_n - \lambda\frac{p_0}{p}d_n - \gamma\cdot c_n\frac{p_0}{p-p_0} + \lambda\frac{p-p_0}{p}\frac{p_0}{p-p_0}c_n\right] \\ &\quad + (1-p)\max\left[0, v(p-p_0) - c_n\frac{p}{p-p_0} - \delta d_n - \lambda\frac{1-p_0}{1-p}d_n - \gamma\cdot c_n\frac{p_0}{p-p_0} - \lambda\frac{p-p_0}{1-p}\cdot\frac{p_0}{p-p_0}c_n\right] \\ s.t. &\quad (p-p_0)\lambda \leq p(1+\gamma)\end{aligned}$$

We first show that the objective always decreases with λ and thus the wage constraint must bind:

- If both continuations in period n are negative, the effect of λ on the objective is $-c_{n-1} < 0$
- If only the continuation in period n after success is positive, the effect of λ on the objective is

$$-c_{n-1} - p\cdot\frac{p_0}{p}d_n + \frac{p_0}{p}c_n = -(1-p_0)c_{n-1} < 0$$

- As the value after success is higher than the value after failure (lemma 5), if the value after failure is positive the value after success is also positive. Thus, the last case to consider is that both continuations are positive. The value of the continuation after failure is decreasing in λ by $\frac{1-p_0}{1-p}d_n + \frac{p_0}{1-p}c_n$ whenever it is non-zero. Thus, the overall effect of λ on the objective must be negative.

Plugging in $\lambda = \frac{p}{p-p_0}(1+\gamma)$ in the problem for period $n-1$ obtains the result in the lemma. *Q.E.D.*

PROPOSITION 6*In a two period problem, for any $\varepsilon > 0$ consider increasing c_1 by ε and decreasing c_2 by ε. The change strictly increases expected profits if and only if:*

- *The optimal contract asks the agent to work after failing in the first period; or*
- *The optimal contract asks the agent to work only after success in the first period and $p+p_0 > 1$*

PROOF: The first condition is equivalent to

$$v(p-p_0) \geq \overline{v}$$

The second condition is equivalent to

$$v\left(p-p_0\right)\in\left[c_2+\frac{p_0}{p-p_0}\left(c_2-c_1\right),\overline{v}\right]$$

With

$$\overline{v}\equiv c_2\frac{p}{1-p}\frac{2-p}{p-p_0}-c_1\frac{1-p_0}{p-p_0}\frac{p}{1-p}$$

By the lemma, if $v\left(p-p_0\right)\geq\overline{v}$, the dual value after failure is positive and the contract asks the agent to work in the second period for any continuation. The proof for proposition 4 applies and the payment in both periods is $\frac{c_2}{p-p_0}$. Thus, the payoff to the principal decreases in c_2 and is not affected by c_1. As $\overline{v}$ increases in c_2 and decreases in c_1, applying the change proposed in the proposition will not violate the condition.

If

$$v\left(p-p_0\right)\in\left[c_2+\frac{p_0}{p-p_0}\left(c_2-c_1\right),\overline{v}\right]$$

The dual value is positive after success and negative after failure. In this case the expected profit is

$$\mu=v\left(p-p_0\right)-\frac{p}{p-p_0}c_1+p\cdot\left(v\left(p-p_0\right)-c_2-\frac{p_0}{p-p_0}\left(c_2-c_1\right)\right)$$

The effect of cost changes on profits is:

$$\frac{\partial\mu}{\partial c_1}=-\frac{p}{p-p_0}+p\frac{p_0}{p-p_0}=-\frac{p}{p-p_0}\left(1-p_0\right)$$

$$\frac{\partial\mu}{\partial c_2}=-p\left(1+\frac{p_0}{p-p_0}\right)=-p\frac{p}{p-p_0}$$

Thus, the change proposed in the proposition increases expected profits iff

$$\frac{p}{p-p_0}\left(1-p_0\right)<\frac{p}{p-p_0}p$$

Or simply

$$p+p_0>1$$

Q.E.D.