

Order Number 9522927

Achieving objectives in a stochastic environment in the presence
of moral hazard

Spero, Andrew Edward, Jr., Ph.D.

Carnegie Mellon University, 1994

Copyright ©1995 by Spero, Andrew Edward, Jr. All rights reserved.

U·M·I

300 N. Zeeb Rd.
Ann Arbor, MI 48106

**Achieving Objectives in a Stochastic Environment
in the Presence of Moral Hazard**

Andrew Edward Spero Jr.

May 31, 1994

I would like to thank my dissertation committee—John O'Brien, Shiva Sivaramkrishnan, and especially the Chairman, Srikant Datar—for their help along with Stan Baiman, Jon Glover, and Steve Sung for their comments previous versions of this dissertation. In addition, I would like to thank the seminar participants at GSIA, Harvard, Minnesota, Northwestern, Stanford, UBC, and Washington University for their comments.

CARNEGIE MELLON UNIVERSITY

GRADUATE SCHOOL OF INDUSTRIAL
ADMINISTRATION

DISSERTATION

Submitted in partial fulfillment of the requirements

for the degree of DOCTOR OF PHILOSOPHY in

Industrial Administration (Accounting)

Title ACHIEVING OBJECTIVES IN A STOCHASTIC ENVIRONMENT
OR

If at first you don't succeed, try, try again

Presented by ANDY SPERO

Accepted by Srikant Datar
Professor Srikant Datar, Chairman

May 13, 1994
Date

Approved by the Dean

Robert Sullivan
Robert Sullivan Dean

May 13, 1994
Date

Achieving Objectives in a Stochastic Environment in the Presence of Moral Hazard

by

Andrew Edward Spero Jr.

Abstract: Responsibility accounting is a procedure for implementing a management-by-objectives system in which the subordinate's performance is evaluated in comparison to agreed-upon goals. In many institutional settings, agents are indeed hired to accomplish goals. In a stochastic environment, where failure to achieve the goal is costly and success may be random, the agent may be required to act repeatedly until the task has been successfully completed. In this thesis, I study an agency model in which the agent must accomplish a task in a stochastic environment. I assume that when the task is complete, the contract ends. Under suitable conditions, I find the main result of this analysis: when the agent precommits to completing the task, the first-best contract may have an infinite horizon, whereas the second-best contract will have a finite, maximum length (which depends upon the level of exogenous input costs). I derive some results with respect to the sequences of wages and actions in these contracts and, under specific assumptions, show that wages are decreasing and actions are increasing over time.

I also consider this problem when the agent learns over time. Learning by the agent increases his efficiency (or likelihood to accomplish the task in any period). Learning can occur costlessly—through experience—or may require additional, specific, directed effort. With precommitment by the agent, I find that my main result, as described above still holds, and without precommitment by the agent, I find that when learning occurs costlessly, or through experience, two results are obtained: (1) holding future parameters fixed, more efficient or better trained agents work harder than inefficient agents, and (2) as the agent can learn more within an interval, actions at the beginning of the interval decrease. However, in general, actions are not increasing over time.

Chapter I - Introduction:

In this dissertation, I analyze an agency model of responsibility accounting (and responsibility centers) in which the superior holds the subordinate accountable for meeting a predetermined objective. Responsibility accounting is a procedure for implementing a management-by-objectives system in which the subordinate's performance is evaluated in comparison to agreed-upon goals. The analysis presented here provides a characterization of the optimal contractual arrangement between the superior and subordinate in such a system and thus leads to insights into the planning and control process that management accounting systems are designed to support.


Anthony, Dearden, and Govindarajan (1992) list the design and support of management control systems as one of the three functions of management accounting (along with differential accounting for decisions and cost accounting for valuation purposes), and Horngren, Foster, and Datar (1994) divide the management control process into two components: planning and control. They define planning as "choosing goals ... and then deciding how to attain the desired goals," and control as "both the action that implements the planning decision and performance evaluation of its personnel and its operations."

In the planning phase, the superior—or the superior and the subordinate in a participative budgeting process—determine the objective to be achieved and the strategy or plan for achieving that objective. In this thesis, I abstract away from the planning problem by assuming that the objective—profit maximization—and the production function—the probabilistic relationship between effort and output—are exogenously specified.

Achieving Objectives

In the control phase, the subordinate acts, and the superior evaluates the subordinate and provides feedback. This feedback communicates either a change in prospective goals or prospective actions to the subordinate.¹ In this analysis, where I assume that the goal or objective is specified exogenously, the feedback consists solely of instructions about future activities (or effort levels); however, since both the superior and subordinate are assumed to be far-sighted and rational, both can determine and analyze the extensive form of their common game tree and thus develop their optimal policies ex ante. Thus, in the general case (as many authors, including Anthony, Dearden, and Govindarajan, have noted) one cannot separate the planning process from the control process since the optimal plans and controls are determined simultaneously. The following diagram provides an illustration of the planning and control process:

Exhibit 1: A Simple Illustration of the Planning & Control Process



¹As is diagrammed in Exhibit 1.

Achieving Objectives

In this thesis, I analyze the following cases: (i) when the superior can measure the subordinate's activities and results—the first-best case—and (ii) when he cannot observe the subordinate's actions—the second-best case. In the second-best case, all inferences (and evaluations) of effort must be drawn from the measurements of output.

In a generic responsibility center the subordinate generates output through the combination of his efforts and other inputs. The evaluation of the subordinate's performance and the nature of the feedback communicated can occur on the basis of one or two dimensions—effectiveness and/or efficiency. Effectiveness is the relationship between the center's actual output and its objective (or planned output). Efficiency is the relationship between (the ratio of) the center's actual output and (to) actual input(s).² The choice of these performance measures depends upon the nature of the responsibility center, i.e., the factors for which the subordinate is held accountable.³ In this thesis, I assume that the subordinate must produce one unit of nondivisible output to be effective (and thus is in control of a cost center). Furthermore, I assume that if the subordinate puts forth enough effort, he can be effective in any period; that is, he can accomplish the objective (by producing a unit of output). As shown in Chapter III, the principal uses this fact to design the optimal contract.

The subordinate's efficiency decreases as the time (and materials) used to complete the task increases. In this model, where failure requires the subordinate to begin anew and he does not learn from past failures, the actual efficiency (or actual partial productivity) can be measured by the number of periods elapsed until the subordinate is successful or the actual usage of (number of units of) raw materials. As one would expect, information

²Here, efficiency is equivalent to productivity. One can determine partial productivity factors, or output compared to each individual input, or total factor productivity where output is divided by the aggregate cost of all inputs (including the agent's labor).

³Costs, revenues, or both.

Achieving Objectives

about the subordinate's efficiency (at accomplishing the task) is incorporated into his periodic performance evaluation and thus into his periodic compensation function.

There has been little theoretical work in the accounting field relating to the control process and management-by-objective systems;⁴ however, in many economic relationships, individuals are hired to accomplish goals: be it to properly build, write, sell, buy, or count something according to well-defined specifications. Such goal-based relationships include job-shop operations and government, defense, and residential contracts where the contractor must complete a construction or remodeling task. In these settings, the contractor must work until the goal is accomplished; so, the length of the relationship is random.

I frame my analysis as a principal-agent model in which the agent is hired to achieve an objective and, depending upon the derived contract, there is a possibility that the agent will not immediately accomplish the task. When the agent's failure to achieve the objective is costly to the principal, the agent may be required to act until the task is successfully completed; thus, in general, the optimal contract will have a random length—as many of the relationships mentioned above do.⁵ When the agent's effort is unobservable, an incentive problem exists because the principal desires that the agent work hard (to reduce the probability of paying failure costs), whereas the risk- and work-averse agent prefers to work at a lower level of effort and risk the possibility of taking additional effort at that lower level.

⁴Issues like goal achievability have been addressed in the behavioral accounting literature—see Manzoni and Merchant (1989) for example.

⁵Failure may be costly to the principal if the principal must contribute additional inputs for the next attempt, or if the principal is penalized (each successful period) for not delivering output on time.

Achieving Objectives

An advantage of this modeling approach is that it provides a natural setting to consider issues like quality, productivity, and learning in the presence of incentive problems. This problem is, in fact, a model of stochastic quality and moral hazard where the firm's objective is to produce output at or above an exogenously-specified design standard.⁶ The firm's proficiency at producing at this standard is known as its *quality of conformance*. In this thesis, where the agent's level of effort determines the probability of success, the sequence of optimal actions through time determines the firm's ex ante quality of conformance.

The facts that (1) the length of the optimal contract is endogenous and (2) the length of the realized contract is random make the setting analyzed here fundamentally different from those modeled in other multi-period agency papers, like Lambert (1983) and Ramakrishnan (1988). Here, the principal has a sense of urgency that is absent in those papers. That is, the principal must determine the optimal trade-off between (1) compensating the agent to take high levels of effort in early periods to increase the likelihood of accomplishing the goal sooner to save input costs and (2) the alternative of paying lower levels of compensation over a potentially longer period of time at the risk of higher failure-related costs.

The result that the contract length will be finite holds even when the principal has an infinite planning horizon and even if the first-best contract has a (possible) infinite length. In fact, the maximum length of the second-best contract decreases as the level of exogenous, failure-related costs increase. For high enough costs, a zero defect policy is optimal. A *zero defect policy* (ZDP) guarantees first-period success by requiring the agent to take the maximum level of effort in the first period, and there are situations when

⁶Generally, this design standard is a choice variable since revenue, input costs, and the probability of successfully producing a unit depends upon desired level of functionality or consumer satisfaction.

Achieving Objectives

the second-best contract induces an immediate success through the implementation of a ZDP, but first-best contract has an infinite-horizon.

This result may seem surprising in the context of other multi-period agency papers like the ones listed above. In those papers, the principal determines the optimal way to structure wages and actions over the course of a fixed-length contract given the available information structure (i.e., the informativeness of the available signals). When the length of the contract is fixed, the benefits to the principal increase if the number of periods can be increased. In finite-period problems, like Lambert (1983) and Rogerson (1985), these gains result from the principal's ability to smooth the agent's consumption over a larger number of periods and from the increase in the number of signals about the agent's effort levels, whereas in infinite-horizon problems, like Rubinstein & Yaari (1983) and Radner (1985), the principal benefits by having long samples of past outcomes, i.e., better information, and by having the ability to punish the agent over long periods of time for any observed deviations (at least at low discount rates).

Furthermore, with specific assumptions regarding the agent's utility function and the periodic input costs, I find that when effort is unobservable, periodic wages decrease over time (as failure costs increase) and the level of induced action increases. For example, when success is achieved early (at a low cost), the agent receives a bonus in the form of a high periodic wage, whereas for a late success (associated with high failure-related costs) he is subjected to performance penalties and thus receives a smaller (periodic) payment for his services. Reichelstein (1992) discusses incentive schemes for government contracts which follow a similar pattern. He analyzes a *cost-plus-incentive-fee* based contract in which the agent's net (cumulative) compensation (after deducting out-of-pocket costs) is a decreasing function of total out-of-pocket costs; however, Reichelstein ignores the multi-period nature of such settings—that failures cause cost over-runs and

Achieving Objectives

take time to correct. Thus, he ignores the time dimension of such contracts and the fact that the agent may receive compensation in the interim periods prior to the completion of the task. These interim payment may take the form of salaries, retainers, or progress payments. They are referred to as failure wages here. Although these periodic wages are decreasing over time, total compensation need not be decreasing in total costs (as it is in Reichelstein's paper).

I also extend the results of the basic model by analyzing the case when the agent learns over time. Learning occurs when the agent increases his proficiency over time—when the probability of success for a fixed level of effort (or equivalently, for a fixed level of disutility) increases over time. When the employee is capable of learning, the firm must determine the optimal trade-off between the effort and pay in earlier periods and effort and pay in succeeding periods when the employee is more efficient. Here, again, there is a trade-off between getting done sooner, in expectation, to save input costs and getting done later to reduce wage-related costs. In the standard, fixed-period model, this trade-off—and thus the analysis of learning—is irrelevant because there is no sense of urgency to complete the task. The efficiency gains can result from two types of learning. First, they can arise (almost) autonomously over time when little or no learning effort is required; this type of learning is referred to as learning by experience. Alternatively, they can result from the deliberate attempt by the agent to improve his efficiency; this type of learning is referred to as costly learning.

In Chapter IV, I analyze learning by experience in both the observable and unobservable action cases (I also present an example of a costly learning problem in Appendix 4)⁷. In the case where action is unobservable and the agent commits to completing the task, the

⁷A separate literature review of learning-related papers appears in Chapter IV.

Achieving Objectives

main result of the finiteness of the contract length holds, and because the learning process affects costs in a similar manner in both the observable and unobservable action cases without commitment, I obtain the same results in both cases: (i) holding future efficiency levels constant, agents who are more efficient through past training or experience work harder in a period than inefficient agents, and (ii) as the amount learned over any future interval increases, the optimal action at the beginning of the interval decreases.

The first result of the no commitment case may seem somewhat surprising since one might expect that less efficient agents would be required to work harder to compensate for their inefficiency, yet it means that there are two benefits from hiring more efficient agents: (1) for fixed effort levels, they are more likely to succeed, and (2) in equilibrium, they work harder. Thus, this result may provide at least some explanation as to why domestic corporations are interested in improving the education system in the United States, and also why anecdotal evidence suggests that German firms tend to have high quality but seem to place little emphasis on traditional quality programs.⁸

The second result, that as agents learn more in the future, the optimal level of effort in the current period decreases, is more intuitive. As the agent becomes more efficient in the future, costs in those periods decrease; therefore, marginal expected failure costs in prior periods decrease, and so activities undertaken to reduce those failure costs decrease.

With these two learning results, it is easier to understand the nature of training when the training activity derives no current productive benefit. In fact, I define training to occur when the owner instructs the employee take no productive effort in the initial period(s). The agent must still be hired during the training period and, depending upon his utility

⁸See "The Riddle of German Quality," by Jerry Bowles, in *Across the Board*, published by the Conference Board, (January/February, 1993).

Achieving Objectives

function, may require compensation during that time. It is assumed that the agent learns by watching the process fail or by watching a similarly costly event, like a seminar. Note that training is optimal when the agent is very inefficient initially but overcomes his inefficiency rapidly.

Chapter II. provides an in-depth literature review. Chapter III. contains a general description of the agency setting and analyzes the observable-action and unobservable-action cases without learning. Chapter IV. discusses relevant learning-related papers and the analysis of the problem when the agent can learn by experience and Chapter V. summarizes the results and discusses planned future research. Except where noted, theorems are stated in the text and then restated and proved in Appendix 2.

Chapter II - Literature Review:

The control problem of achieving objectives in a stochastic environment with an infinite planning horizon is obviously a multi-period one. While there are no other papers which consider this specific problem, several papers have analyzed aspects of multi-period agency relationships. In these papers, the agency's duration, whether infinite or finite, is assumed exogenous and then the characteristics of the optimal contract are derived.

Radner (1981) was one of the first papers to present an analysis of a finite, multi-period agency. In it, the author studies noncooperative equilibria of multi-period relationships which produce (what appear to be) cooperative outcomes. Radner's main result is that for a (long) T -period relationship, for any single-period, Pareto-optimal, cooperative equilibrium that dominates a single-period, Nash (noncooperative) equilibrium and any

Achieving Objectives

positive number, $\epsilon > 0$, there exists an ϵ -equilibrium of the T -period relationship that yields both players an average expected utility greater than their utilities in the single-period, cooperative equilibrium minus ϵ .⁹ So, for long, finite relationships, cooperative outcomes can be produced from an approximate noncooperative equilibria—or for any ϵ , there is a (finite) T large enough so that the noncooperative equilibria have cooperative outcomes in each subgame.¹⁰

To sustain such relationships, there must be a method to 1) detect cheating or shirking, and 2) punish such (dysfunctional) behavior. The method used by both parties to punish antisocial behavior is referred to as a *trigger* strategy. It involves cooperating as long as one infers that the other party is cooperating. For the risk-averse agent, such a strategy is easy to formulate: choose the first-best level of effort as long as the principal offers a constant wage (over outcomes); if the principal imposes risk, choose the second-best level of effort for as long as such a contract is offered. Alternatively, the principal offers the first-best wage until he can detect shirking on the part of the agent. If shirking is detected, he then offers the second-best wage scheme for the remaining periods of the contract; thus, such a strategy is “harsh” in the sense that once shirking is inferred from the history of past outcomes, cooperation dissolves for the remainder of the relationship. So, the inference rule must balance the costs of allowing the agent to shirk—by using a wide range of acceptable average outcomes—with the costs of eliminating cooperative behavior (when it should not be eliminated) with a narrow range of acceptable average outcomes.

⁹An ϵ -equilibrium is an approximate equilibrium when each player's sequential move is within ϵ in utility of being the player's best response.

¹⁰Backward induction fails to unravel such an equilibrium because the parties are assumed to evaluate long-term average utility over the entire length of the contract—not just the remaining length; so, the marginal benefits of defecting in the last period are small for a large T , and cooperation is maintained.

Achieving Objectives

Rubinstein and Yaari (1983) also consider a multi-period moral hazard model. They attempt to explain the observation that insurance companies offer discounts to clients with good records (with respect to past claims) by arguing that the discounts provide a mechanism to allow both parties to counteract the inefficiency which arises from the moral hazard problem.

They try to de-emphasize the role of periodic (incentive) wage schemes by concentrating on the time structure of the problem and on the case where the agent's actions in period t do not affect compensation in period t ; thus, they study full-indemnity insurance contracts, i.e., contracts without deductibles. They do allow the price of the contract to vary over time; so, although the agent's actions do not have short-term implications, they do have long-term implications. They find that giving the insurer (long-term) pricing flexibility is sufficient to eliminate the inefficiency created by the moral hazard problem—the unobservable level of care that the insured should undertake—as long as the insured is risk-averse. Thus, their problem is similar to Radner's (1981), but, as will be described below, their solution provides more flexibility.

Their result is obtained by assuming that utility for both parties is calculated (and evaluated) as the long-term expected periodic payoff, and that the insurer commits to an announced, long-term strategy. This strategy requires him to analyze the history of past claims to make certain that, on average, the percentage of claims is consistent with the percentage expected if the proper (first-best) level of care was exercised. If it is, then the insurer offers the agent a “no claims” discount price—or more precisely—an expected claims, discount-off premium. If the average level of claims is too high, then the agent is charged a higher (penalty) premium until the history of claims is consistent with the proper level of care. So, unlike Radner (1981) in which his trigger strategies were harsh—and lasting—in Rubenstein & Yaari, the principal has the flexibility to reintroduce a cooperative pricing policy after a period in which the discounts had been

Achieving Objectives

eliminated. In the long-run, the insurer will have sufficiently large samples on which to make inferences about the agent's average level of care as well as the flexibility of an infinite horizon over which to penalize the agent. Since the agent does not discount his periodic utilities, such penalties—which can only be applied in the long-term—induce the agent to take the first-best level of effort.

Radner (1985) studies an infinite-horizon, repeated game where the parties make decisions based upon their long-run average expected utility, and the agent discounts periodic utility. He finds that efficient (first-best) behavior is a Nash equilibrium in the infinite-horizon game if it is Pareto-superior to the one-period Nash equilibrium. Employing the same ϵ -equilibrium concept that he applied in Radner (1981), he finds that for sufficiently low discount rates—where the parties discount their future expected utilities—there is an ϵ -equilibrium which approximates the first-best outcome.

Again, in this paper, the principal can eventually detect any systematic shirking on the part of the agent by comparing the agent's average output with what would be expected if the agent had been selecting the first-best level of effort in each period. A simple dichotomous contract is optimal in which the agent is offered the first-best sharing rule in period t if his average performance in periods one through $t-1$ has been acceptable and a penalty contract if his performance has been unacceptable. With low discount rates, the future penalties that can be imposed upon detection and during the infinite horizon are sufficient to prevent the agent from shirking. Likewise, with a long history of past outcomes, the uncertainty about the agent's hidden actions can be completely diversified away and so the incentive problem can be completely eliminated.

Lambert (1983) considers finite, multi-period relationships where the number of periods are not large enough to induce cooperative behavior. He examines the role of long-term

Achieving Objectives

contracts in controlling moral hazard problems when both the risk-averse principal and risk- and work-averse agent precommit to the contract.¹¹ As in all these papers, the moral hazard problem exists because the principal cannot perfectly infer the level of effort from observations of periodic outcomes. Like Radner and Rubenstein and Yaari, Lambert finds that periodic wages will depend upon the current outcome as well as the sequence of past outcomes since such a contract optimally reduces the risk imposed on the agent by allowing the principal to make inferences based upon a larger sample of the agent's chosen actions (rather than just on one period's). Lambert labels this as the diversification effect since if the agent chooses a constant level of effort each period, periodic outcomes would be independently and identically distributed, and the sample variance would decrease as the number of periods increase thereby allowing the principal to make more precise inferences about the agent's average level of effort. Because wages are a function of all outcomes to date, periodic actions will also be parameterized by the past sequence of outcomes.

Unlike in the cases where the relationship lasts a (relatively) long period of time—even an infinite period of time—in a problem with T (finite) periods, the principal cannot (completely) eliminate the uncertainty regarding the agent's sequence of effort choices, and at any time t , the principal has only $T - t$ in which to penalize the agent for shirking. Thus the optimal contract is more complicated than a simple dichotomous one, and it will not be simply a function of the average outcome through period t .

Somewhat related to this thesis, Lambert notes that, "The diversification effect suggests that the more periods the agency relationship lasts, the more the incentive problem is alleviated. In fact, we can show that there is value to extending the contract." These

¹¹He also investigates the problem that arises when the agent does not precommit to work for the duration of the contract.

Achieving Objectives

gains result from the principal's ability to smooth the agent's consumption over a larger number of periods and from the increase in the number of signals about the agent's effort levels. However, this result does not obtain here where the agent must accomplish an objective and where failure to do so is costly to the principal; here, the principal must determine the optimal trade-off between (1) compensating the agent to take high levels of effort in early periods to increase the likelihood of accomplishing the goal sooner to save input costs and (2) the alternative of paying lower levels of compensation over a potentially longer period of time at the risk of higher failure-related costs.

Rogerson (1985) also studies a multi-period moral hazard problem and analyzes the relationship between wages and expected wages in consecutive periods. He finds that if the inverse marginal utility function is convex (concave), then conditional on the first period's outcome, the first period's wage is greater than or equal to (less than or equal to) the expected second period wage; therefore, the unconditional expected first wage is greater than or equal to (less than or equal to) the expected second period wage.

Lambert (1985) investigates a similar issue related to the sequence of wages and actions during a contract. He shows that income smoothing may be optimal where the agent "smooths" real earnings in a two-period model by trying to hit a two-period ex ante target total; thus, he notes that an income smoothing strategy would make second-period effort a decreasing function of the first period's outcome, and that this would induce a negative correlation between the first and second period outcomes. By assuming that the agent has a square root utility function, he provides an example where the optimal compensation scheme offered by the principal causes the agent to 'smooth' output.

However, supplementing the work of Rogerson, Ramakrishnan (1988) shows that this smoothing result is not true in general. In fact, income acceleration—inducing the agent

Achieving Objectives

to take a higher level of effort in the second period of a two-period contract to increase expected periodic income when realized first period income is (relatively) high—may be optimal. Ramakrishnan notes two costs for the principal associated with inducing increased effort: (i) a *wealth effect*, or the compensation for the additional effort, and (ii) the *risk effect*, or the risk premium for the additional risk introduced by making compensation scheme more pronounced. He notes that if the agent's utility function exhibits decreasing absolute risk aversion (DARA), then an increase in utility (for wealth only, not net utility) from paying (expected) higher wages may in fact reduce the risk premium that needs to be paid to the agent. If these risk effects are greater than the wealth effects, then acceleration could be optimal. In fact, his main result shows that the optimality of smoothing or acceleration depends upon the shape of the agent's marginal inverse utility function. If this function is convex, income smoothing is optimal. If it is concave, then there exists a problem where income acceleration is optimal.

Along a different vein of the theoretical accounting literature, Reichelstein (1992) discusses incentive schemes for government contracts, specifically a *cost-plus-incentive-fee* based contract in which the agent's net (cumulative) compensation (after deducting out-of-pocket costs) is a decreasing function of total out-of-pocket costs. The purpose of his article is to show that such an example is an application of agency theory to contract design.

Such a contract turns out to be (somewhat) similar to the optimal contracts found in this thesis; however, in his application, Reichelstein ignores the multi-period nature of such settings—that one cause of cost over-runs is the failure to complete subgoals, or specific tasks associated with the project, and that such failures take time to correct. Thus, the time dimension, which is a contractible variable, is ignored for evaluation purposes as well as consumption purposes. In fact, we observe that in many settings, the agent may

Achieving Objectives

receive compensation in the interim periods prior to the completion of the task. These interim payments, which may take the form of salaries, retainers, or progress payments, are referred to as failure wages here, and although they are decreasing over time, total compensation need not be decreasing in total costs (as in Reichelstein's paper).

Like my analysis, Fudenberg, Holmstrom, and Milgrom (1990) provide a possible explanation for the observed variation in the length of long-term contracts. They show that long-term contracts are valuable only if optimal contracting requires commitment to a plan today that would not otherwise be adopted tomorrow. They show that such commitment is unnecessary if the following conditions are met: (1) all public information can be used in the contract, (2) the agent can access a bank on equal terms with the principal, (3) recontracting takes place with common knowledge about technology and preferences, and (4) the frontier of expected utility payoffs generated by the set of incentive compatible contracts is downward sloping at all times.

The motivation for their analysis is to attempt to explain the observed variation in the length of employment contracts. They note that long-term contracts are advantageous if these contracts increase a risk-averse agent's ability to smooth consumption over time; however, they observe that (it seems that) higher paid workers have greater access to capital markets than lower paid workers do; so, lower paid employees should be involved more often in long-term contracts than higher paid employees, but this hypothesis seems to contradict empirical observations. Thus, they conclude that long-term contracts must occur to alleviate problems caused by information asymmetries, and that these asymmetries are more prevalent at higher levels than at lower ones. That is, the authors note.

Achieving Objectives

“Many more of the activities of managers than of factory workers or salesmen contribute directly to future [periods'] production in ways that are not reflected in current performance measures. Long-term contracts, which await the arrival of additional information on current activities, are important in managerial contracting but not in contracting with workers for whom current observations are sufficient for evaluating current performance.”

So, the authors conjecture that the benefits of extending the contract's length are positively related to the length and extent of the information lag, and so they note that one must balance the costs of increasing the contract's length with the benefits of incorporating additional information. However, the determination of the optimal length of the contract is outside the scope of their analysis.¹²

Joskow (1987) studies the variation in contract lengths between coal suppliers and electric utilities. He provides empirical evidence that supports the hypothesis that these parties make longer commitments (and rely less on repeated negotiations) when relationship-specific assets are important. Such commitments guard against one party's ability to act opportunistically once the other has sunk costs into an asset which has limited alternative uses and mobility. In this thesis, Corollary 3.2 provides an alternative explanation for varying contract lengths: the optimal maximum length of a goal-based relationship is a nonincreasing function of exogenous failure costs.

¹²In this thesis, it is shown to be an inverse function of the periodic input cost.

Chapter III - Analysis of the Problem (without learning):

III.A. Basic Assumptions:

I frame my analysis as a multi-period, principal-agent model in which the agent is hired to achieve an objective and, depending upon the derived contract, there is a possibility that the agent will not immediately accomplish the task. When the agent's failure to achieve the objective is costly to the principal, the agent may be required to act until the task is successfully completed: thus, in general, the optimal contract will have a random length. When the agent's effort is unobservable, an incentive problem exists because the principal desires that the agent work hard (to reduce the probability of paying failure costs), whereas the strictly work-averse agent prefers to work at a lower level of effort and risk the possibility of taking additional effort at that lower level.

Formally, I assume that the principal desires to hire an agent to produce one unit of good output and that there are two possible levels of quality: good and bad. If the principal's expected utility in any period is nonnegative, then he wants the agent to act until one unit of (good) output is produced.

For simplicity I assume that the stochastic production function in period t —the probability of producing a good unit in period t —is the agent's action, $a_t \in [0, 1]$ and that it is possible for a_t to equal 1. $a_t = 1$ in the action set means that the principal knows (with certainty) that the agent is capable of accomplishing the task.¹³ A zero defect policy occurs when the optimal first-period action is equal to one, when $a_1 = 1$.

¹³When it is optimal for the agent to do so.

Achieving Objectives

The agent's total utility function is assumed to be additively separable in wealth and effort and is constant and additive over time. Let $U(\cdot)$ denote his periodic utility for wealth and $V(\cdot)$ his periodic disutility for effort. $U(\cdot)$ is strictly concave, and, unless otherwise noted, $V(\cdot)$ is strictly convex. Neither the principal nor the agent discounts future utility although the results are unchanged if they do.

When effort is observable, the agent receives a periodic wage, w_t . When effort is unobservable, if the agent produces a good unit in period t , he receives a success wage, $g_t \geq 0$, which depends upon t ; and the contract ends. When he fails (by producing a bad unit in period t) he receives a time-dependent failure wage, $b_t \geq 0$; the contract continues; and the principal must supply another unit of input at a cost of $c > 0$.¹⁴ The incentive problem exists because the risk-neutral principal wants the agent to work hard to reduce the probability of paying an additional c dollars next period, but, *ceteris paribus*, the strictly risk- and work-averse agent desires to take low levels of effort and risk the possibility of repeating the task at those lower levels of effort.

III.B. Observable Action Case:

When effort is observable and the worker acts as directed, then he is not responsible for the quality of output, and thus must be adequately compensated for performing failure-related work, i.e., for work after the first production attempt. If the agent is strictly work-averse, then his commitment status is irrelevant, and either a zero defect policy is optimal or a stationary policy is optimal. The following theorem presents characterizations of the optimal contracts in both the committed and uncommitted agent cases. Here commitment by the agent means that the agent promises to work until a success is achieved, and

¹⁴The assumption that exogenous costs are constant over time is for simplicity, only. In general, periodic failure-related costs may increase, hold constant, or decrease over time.

Achieving Objectives

commitment by the principal means that the principal promises to employ the agent only until a success is achieved (and no later).¹⁵

Theorem 1: In the first-best case—with or without precommitment by the agent—if the agent is strictly risk- and work-averse, (a) a zero defect policy is optimal if c , the constant periodic input cost, is greater than or equal to c^f where c^f satisfies

$$c^f = w'(1) - w(1), \quad (1)$$

and $w(a)$, as described below, is the function which describes the wage required to induce the agent to participate and take action level a as directed.

(b) if an interior solution obtains, the optimal contract is a stationary policy.

(i) Without commitment, it is described by:

$$w'(a^*) = \frac{c + w(a^*)}{a^*}. \quad (2)$$

where $\{a^*, w^*\}$ is the optimal periodic action-wage pair and both variables are independent of t , and $w(a)$ is a composite function described below.

(ii) With commitment, the optimal pair $\{a^*, w^*\}$ is described by:

$$\lambda = \frac{1}{U'(w^*)} \quad (3)$$

¹⁵The necessity of this parenthetical qualification is explained in the next section.

$$\lambda = \frac{\left(\frac{w^* + c}{a^*} \right)}{V'(a^*)} \quad (4)$$

Please note that equation (2) or both (3) and (4) combine to characterize the solution.

The part (a) of the above theorem follows directly from part (b); thus, I begin by discussing part (b.i). Theorem 1.b.i is proved by recognizing that without commitment, the principal's problem is a simple stochastic dynamic programming one. The environment is stationary and Markovian; after a failure occurs, the principal's costs, benefits, and available actions are all unchanged, and the past history of performance has no effect on the current probability of success. So, the optimal action-wage pair, $\{a^*, w^*\}$, must satisfy the following periodic participation constraint:

$$U(w^*) - V(a^*) \geq 0. \quad (5)$$

Since the principal will never overcompensate the agent, constraint (5) will always hold with equality; thus, we can define a function $w(a)$:

$$w(a) = U^{-1}[V(a)]. \quad (6)$$

Since the agent's inverse utility function, $U^{-1}(\cdot)$, and disutility function, $V(a)$, are both increasing and strictly convex, the composite function $w(a)$ is increasing and strictly convex. Thus, $w(a)$ is unique for every a , and constraint (5), through the function $w(a)$,

Achieving Objectives

may be substituted into the principal's objective function to yield the following problem:¹⁶

$$\max_a \left\{ r - \frac{[w(a) + c]}{a} \right\}. \quad (7)$$

Differentiating this problem with respect to a gives condition (2). Condition (2) is similar to the first-order conditions which describe the *economic conformance level* in a quality cost analysis (see Fine (1986) for example). Quality costs are usually divided into four categories: prevention, appraisal, internal failure, and external failure.¹⁷ Here, the action, a , is both a prevention activity and the economic conformance level, whereas expected failure costs are easily identified by rearranging the principal's objective function:¹⁸

$$w(a) + c + \underbrace{(1-a) \frac{c + w(a)}{a}}_{\text{expected failure costs}}. \quad (8)$$

Within a quality-cost framework, condition (2) states that at an (interior) optimum, marginal prevention costs equal marginal failure costs. In one respect, this model is more general than the typical economic conformance level model because there is no guarantee that first-period failures can be corrected immediately at a certain cost—there is a chance that the failure-related activities will fail which is why the problem has an infinite horizon.

¹⁶As developed in the proof of Theorem 1 in Appendix 2.

¹⁷Prevention costs relate to those activities which increase the likelihood of success (i.e., reduce the probability of failure); appraisal costs relate to the inspection activities which attempt to discover defective items (failures); internal failure costs involve those activities concerned with correcting or disposing of failures discovered before shipment; and external failure costs relate to correcting or disposing of failures which were discovered after shipment.

¹⁸Alternatively, a can be interpreted as the level of inspection work, and c as the cost of a Type II error.

Achieving Objectives

Theorem 1.b.ii shows that if the agent is strictly work-averse, then the solution to the problem with commitment is identical to the solution with commitment. Thus, commitment by the agent is not valuable to the principal; that is, the principal cannot take advantage of the additional degrees of freedom. This result occurs because the agent's total utility function is concave in both wealth and effort; thus, requiring the agent to take a constant action and paying him a constant amount each period minimizes the principal's expected costs to succeed.

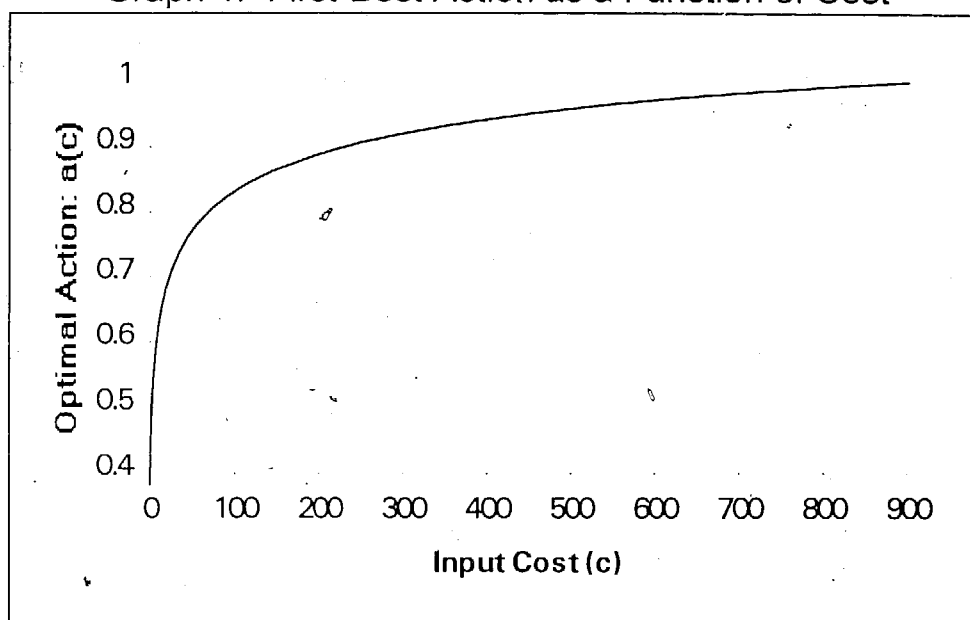
The validity of part (a) of Theorem 1 follows from the following (very intuitive) observation: as the principal's exogenous failure costs—that is, the nonwage costs—~~increase~~, his desires for early success increase (and thus he instructs the agent to take a higher level of effort), and thus we have lemma 1.1 which is shown for the case of constant (failure-related) input costs.

Lemma 1.1: As the level of (future) production costs increases, the optimal prevention action increases, i.e.,

$$a'(c) = \frac{1}{a(c)w''(a(c))} > 0. \quad (9)$$

The following graph provides an example of lemma 1.1 where the agent's utility functions are defined to be $U(w) = \ln(w)$ and $V(a) = ae^{1.5a}$.

Graph 1: First-Best Action as a Function of Cost



Note that in Graph 1, a zero defect policy is optimal at $c = 901.89$.¹⁹ This is the point, c^f , described in part (a) of the theorem for which a ZDP is optimal. In other words, to paraphrase quality “guru” Philip Crosby, quality—in terms of quality of conformance—may indeed be free with high enough failure-related input costs.

Theorem 2: When a ZDP is optimal in the first-best case, it is also optimal in the second-best case; however, the converse is not true. (With unobservable effort, a ZDP is optimal at input costs of less than c^f .)

The validity of this theorem is easy to see. The first-best contract must (at least) weakly dominate the second-best contract. When a zero defect policy is optimal in the first-best case, the same contract can be induced in the second-best case since the support is not fixed at $a_1 = 1$ (any deviations can be discovered with positive probability). Every other


¹⁹Although the wage to induce a ZDP is only 88.38, because the principal is a risk-neutral, expected-utility maximizer a ZDP is not optimal until periodic costs are greater than or equal to 901.89.

Achieving Objectives

feasible, unobservable-action contract imposes risk on the agent since, in these contracts, the agent's wage is outcome-dependent, and there is a chance of failure, $a_1 < 1$. Since the agent is risk-averse, and the principal is risk-neutral, each of these contracts is dominated by the first-best contract which, in this case, is a ZDP (and which is attainable in the second-best case). That the converse is not true can be seen by noting that in the second-best case, a ZDP is optimal at input costs of less than c^* because at these levels of cost it is cheaper for the principal to compensate the agent with a fixed wage for taking the highest level of effort than it is for him to impose risk on the agent (to induce the agent to take a relatively high level of effort but less than the maximum level of one).

The following graph shows c^* , the minimum level of cost to induce a ZDP in the first-best case, when the agent's utility functions are $U(w) = \ln(w)$ and $V(a) = ae^{\beta a}$ and $\beta \in [0, 1.5]$. When $\beta = 1.5$, a ZDP is optimal when input costs are greater than 901.89 as in Graph 1.

Graph 2: Cost Required to Induce a ZDP



In summary, in the first-best case, we observe that when the agent is strictly risk-averse and strictly work-averse, (1) commitment status is irrelevant; (2) an interior optimum is stationary so the optimal contract consists of a series of identical, single-period contracts which may go on forever; (3) there exists a high enough level of cost to induce a ZDP;

Achieving Objectives

and (4) when a ZDP is optimal in the first-best case, it is also optimal in the second-best case, but the converse is not true.

III.C. Unobservable Actions (Committed Agent):

When effort is unobservable and the agent does not precommit to complete the task, either a ZDP or a series of identical single-period contracts is optimal; this result occurs because there are only two possible outcomes and because the relationship ends when a good unit is produced. Examples of such contracts are piece-work or sales commission contracts. Because these contracts are derived in a manner similar to the first-best contracts, they are not presented here; instead, an example of one is presented in Appendix 3.1.

With precommitment by the agent, the agent agrees to work until the task is successfully completed. This differs from the other multi-period agency papers which focus on the agent either (1) committing to work for a particular number of periods or (2) agreeing to participate in a sequence of contracts. With this type of precommitment, the agent is responsible—held accountable—for completing the task. Because the contractual relationship ends once the objective has been met, the actual length of the relationship is random; so, unlike in other multi-period models, the principal is not responsible for the agent's welfare for fixed length of time.

Precommitment by the principal means that: (1) the principal commits to employing (and paying) the agent a periodic wage—which will vary over time and outcome—until the project is completed, but (2) he, the principal, cannot credibly commit to compensate the agent in periods after a success has been achieved.²⁰

²⁰The above definition of precommitment allows for the existence of a solution to the problem. Without this specification, there is no optimal contract because there is no minimum expected cost for the contract.

Achieving Objectives

The principal's profit maximization problem—written in recursive form—is:

$$\max_{\{a_t, g_t, b_t\}_{t=1}^T} \{a_t(r - g_t) + (1 - a_t)(\Psi_t - b_t) - c\} \quad (10)$$

subject to:

Participation:

$$a_t U(g_t) + (1 - a_t)(U(b_t) + \Gamma_t) - V(a_t) \geq 0 \quad (11)$$

Incentive Compatibility:²¹ (for every t)

$$(g_t) - U(b_t) - \Gamma_t - V'(a_t) = 0 \quad (12)$$

Where a_t is the agent's action in period t , $c > 0$ is the periodic input cost, and $g_t \geq 0$ and $b_t \geq 0$ are period t 's success and failure wages for producing good and bad output in period t , respectively, and

$$\Gamma_{t-1} = a_t U(g_t) + (1 - a_t)(U(b_t) + \Gamma_t) - V(a_t) \quad (13)$$

$$\Psi_{t-1} = a_t(r - g_t) + (1 - a_t)(\Psi_t - b_t) - c \quad (14)$$

are the agent's and principal's conditional expected future utilities given that failures have occurred in all periods prior to period t , respectively.

The main result of this section is the following characterization of the optimal contract.

To see this fact, consider the case where the principal may compensate the agent for all periods after a success. Suppose he instructs the agent to succeed in the first period (which minimizes input costs) and then compensates the agent some amount $w(T)$ for T periods to ensure participation:

$$T \cdot U(w(T)) = V(1).$$

When the agent's periodic reservation utility is zero, no finite T solves this problem since as $T \rightarrow \infty$, we know that $U(w(T)) \rightarrow 0$ which means that $w(T) \rightarrow 0$; so, the cost to induce a ZDP goes to zero.

Besides using the above definition of precommitment, setting a positive reservation utility, adding discounting to the problem or allowing for negative utilities will eliminate this nonexistence problem, and thus a solution will exist when either (1) the agent has a higher discount rate than the principal, or (2) the lower bound of the agent's utility function is less than the periodic reservation utility.

²¹Using the first-order approach.

Theorem 3: Assume that wages are constrained to be nonnegative, that periodic failure costs are positive, and that the strictly risk-averse and strictly work-averse agent's utility function satisfies: $U(0) \geq V(0) > -\infty$, then the optimal contract will have a finite maximum length, and success will be assured by the end of the contract, i.e., there exists T^* such that $a_{T^*} = 1$.²²

This result tells us that even with an infinite horizon, the maximum contract length is finite, and when the principal knows for certain that the agent is capable of accomplishing, he uses this information to set a deadline by which time the agent must succeed. Do we observe such truncated contracts in practice? Yes. Some real-estate agencies guarantee to purchase a property (at a specified price) if their agents cannot sell it within a particular period of time, and similarly, for some underwriting activities, investment bankers commit to buying all unsold shares. Similarly, we see that when some automobiles are deemed lemons, their manufacturers repurchased them. For example, after four recalls of its 1990 minivans, Nissan repurchased and destroyed 900 of them.²³ The following corollary extends the second result in this theorem to situations in which the agent's utility is unbounded below.

Corollary 3.1: If a finite-length contract is optimal when the agent's utility function is unbounded below, then $a_{T^*} = 1$.

Thus, at some point in time the agent is held accountable for his actions and, therefore, he must ensure success by that time, denoted T^* . As Corollary 3.2 shows, T^* will depend upon the periodic failure cost, c .

²²This result also holds when the agent is paid only a wage upon success, i.e., when there are no interim failure wages.

²³The Wall Street Journal (12/9/93).

Achieving Objectives

Corollary 3.2: As c increases, T^* , the maximum length of the contract, decreases.

This result is rather intuitive. Consider a cost-level χ which induces a contract of length T^* . As c increases above χ , the expected costs of failure increase. This increases the principal's desire (and his willingness to pay the agent) for the agent to exert more effort in earlier periods to reduce these expected costs. Since the principal operates in an infinite horizon, one can consider the agent's infinite vector of actions. At cost level χ , the vector of actions is:

$$[a_1, a_2, a_3, \dots, a_{T^*-1}, 1, 1, 1, \dots]$$

As c increases, $a_{T^*-1} \rightarrow 1$. As costs increase further, this process continues until a ZDP is optimal. Thus, this result provides an alternative explanation for Fudenberg, Holmstrom, and Milgrom's (1990) observation that, "Long-term contracts enjoy an obvious advantage if they expand the agent's ability to smooth consumption over time... But such a rationale can hardly explain the observed variation in the length of employment contracts."


Because of this cost of extending the maximum length of the contract, the results presented here differ substantially from those found in other multi-period, moral hazard papers. As was mentioned in Chapter II, in those papers, the principal takes advantage of additional periods to try to eliminate the agency costs. For example, Holmstrom (1979) notes, "When the same situation repeats itself over time, the effects of uncertainty tend to be reduced and dysfunctional behavior is more accurately revealed, thus alleviating moral hazard." And Lambert (1983) states, "The diversification effect suggests that the more periods the agency relationship lasts, the more the incentive problem is alleviated. In fact, we can show that there exists value in extending the contract." Indeed, this

Achieving Objectives

argument extends to the infinite-horizon agency problems, where Radner (1985) and Rubinstein & Yaari (1983) show that the first-best outcome can be attained by taking advantage of the entire horizon.

Using the same utility functions as previous examples in Sections III.B and III.C, the following graph shows the maximum length of the optimal contract for a series of costs $c \in [1, 901]$.²⁴ Recall that the periodic cost to induce a ZDP in the first-best case was 901.89, whereas here—in the second-best case—the minimum cost which induces a ZDP is 224.25; so, in the range $[224.25, 901.89]$ a ZDP is optimal in the presence of moral hazard, whereas without moral hazard the optimal contract has a possible infinite length.

Graph 4




²⁴Given that the principal's revenue is greater than total expected costs and so the project is worth undertaking.

Achieving Objectives

Here, however, it is the finiteness of the contract term and the threat of being forced to take the highest level of effort at the end of the contract—when $a_{T^*} = 1$ —which in my examples induces the agent to work harder in earlier periods to try avoid that high cost (i.e., disutility). It is this mechanism which helps alleviate the moral hazard problem since it aligns the agent's desires for an early success with the principal's. Thus, the agent responds not only to the prospective wage schedule, but also to the prospect of high levels of future disutility from effort. One expects that as the deadline approaches, the agent is willing to work progressively harder to avoid the ultimate penalty of $-V(I)$. The following graph shows the sequence of second-best actions and wages for the same agency that was analyzed in Section III.B of this Chapter.

Graph 5



We observe in the numerical example that periodic success and failure wages are decreasing over time. In fact, this leads to the following theorem, which shows that

Achieving Objectives

periodic failure wages are (always) a decreasing sequence in time, and the following conjecture.

Theorem 4: When $V(\cdot)$ is strictly convex, failure wages, b_t , are a decreasing sequence in time (and $g_t > b_{t-1}$).

Conjecture 4.1: Success wages are decreasing over time.

Conjecture 4.2: Actions are increasing over time.

In practice we observe that success wages are decreasing in time since it is often the case that bonuses are paid for early successes (or, equivalently, for keeping costs below expectations), and performance penalties are imposed for completion delays or cost overruns. However, decreasing failure-related wages are more difficult to observe in practice since they take the form of salaries or of retainers or, depending upon the nature of a long-term project, of progress payments. The following corollary shows that for many common utility functions, expected wages are also a decreasing sequence in time.

Corollary 4.1: If the agent's inverse marginal utility function, $\frac{1}{U'(\cdot)}$, is linear or convex, then expected periodic wages are decreasing over time.

Among the functions for which this corollary holds are the logarithmic utility and square root utility. If we consider the class of HARA utility functions which is defined by

$$(w) = \frac{(1-\gamma)}{\gamma} \left[\frac{\beta w}{1-\gamma} + \eta \right]^\gamma \quad (15)$$

Achieving Objectives

where w is the agent's wealth and $\gamma \neq 1$ and $\beta > 0$, then Corollary 4.1 holds when $\gamma \leq \frac{1}{2}$ or $\gamma > 1$. This set of parameters corresponds to those which induce the investigation of lower-tailed performance in Baiman and Demski (1980) and the optimality of income smoothing in Ramakrishnan (1988).

Showing that expected periodic wages are decreasing for any (general) utility function is not as innocuous as it might seem (given Theorem 4 and Conjecture 4.1). To see this, recall the definition of Γ_t , the agent's conditional expected utility at time t :

$$\Gamma_{t-1} = a_t^* U(g_t^*) + (1 - a_t^*) (U(b_t^*) + \Gamma_t) - V(a_t^*).$$

From Lemma 3.1, we know that the agent's periodic expected utility, Γ_t , is decreasing in t . Theorem 3 tells us that since the highest possible level of effort occurs in the last period, T^* , ceteris paribus, as t increases, Γ_t will decrease. So, it is possible for expected wages to increase, but for expected utility to decrease since expected utility includes the disutility of effort and the likelihood that a higher disutility of effort will be incurred in the future. In fact, it these above features of the solution which make conjectures 4.1 and 4.2 conjectures rather than theorems.²⁵

The results in this section depend critically on the assumption that the agent is strictly work-averse—that his disutility function, $V(a)$, is strictly convex. The following theorem illustrates that the principal's problem (with respect to the agent) is not just a consumption smoothing problem; it is also an effort-smoothing problem when the agent has convex disutility.

²⁵Appendix 3 provides examples using a variety of utility functions which provide some support for these conjectures.

Theorem 5: Assume $V(a) = a$, then the optimal contract has the following characteristics:

1. $b_1 = b_2 = \dots = b_{T^*-1} = g_{T^*-1} = w(t)$.
2. $a_1 = a_2 = \dots = a_{T^*-1} = 0$, and $a_{T^*} = 1$, and so
3. T^* solves the following program:

$$\min_t \{t[c + w(t)]\} \quad (16)$$

subject to:

$$tU(w(t)) - 1 \geq 0. \quad (17)$$

where $w(t)$ is the periodic wage required to ensure the agent's participation when no effort is taken until period t and then the maximum effort is required.

However, as long as the agent is strictly risk-averse, the following theorem holds, and this result allows us to prove conjectures 4.1 and 4.2 when $T^* = 2$ —which would occur at costs slightly below the minimum level which induces a ZDP in the second-best case.

Theorem 6: If $a_{T^*} = 1$, then $b_{T^*-1} = g_{T^*}$. (b_{T^*-1} is the failure wage paid in period $T^* - 1$ and g_{T^*} is the wage paid in the final period, T^*).

In other words, Theorem 6 states that the wage paid immediately prior to the deadline (in period T^*-1), the wage paid at the deadline (in the final period, T^*) are equal because at this point there is no reason to impose risk on the agent; success is assured in period T^* .

Corollary 6.1: If the maximum length of the optimal contract is two periods, then conjectures 4.1 and 4.2 are true.

Achieving Objectives

The proof of this corollary is simple: if given the agent's utility function and the periodic costs of failure, c , the maximum length of the optimal contract is two periods, i.e., $T^* = 2$, then it must be that $a_1^* < a_2^* = 1$, and from the first period's incentive compatibility constraint (equation 11) and Theorem 6, we have $g_1^* > b_1^* = g_2^*$. Note that for any risk- and work-averse agent, we can find a sufficiently high cost parameter, c , to make $T^* = 2$. Thus, it seems that the relationships we observe in the examples should be generalizable to cases when $T^* > 2$.


Regardless of the generalizability of Corollary 6.1, an application of the envelope theorem yields that the principal's expected cost to complete the project—ignoring sunk costs—will decrease over time. In the first-best case the principal's expected cost to completion is constant over time, and it is, of course, less than or equal to the principal's ex ante expected cost in the second-best case (as in the example in Appendix 3.1 where $15.42 < 18.44$). However, when a ZDP is not optimal, it may be the case that for some period t and all periods thereafter, the expected cost to completion in the second-best case is less than the expected cost to completion in the first-best case. In the example in Appendix 3.1, this occurs after the first failure.

Because failure wages are paid in the interim, depending upon the agent's utility function, total compensation after several failures and then a success may be higher than, say, after an immediate success (although the agent's net utility would still be lower in the former case). This observation has implications for empirical accounting research. If outsiders conjecture that firms pay for performance and they attempt to test this hypothesis by measuring the relationship between, say, earnings and executive compensation, the support they would find for their hypothesis would depend upon the portion of compensation related to profit levels, the portion related to accomplishing other non-profit related goals, and the executive's particular utility function since goal-based


Achieving Objectives

compensation can be positively- or negatively-correlated with the costs associated with achieving these non-profit-based objectives. Thus, claims by some in the popular press that show that incentive problems exist because executive pay is not correlated with financial performance ignore the fact that executives are rewarded for achieving other goals as well. In fact, as the following graphs show, one can hold all other assumptions constant and only change the periodic input cost to derive examples where total compensation is either decreasing or increasing in time. These graphs are derived using the same assumptions as before; however in Graph 6.A, exogenous costs are five per period, and in Graph 6.B, costs are two per period. In Graph 6.A, we see a positive correlation between *profits* and pay, whereas in Graph 6.B, we see a negative correlation between the two.

Graph 6.A



Graph 6.B



Thus, as a summary to this section, we see that when the agent's utility functions are bounded below and the principal knows for certain that the agent is capable of accomplishing the task, then he will impose a deadline by which time the agent must succeed. The maximum length of the contract, or the length of the interval before the deadline is imposed, is a function of the exogenous costs of production; that is, as costs increase, the length of the contract decreases (weakly). Furthermore, the agent's expected utility is decreasing over time, and for many common utility functions, the agent's periodic expected wages are decreasing over time.

Chapter IV - Learning by the Agent:

Learning is the act of gaining knowledge, understanding, or skill by study, instruction, or experience. In a productive setting, it results in the agent becoming more efficient (at accomplishing tasks) over time—more efficient in the sense that for a given level of effort or disutility, the probability of success increases over time. Learning can be the consequence of a deliberate attempt by the agent to improve his efficiency, or it can occur (almost) automatically through time (when little to no effort is required). The former type of learning is referred to as *costly learning* while the latter type is referred to as *learning by (or through) experience*.

Fine (1986) also discusses these two types of learning. He refers to costly learning as *induced learning* and notes that such learning depends upon “conscious actions and efforts by management...to increase the efficiency of the production system.” Costly learning occurs when the agent can improve his future proficiency (at accomplishing a task) by undertaking additional current-period effort from which he suffers disutility. In the context of this thesis, the costly learning problem becomes an interesting one when (1) the opportunity to learn occurs randomly and (2) the principal is ignorant of the opportunity when it arises for the agent. Without these conditions, there are no incentive issues regarding the agent’s decision to learn. With the exception of Amershi and Datar (1991) there has been little theoretical work on this topic. Amershi and Datar consider costly learning with moral hazard; in their model, the expected long-term benefits of learning are difficult to measure so the analysis is imbedded in an incomplete contracts setting.

Fine refers to learning by experience as *autonomous learning* and notes that it involves “quasi-automatic improvements.” In this chapter, I analyze the situation when the agent

Achieving Objectives

learns in this costless fashion. I show that, with commitment and unobservable actions, if the principal knows that the agent will eventually learn the task well enough to succeed at will, then the main result of Chapter III holds. In addition, I study both the observable and unobservable action cases when the agent does not commit to complete the task.

Learning by experience occurs in many jobs where workers discover short-cuts and process improvements as by-products of their productive effort (and costless observations). These improvements allow the workers to be more efficient over time. Other than in contrived examples, it would be difficult to find actual environments in which learning-by-experience does not occur. Regardless of the task—from digging ditches to performing surgery to writing research papers—experience matters, and we see that firms consider this variable in their decision-making. As Horngren, Foster & Datar (1994) write, “Predictions of costs should allow for learning....the effects of the learning curve could have a major influence on decisions.” For example, as the three authors note, the Kaizen (or continuous improvement) budgeting approach bases budgeted amounts “on process improvements that are *yet to be implemented*. [their emphasis]” Such decisions involve the consideration of initial levels of efficiency and expected levels of efficiency over time.

In the first-best case, such settings have been studied in the economics literature. For example, Spence (1981) considers the effect of learning on production decisions over time. In his thesis, the benefits from learning in any period are endogenous since they depend upon cumulative production to that point; that is, he relates unit costs to accumulated volume to show that additions to output reduce future costs. Thus the optimal level of short-term production with learning is higher than it would be in the absence of learning. In the stochastic environment studied in this thesis, the production process ends once a single unit of good output is produced; so, it is not possible to link

Achieving Objectives

efficiency to cumulative output; instead, it is assumed that the agent learns in a deterministic fashion over time. This approach is similar to Zeckhauser's (1974). Zeckhauser investigates the optimal interval for which a social planner should fix the level (or type) of technology when an employee learns how to become more efficient with any one technology over time. So, he studies the trade-off of gaining technological efficiency (by using a new process) with the cost of lost efficiency which results because the employee is less familiar with the new technology.

Fudenberg and Tirole (1983) also study learning-by-experience. They refer to it as learning-by-doing and note that practice makes perfect or that "through the repetition of an activity one gains proficiency." They investigate the effects of learning on market conduct and performance by considering both a monopolist's problem and a problem with strategic interaction.

Learning by experience actually depends upon two factors: being present during the phenomenon and exerting effort to understand what was observed. As Arrow (1962) notes: "Learning is a product of experience. Learning can take place only through the attempt to solve a problem and therefore takes place only during activity." I model learning by experience as simply being present and thus assume that it is costless for the agent to improve his efficiency each period; with this type of learning, repetitive exposure to the problem increases the agent's likelihood of success.²⁶

The benefits of learning—the increase in the agent's probability of success over time—are modeled in a simple multiplicative fashion. I assume that the probability of success in (learning) period t is $p_t a_t$ where each $p_t \in (0, 1)$, and the sequence

²⁶Since we cannot link it to cumulative output (of zero).

$$\{p_t\}_{t=1}^T$$

is increasing in t ($p_s = 1$ for $\forall s \geq T + 1$). Thus, each p_t can be viewed as a representation of the agent's cumulative knowledge to date (to period t)

I assume that learning ends after a finite, or T , number of periods. In problems without commitment, this assumption allows me to solve the principal's problem through backward induction (in both the observable and unobservable action cases). With the proper choice of parameters, $\{p_t\}$, it approximates the concave nature of the benefits associated with learning over time—as Arrow (1962) notes, “Learning associated with repetition of essentially the same problem is subject to sharply diminishing returns.”²⁷

When the learning horizon is infinite, optimal periodic actions and wages are difficult to characterize because the problem is a nonstationary one (the transition probabilities from the failed state to the good state change each period). By augmenting the state space—the two levels of quality—by the iteration number, the problem can be converted into a stationary one; however, it is difficult to characterize the optimal periodic policies.

When learning ends after T iterations, we can derive characteristics of the optimal contract. In the following section, I show that the main result of Chapter 3 holds under certain conditions, whereas in Section IV.B. I provide characterizations of the optimal contracts in cases without commitment.

²⁷Although I do not place any such restrictions on the parameter path in this version of the paper. The concave nature of the benefits of learning over time can be derived from typical assumptions and observations of learning curves—see Horngren, Foster, & Datar (1994), for example.

IV.A. Committed Agent (Unobservable Action)

With precommitment, moral hazard, and learning-by experience a finite contract which ensures success by its deadline—as in Theorem 3—is optimal if and only if the agent eventually learns how to accomplish the task in any period. That is, regardless of the length of the learning interval, if at the end of the interval the agent is a fully capable individual, then the principal will (eventually) impose a deadline on the agent.

IV.B. Learning Without Precommitment

Without precommitment there are two main phenomena to consider: (1) how does the optimal action change with respect to the agent's efficiency in the current period, and (2) how does the optimal action change with respect to the agent's efficiency in future periods?

By assuming that agent's learning ends after a finite number of periods, say T , the solution can be characterized because from period T onward the problem is stationary; so, the value function at iteration T is the salvage value in a finite-horizon problem. By working backward, each learning period's optimal action can be determined. The level of this effort depends upon the agent's current and prospective levels of efficiency. The following theorem shows that in both the observable and unobservable action cases, as the agent's efficiency increases within a period, the optimal action increases.

Theorem 7: As the agent's efficiency parameter increases in a period, the optimal action increases within the period, i.e., $a_t'(p_t) > 0$.


Achieving Objectives

In the observable-action case, in any period t , the following optimality condition holds:

$$w'(a_t) = p_t[\Psi_t] \quad (18)$$

where Ψ_t represents failure costs (expected production costs after period t). As in Spence (1981), on the optimal path in any learning period t , marginal short-run costs (the LHS) equal the present value of the total expected future-period cost reduction of an additional unit of effort in period t (the RHS).


One might expect that less skilled workers would be required to work harder to compensate for their deficiencies, but this does not happen because the benefits to the principal associated with inducing increased effort—the reduction in expected failure costs—are greater for agents with higher efficiency parameters than ones with lower efficiency parameters. Thus, agents who are more efficient as the result of past training or experience work harder in equilibrium than less efficient agents (when future levels of efficiency are held fixed). The following graph, where efficiency parameter $p(H) > p(L)$, illustrates this point:



Achieving Objectives

Theorem 7 implies that as the agent becomes more efficient, the probability of success increases in two ways: first, there is a higher chance of success for a given level of effort, and second, the agent actually exerts a higher level of effort in equilibrium. This observation can be seen in the following graph which is based on a single-period learning problem ($T = 1$).

Graph 8: The Optimal Action and Probability Success when $T = 1$.



Graph 8 also illustrates Corollary 7.1. For any $p_1 < 1$, the first-period's optimal action is less than the constant optimal action from period two onward.

Corollary 7.1: The optimal action and wage from the $(T+1)th$ period onward—after learning has stopped—is greater than the optimal action and wage in the Tth period.

Achieving Objectives

However, actions need not increase in time over a sequence of learning periods.²⁸ Although Theorem 7 states that, *ceteris paribus*, optimal effort within a period increases in efficiency this need not occur over the sequence of learning periods, because as the following theorem shows, increases in future efficiency parameters cause decreases in current-period actions.

Theorem 8: The higher the future efficiency parameter, p_s , for $s > t$, the lower the optimal action in period t , i.e., $a_t'(p_s) < 0$

In other words, the more that the agent learns in any future era—the larger the gains from future learning—the lower the level of effort at the beginning of that era. This result is intuitive: at the (interior) optimum, marginal prevention costs equal marginal failure costs. As increases in future efficiency reduce expected failure costs, prevention activities decrease.

If the efficiency gains from learning are large enough, it may be optimal to instruct the agent to take no action in a particular period—that is, to just watch the experiment. This occurs when input costs are relatively low and agent is initially inefficient but improves rapidly over time. Such an outcome can be interpreted as on-the-job training. It is important to note, though, that the agent must be hired during the training period even if he does not work. Depending upon his utility and disutility functions, he may be paid a positive wage even without taking effort. For example, if the agent has logarithmic utility, if $V(0) = 0$, and if his periodic reservation utility is zero, then he would be paid a wage of one unit since $\ln(1) = 0$.

²⁸They do appear to be increasing in time when the benefits from learning are concave in time—when $p_{t+1} - p_t < p_t - p_{t-1}$ for every t . I plan to prove this fact in the near future.

Chapter V - Conclusion and Future Research:

In this thesis, I analyze a setting in which the agent is hired to accomplish a task for the principal. Such a problem provides an accurate model of the planning and control problem that superiors must solve in decentralized organizations. In addition, the problem models similar goal-based relationships which exist within the economy, especially in job-shops, manufacturing departments, and among contractors and their clients.

This problem is fundamentally different from other multi-period agency problems, and a major advantage of studying it is that it provides a natural setting to analyze issues like quality, productivity, and learning in the presence of incentive problems. Analysis of the problem has led to several interesting results.

First, when effort is unobservable, the agent precommits to accomplishing the task, the maximum length of the optimal contract is finite and endogenous. This result may explain the fact that observed contracts come in a variety of lengths, because, here the maximum length depends upon failure-related input costs.

Second, the optimal contract requires the agent to take the highest possible level of effort in the final period of the contract if it is reached. It seems that the agent is induced to work harder in early periods, not solely through wage incentives, but also by the "threat" having to work even harder in the future. Thus, it is this threat of high disutility from effort combined with decreasing periodic wages over time which provides the mechanism by which the principal aligns the agent's interests with his own. One can view such a result as a *theory of deadlines*; when one knows for certain that the agent is capable of succeeding, the imposition of a deadline increases the agency's efficiency by inducing the

Achieving Objectives

agent to attempt to succeed in earlier periods (to avoid the pain of working extremely hard at the deadline. In addition, I have shown optimal failure wages are decreasing in time, and that for many common utility functions, expected periodic wages are also decreasing. Under specific assumptions, I have shown that optimal actions are increasing in time and success wages are decreasing in time. Furthermore, I have shown this to be true for general utility functions when the length of the optimal contract is two periods.

Third, when learning is introduced, within a period, more efficient agents (with higher levels of cumulative knowledge) work harder than less efficient agents. Additionally, as the benefits of learning in the future increase, the optimal actions in previous periods decrease. However, without additional assumptions, actions need not increase over time.

The research presented in this dissertation can be extended along at least two lines. First, a further investigation of costly learning opportunities in the presence of incentive problems can provide insights into the nature of performance evaluation in total quality programs. As was mentioned in Section IV., costly learning occurs when employees must make sacrifices—which are non-productive or even detrimental towards current period output—to earn long-term benefits. A costly learning problem is presented in Appendix 4.

Second, the results can be generalized to the case where, *ex ante*, achievability of the goal is unknown; that is, where neither the principal nor agent know whether it is possible to achieve the objective. Because they are both uninformed this is not an adverse selection problem. A paper somewhat along these lines is Hirao (1993) which investigates a multi-period agency in which the agent's first-period effort provides information about the long-term profitability of an investment. His paper does not consider this learning problem in an environment in which the agency attempts to accomplish a goal which is

Achieving Objectives

what I suggest. For example, such a model can be operationalized by multiplying each period's level of effort, a_t , by a parameter, $p \in \{0, 1\}$ to attain the probability of success. Of course, if $p = 0$, then the objective could never be achieved; so, it would be in the best interests of the principal to cancel the project and terminate the agent's employment.²⁹ In such a situation, the agency problem is compounded because the principal cannot infer whether failure is due to selecting an unattainable goal or because the agent is shirking.

²⁹If $p = 1$, then the goal is achievable.

Appendix 1

Notation

$U(\cdot)$	agent's strictly concave utility function
$V(\cdot)$	agent's strictly convex disutility function
c	cost (of exogenous inputs) per iteration
a_t	action taken in period t ($\in [0, 1]$)
b_t	wage paid upon failure in period t
g_t	wage paid upon success in period t
p_t	efficiency parameter $\in [0, 1]$; an increasing sequence, $\{p_t\}$, indicates that the agent learns through experience

Appendix 2: Theorems & Proofs

Theorem 1: In the first-best case—with or without precommitment by the agent—if the agent is strictly risk- and work-averse, (a) a zero defect policy is optimal if c , the constant periodic input cost, is greater than or equal to c^f where c^f satisfies

$$c^f = w'(1) - w(1). \quad (1.1)$$

(b) if an interior solution obtains, the optimal contract is a stationary policy.

(i) Without commitment, it is described by:

$$w'(a^*) = \frac{c + w(a^*)}{a^*} \quad (1.2)$$

where $\{a^*, w^*\}$ is the optimal periodic action-wage pair and both variables are independent of t , and $w(a)$ is a composite function described below.

(ii) With commitment, the optimal pair $\{a^*, w^*\}$ are described by:

$$\lambda = \frac{1}{U'(w^*)} \quad (1.3)$$

$$\lambda = \frac{\left(\frac{w^* + c}{a^*} \right)}{V'(a^*)} \quad (1.4)$$

Proof of Theorem 1.b.i: To prove the case without commitment, I use stochastic dynamic programming, and note that the principal's value function, $\phi(0)$, is defined as the value of the optimal policy in the failure state of $q = 0$.¹

¹The game ends when the objective, q^* , is realized; so $\phi(q^*) = 0$. (This problem is a discrete version of a “first crossing time problem”; so, for any constant action, $a > 0$, the game ends with probability one.)

Appendix 2: Theorems & Proofs

$$\phi(0) = \max_a \left\{ ar - w(a) - c + (1-a)\phi(0) \right\}, \quad (1.5)$$

or, since r will be received with probability one, we can rewrite this as:

$$\phi(0) = r - \min_a \left\{ \frac{[w(a) + c]}{a} \right\}. \quad (1.6)$$

differentiating with respect to a gives,

$$w'(a^*) = \frac{c + w(a^*)}{a^*}. \quad (1.2)$$

Proof of Theorem 1.b.ii: With commitment, the principal's optimization problem can be written as:

$$\max_{\{a_i, w_i\}_{i=1}^I} \left\{ a_i r - w_i - c + (1-a_i)\Psi_i \right\} \quad (1.7)$$

subject to:

Participation:

$$U(w_i^*) - V(a_i^*) + (1-a_i^*)\Gamma_i \geq 0 \quad (1.8)$$

Where the agent's and principal's conditional expected utilities are defined, respectively, as:

$$\Gamma_{i-1} = U(w_i^*) - V(a_i^*) + (1-a_i^*)\Gamma_i \quad (1.9)$$

$$\Psi_{i-1} = a_i r - w_i - c + (1-a_i)\Psi_i \quad (1.10)$$

Appendix 2: Theorems & Proofs

The objective function is (weakly) concave with respect to the choice variables and the participation constraint is strictly concave (so we are maximizing over a convex set). The first-order conditions with respect to the periodic wage, w_t ,

$$\lambda = \frac{1}{U'(w_t^*)} \quad \text{for } \forall t \quad (1.11)$$

are identical; so, the wage is constant over time. Thus, all that is left to show is that the action is constant over time. Since the wage is constant, I rewrite the maximization problem as:

$$\max_{\{a_t, w\}_{t=1}} \{a_t r - w - c + (1 - a_t) \Psi_t\} \quad (1.12)$$

subject to:

Participation:

$$U(w) - V(a_t^*) + (1 - a_t^*) \Gamma_t \geq 0. \quad (1.13)$$

Now, assume that a constant action over time is optimal, then the first-order conditions with respect to each a_t are

$$(1 - a^*)^{t-1} [r - \Psi_t] - \lambda [V'(a_t) + \Gamma_t] = 0, \quad (1.14)$$

and the matrix of cross-partials is a negative, diagonal matrix where the diagonal terms are

$$\frac{\partial^2 L}{\partial w^2} = \lambda U''(w) < 0 \quad \text{and} \quad \frac{\partial^2 L}{\partial a_t^2} = -\lambda V''(a_t) < 0 \quad \text{for } \forall t, \quad (1.15)$$

Appendix 2: Theorems & Proofs

and the off-diagonal terms equal zero:

$$L_{wa_t} = L_{a_t a_t} = 0 \text{ for } t \neq s. \quad (1.16)$$

Consider the naturally-ordered principal minors of such a diagonal matrix where each term is negative. The sign of the derivative of the k th-order principal minor is negative when k is odd and positive when k is even; this is exactly the condition which describes a local maximum, and since we are maximizing a weakly concave function over a convex set, it is a global maximum.

Proof of Theorem 1.a: first lemma 1.1 is proven.

Lemma 1.1: as the level of (future) production costs increases, the optimal action increases, i.e., $a'(c) > 0$.

Proof of Lemma 1.1: Since condition 1.2 must hold for every $c > 0$,

$$w'(a^*(c)) - \frac{c + w(a^*(c))}{a^*(c)} = 0, \quad (1.1.1)$$

differentiation with respect to c yields:

$$a'(c) \left[w''(a(c)) - \frac{1}{a(c)} \underbrace{\left(w'(a(c)) - \frac{c + w(a(c))}{a(c)} \right)}_{= 0 \text{ by first-order condition}} \right] - \frac{a(c)}{(a(c))^2} = 0. \quad (1.1.2)$$

Appendix 2: Theorems & Proofs

When evaluated at the optimal a^* , the term in the large parentheses in the brackets disappears; so, rearranging gives:

$$a'(c) = \frac{1}{a(c)w''(a(c))} > 0. \quad (1.1.3)$$

This is positive since $a(c) > 0$ and $w''(a) > 0$.

Now, rearranging the optimality condition (1.2), which holds for any $c < c^f$, gives

$$c = a^* w'(a^*) - w(a^*). \quad (1.17)$$

Since $w(a)$ is strictly convex, the RHS is strictly increasing in a ; so, for any a^* we can find a level of cost, c , which satisfies (1.17). Thus, there exists a c^f such that

$$c^f = w'(1) - w(1), \quad (1.18)$$

and for every $c \geq c^f$, a ZDP is optimal.

Theorem 3: (Restated for the case where $U(0) \geq V(0) = 0$.) If all of the assumptions stated in the text hold and if $U(0) \geq 0$ and $V(0) = 0$, then there exists T^* such that $a_{T^*} = 1$.

Proof of Theorem 3: The proof is given through a trio of lemmas.²

²The proof given here is more general than the one given in previous versions of this paper which required the assumption that once the project was started, the principal was obligated to complete it, i.e., the principal was desperate.

Appendix 2: Theorems & Proofs

Lemma 3.1: The optimal contract will have a finite length.

Proof of Lemma 3.1: Recall the definition of Γ_{t-1} (equation (12) in the text). Rearranging this definition gives

$$\Gamma_{t-1} = U(b_t^*) + \Gamma_t + a_t^* (U(g_t^*) - U(b_t^*) - \Gamma_t) - V(a_t^*). \quad (3.1.1)$$

Using period t 's incentive compatibility constraint, we can substitute for the bracketed terms in (3.1.1) to yield

$$\Gamma_{t-1} - \Gamma_t = U(b_t^*) + a_t^* V'(a_t^*) - V(a_t^*) \quad (3.1.2)$$

Since $V(\cdot)$ is convex and $V(0) = 0$, $a_t V'(a_t) - V(a_t) \geq 0$ for any nonnegative a_t . This fact, combined with the assumption that $U(b_t) \geq 0$, means that the left-hand side of (3.1.2) is nonnegative for every t . This means that the sequence of the agent's (conditional) expected utilities, $\{\Gamma_t\}$, is nonincreasing.

The sequence, $\{\Gamma_t\}$, is bounded below by $-V(1)$; if there existed a τ such that $\Gamma_\tau < -V(1)$, the agent could improve his utility by selecting $a_\tau = 1$ and get utility of at least $-V(1)$ since $U(\cdot) \geq 0$. Since it is nonincreasing and bounded below, it must converge to some point, $\Gamma_\infty \geq -V(1)$. Thus, for every $\varepsilon > 0$, there exists an $T(\varepsilon)$ such that for $t - 1 \geq T(\varepsilon)$,

$$\varepsilon > \Gamma_{t-1} - \Gamma_\infty > \Gamma_{t-1} - \Gamma_t = U(b_t^*) + a_t^* V'(a_t^*) - V(a_t^*) \geq 0. \quad (3.1.3)$$

Appendix 2: Theorems & Proofs

Now $[a_t V'(a_t) - V(a_t)]$ is increasing in a , and both that term and $U(b_t)$ are nonnegative by hypothesis, as $\varepsilon \rightarrow 0$, $[U(b_t) + a_t V'(a_t) - V(a_t)] \rightarrow 0$. Thus, both $U(b_t) \rightarrow 0$ and $a_t \rightarrow 0$ as $t \rightarrow \infty$.

Now consider the principal's expected profits (in non-recursive form) after period t :

$$\Psi_t = \sum_{k=t+1}^{\infty} \left(\prod_{j=t+1}^{k-1} (1 - a_j) \right) [a_k (r - g_k) - (1 - a_k)(b_k - c)] \quad (3.1.4)$$

For any positive r , and nonnegative sequences $\{g_k\}$, $\{b_k\}$, the above term is certainly less than the following term:

$$r - \sum_{k=t+1}^{\infty} \left(\prod_{j=t+1}^{k-1} (1 - a_j) \right) [(1 - a_k)c] \quad (3.1.5)$$

(This is true because (3.1.5) provides the benefit with probability one and requires no additional wage costs after period t .)

From the following steps, it is easy to see that the expression in (3.1.5) is less than zero in some period t where a_t is sufficiently small such that $c/a_t > r$.

$$r - \sum_{k=t+1}^{\infty} \left(\prod_{j=t+1}^{k-1} (1 - a_j) \right) [(1 - a_k)c] = r - c \sum_{k=t+1}^{\infty} \left(\prod_{j=t+1}^k (1 - a_j) \right) < r - c \sum_{k=t+1}^{\infty} (1 - a_j)^{k-t} \quad (3.1.6)$$

(The inequality arises because the sequence of actions is decreasing.) Finally, we see that the negative expression on the RHS of the statement is a geometric series; so,

Appendix 2: Theorems & Proofs

$$r - c \sum_{k=1}^{\infty} (1 - a_1)^{k-1} = r - \frac{c}{(1 - (1 - a_1))} = r - \frac{c}{a_1} \quad (3.1.7)$$

Thus, if the contract were of an infinite length, the principal would always find it in his best interests to terminate the contract. Therefore, the contract cannot be of an infinite length.

Lemma 3.2: If a finite-period contract does not ensure success, then under the assumptions given above the participation constraint does not bind— $\Gamma_0 > 0$.

Proof of Lemma 3.2: Note that after the last period of a contract with a finite maximum length, the agent's expected utility is zero. That is, if the contract lasts a maximum of T periods, then $\Gamma_T = 0$.

If a success is not required by the last period of the contract, then Γ_{T-1} , the agent's expected utility prior to the start of the last period, must be nonnegative:

$$\Gamma_{T-1} = a_T^* U(g_T^*) + (1 - a_T^*) (U(b_T^*) + 0) - V(a_T^*) \geq 0 \quad (3.2.1)$$

If not, the optimal last period action, a_T , would be zero thereby ensuring that $\Gamma_{T-1} \geq 0$. Now, assume that the participation binds (so that $\Gamma_0 = 0$). Since $\{\Gamma_t\}$ is a nonincreasing sequence (see the beginning of the proof of Lemma 3.1), then it must be that $\Gamma_{T-1} = 0$. Thus, since $\Gamma_0 = \Gamma_{T-1} = \Gamma_T = 0$, the agent's expected utility, Γ_t , must be equal to zero at every period $t = 0, 1, \dots, T$.

This means that periodic failure wages, b_t , must be equal to zero each period. To see this note that $\Gamma_{T-1} = 0$, and note that if $b_T > 0$, then the agent could ensure strictly positive

Appendix 2: Theorems & Proofs

utility in the last period by taking no action, i.e., $U(b_T) > 0$. By backward induction, we get that $b_t = 0$ for every t .

Since $\Gamma_t = b_t = 0$ for every t , the incentive compatibility constraints and the participation constraint can be written as:³

$$U(g_t) - V'(a_t) = 0 \quad \forall t \quad (3.2.2)$$

$$a_t U(g_t) - V(a_t) = 0 \quad (3.2.3)$$

Similarly, the facts imply that:

$$a_t U(g_t) - V(a_t) = 0 \quad \forall t \quad (3.2.4)$$

Now combining the incentive compatibility conditions (3.2.2) with the equations in (3.2.4) gives that:

$$U(g_t) = \frac{V(a_t)}{a_t} = V'(a_t) \quad \forall t \quad (3.2.5)$$

This means that $V(a_t) = a_t V'(a_t)$, but this is a contradiction since $V(\cdot)$ is strictly convex and $V(0) = 0$ by assumption. Thus, there is no solution which causes the participation constraint to bind for which $a_T < 1$ is optimal.

Lemma 3.3: The optimal finite-period contract will ensure success in the last period. (If the contract lasts a maximum of T^* periods, where $T^* < \infty$, then $a_{T^*} = 1$.)

³Recall we are assuming that the participation constraint binds.

Appendix 2: Theorems & Proofs

Proof of Lemma 3.3: Suppose not. Suppose that a T -period contract is optimal and that $0 < a_T < 1$. (If $a_T = 0$, then a $(T-1)$ period contract would be optimal.) Since the T th period incentive compatibility constraint must bind, it must be that $\mu_T > 0$. Similarly, for each period $t = 0, 1, \dots, T-1$, $\mu_t > 0$.

Now consider the first-order conditions with respect to b_t :

$$\frac{1}{U'(b_t^*)} = \lambda - \sum_{j=1}^{t-1} \left[\frac{\mu_j}{\prod_{m=1}^j (1 - a_m^*)} \right] - \frac{\mu_t}{\prod_{n=1}^t (1 - a_n^*)} \quad (3.3.1)$$

Note that since $U'(\cdot) > 0$, the right-hand side is always positive (for every t). Now, since $\Gamma_0 > 0$, then $\lambda = 0$, and since $\mu_t > 0$ for every $t = 0, 1, \dots, T$, the left-hand side is negative for every t . Thus, we have a contradiction that $a_T < 1$. Thus, it must be that $a_T = 1$, and so our proof is complete.

Corollary 3.1: If a finite-length contract is optimal when the agent's utility function (for wealth) is unbounded below, then $a_{T^*} = 1$.

Proof of Corollary 3.1: Suppose not. Suppose that the optimal contract has a finite maximum length of T periods and that $a_T < 1$. Let Π_T equal the ex ante probability of success (over the length) of the contract. So,

$$\Pi_T = \sum_{t=1}^T \left[\left(\prod_{j=1}^{t-1} (1 - a_j) \right) a_t \right], \quad (3.C1.1)$$

Appendix 2: Theorems & Proofs

and $(1 - \Pi_T)$ is the probability of failure through period T .

Since the agent's utility is unbounded below in wealth, the participation constraint will bind in the optimal contract; so, under the assumption that the agent's reservation utility is zero, his ex ante expected utility is zero. Similarly, if a failure occurs in the last period, the agent's conditional expected utility is zero.

Now if it is worthwhile for the principal to undertake the project, it must be that his expected profits, Ψ_0 , are greater than zero. By appending the optimal contract with another optimal contract, the agent is no worse off, his incentives are unchanged, and the principal is strictly better off (by $(1 - \Pi_T)\Psi_0$) which contradicts the hypothesis that the optimal contract has a finite maximum length and does not require success in the last period. Therefore, the corollary is proved.

Theorem 4: Failure wages, b_t , are a decreasing sequence in time (and $g_t > b_{t-1}$).

Proof of Theorem 4: consider the first-order conditions with respect to b_t and g_t for every t :

$$\frac{1}{U'(b_t^*)} = \lambda - \sum_{j=1}^{t-1} \left[\frac{\mu_j}{\prod_{m=1}^j (1 - a_m^*)} \right] - \frac{\mu_t}{\prod_{n=1}^t (1 - a_n^*)} \quad (4.1)$$

$$\frac{1}{U'(g_t^*)} = \lambda - \sum_{j=1}^{t-1} \left[\frac{\mu_j}{\prod_{m=1}^j (1 - a_m^*)} \right] + \frac{\mu_t}{a_t^* \prod_{n=1}^{t-1} (1 - a_n^*)} \quad (4.2)$$

Appendix 2: Theorems & Proofs

Now, consider the sequence of conditions in (4.1). Since an additional nonnegative number is subtracted each period, it is easy to see that this sequence of inverse marginal utilities for failure wages is decreasing, i.e.,

$$\left\{ \frac{1}{U'(b_t^*)} \right\}_{t=1}^{\infty} \quad (4.3)$$

is decreasing; thus, failure wages are decreasing over time.

Corollary 4.1: If $\frac{1}{U'(\cdot)}$ is linear or convex, then expected periodic wages are decreasing.

Proof of Corollary 4.1: The following condition (4.4) is a special case of Proposition 1 in Rogerson (1985). It is derived by multiplying the right-hand-sides of the inverse marginal utility conditions (4.1 and 4.2) by their respective probabilities. However, unlike in Rogerson, where this condition holds along any path, here, it holds only along the failure path (since the game ends upon any success).

$$\frac{1}{U'(b_{t-1}^*)} = a_t^* \frac{1}{U'(g_t^*)} + (1 - a_t^*) \frac{1}{U'(b_t^*)} \quad (4.4)$$

If $1/U'(\cdot)$ is linear or strictly convex, then

$$\frac{1}{U'(b_{t-1}^*)} = a_t^* \frac{1}{U'(g_t^*)} + (1 - a_t^*) \frac{1}{U'(b_t^*)} > \frac{1}{U'(a_t^* g_t^* + (1 - a_t^*) b_t^*)} \quad (4.5)$$

and

$$a_t^* g_t^* + (1 - a_t^*) b_t^* > b_t^* \geq a_{t+1}^* g_{t+1}^* + (1 - a_{t+1}^*) b_{t+1}^*. \quad (4.6)$$

Appendix 2: Theorems & Proofs

The LHS and the RHS of (4.6) give that expected periodic wages are decreasing over time.

Theorem 5: Assume $V(a) = a$, then the optimal contract has the following characteristics:

1. $b_1 = b_2 = \dots = b_{T^*-1} = g_{T^*-1} = \lambda$.
2. $a_1 = a_2 = \dots = a_{T^*-1} = 0$, and $a_{T^*} = 1$, and so
3. T^* solves the following program;

$$\min_t \{t[c + w(t)]\} \quad (5.1)$$

subject to:

$$tU(w(t)) - 1 \geq 0. \quad (5.2)$$

where $w(t)$ is the periodic wage required to ensure the agent's participation when no effort is taken until period t and then the maximum effort is required.

Proof of Theorem 5: With linear disutility of effort, the constraints to the principal's profit maximization problem (found on page 14 in the text) can be rewritten as:

$$a_i^* U(g_i^*) + (1 - a_i^*) (U(b_i^*) + \Gamma_i) - a_i^* \geq 0 \quad (5.3)$$

$$U(g_i^*) - U(b_i^*) - \Gamma_i - 1 = 0 \quad (5.4)$$

and the agent's conditional expected utility after a failure has occurred in period $t-1$, can be rewritten as:

$$\Gamma_{t-1} = a_i^* U(g_i^*) + (1 - a_i^*) (U(b_i^*) + \Gamma_i) - a_i^* \quad (5.5)$$

Appendix 2: Theorems & Proofs

Now substituting (5.2) into (5.3) yields:

$$\Gamma_{t-1} = U(b_t^*) + \Gamma_t. \quad (5.6)$$

If the incentive compatibility constraint (5.4) binds, then the optimal action is zero which is a contradiction; thus, the constraint cannot bind for any period t ; so, either periodic effort is equal to zero or to one. Also, since the constraint does not bind, the lagrangian, μ_t , is equal to zero for every t . From conditions (4.1) and (4.2) in the proof of Theorem 4, the periodic failure and success wages must be equal; so the principal pays a fixed wage each period. Since $a_t = 0$ or 1 for every t , and since it is suboptimal for the principal to pay a wage each period if a success will never occur, then there must exist a period T^* in which the agent takes $a_{T^*} = 1$. So, the principal must find the number t for which $a_1 = a_2 = \dots = a_{t-1} = 0$, and $a_t = 1$, which minimizes expected costs while ensuring that the agent participate in the contract. In other words, the principal solves the following program:

$$\min_t \{t[c + w(t)]\}$$

subject to:

$$tU(w(t)) - 1 \geq 0.$$

Theorem 6: Let $a_{T^*} = 1$, then $b_{T^*-1} = g_{T^*}$. (b_{T^*-1} is the failure wage paid in period $T^* - 1$ and g_{T^*} is the wage paid in the final period, T^*).

Proof of Theorem 6: Suppose not. W.L.O.G. assume $b_{T^*-1} > g_{T^*}$. Since $\dot{U}(\cdot)$ is strictly concave, we know from Jensen's Inequality that:

Appendix 2: Theorems & Proofs

$$\left[\frac{1}{2} U(g_T) + \frac{1}{2} U(b_{T-1}) \right] < \left[U\left(\frac{1}{2} g_T + \frac{1}{2} b_{T-1} \right) \right]. \quad (6.1)$$

Thus, the principal could provide the same utility as $[U(b_{T^*-1}) + U(g_{T^*})]$ at a cost less than $(b_{T^*-1} + g_{T^*})$ by offering a constant wage each period. Since total utility after a T^*-1 period failure does not change, the agent's incentives do not change; so, setting $b_{T^*-1} = g_{T^*}$ does not change the agent's optimal actions. Therefore, it is not optimal for $b_{T^*-1} \neq g_{T^*}$. Thus, $b_{T^*-1} = g_{T^*}$.

Theorem 7: As the agent's efficiency parameter increases within a period, the optimal action increases within the period, i.e., $a_t'(p_t) > 0$.

Proof of Theorem 7a (observable action case): the principal's problem is:

$$\min_{\{a_t\}_{t=1}^T} \left\{ w(a_1) + c + \sum_{i=2}^{\infty} \left[\left(\prod_{j=1}^{i-1} (1 - p_j a_j) \right) (w(a_i) + c) \right] \right\}, \quad (7.1a)$$

where each period's participation constraint has been substituted for as in Theorem 1. The optimality condition for the principal's problem in learning-period t ($t = 1, \dots, T$) is:

$$w'(a_t) = p_t \left[w(a_{t+1}) + c + \sum_{j=t+2}^{\infty} \left[\left(\prod_{m=t+1}^{j-1} (1 - p_m a_m) \right) (w(a_j) + c) \right] \right]. \quad (7.2a)$$

Implicit differentiation of the above condition with respect to p_t and simple rearrangement yields (7.3a):

Appendix 2: Theorems & Proofs

$$\frac{\partial a_t}{\partial p_t} = \frac{\left[w(a_{t+1}) + c + \sum_{j=t+2}^T \left[\left(\prod_{j=1}^{j-1} (1 - p_m a_m) \right) (w(a_j) + c) \right] + \left(\prod_{m=t+1}^T (1 - p_m a_m) \right) \left(\frac{w(a_{T+1}) + c}{a_{T+1}} \right) \right]}{w''(a_t)} > 0$$

The numerator, which is the expected cost after period t , and the denominator are both positive; thus, $a_t'(p_t) > 0$.

Proof of Theorem 7b (unobservable action case): With an uncommitted agent, we can backwardly solve each period's contract, and in any period, failure costs are exogenous. Fix all parameters after period $t + 1$, and let Ψ_t represent expected costs after period t (from period $t + 1$ onward). In period t , the principal must solve the following problem:

$$\min_{\{a_t, g_t, b_t\}} \{p_t a_t g_t + (1 - p_t a_t) [b_t + \Psi_t]\} \quad (7.1b)$$

subject to:

$$p_t a_t^* U(g_t^*) + (1 - p_t a_t^*) [U(b_t^*) + \Gamma_t] - V(a_t^*) \geq 0 \quad (7.2b)$$

$$p_t [U(g_t^*) - U(b_t^*) - \Gamma_t] - V'(a_t^*) = 0. \quad (7.3b)$$

There are five endogenous variables in this single-period problem: g_t , b_t , a_t , λ_t , and μ_t .

The five corresponding first-order conditions are

$$L_{g_t} = L_{b_t} = L_{a_t} = L_{\lambda_t} = L_{\mu_t} = 0 \quad (7.4b)$$

where L is the Lagrangian associated with (7.1b - 7.3b). Differentiating this system of equations with respect to p_t and rearranging, generates the following system of equations:

Appendix 2: Theorems & Proofs

$$\begin{bmatrix} L_{gg} & L_{gb} & L_{ga} & L_{g\lambda} & L_{g\mu} \\ L_{gb} & L_{bb} & L_{ba} & L_{b\lambda} & L_{b\mu} \\ L_{ga} & L_{ba} & L_{aa} & L_{a\lambda} & L_{a\mu} \\ L_{g\lambda} & L_{b\lambda} & L_{a\lambda} & L_{\lambda\lambda} & L_{\lambda\mu} \\ L_{g\mu} & L_{b\mu} & L_{a\mu} & L_{\lambda\mu} & L_{\mu\mu} \end{bmatrix} \begin{bmatrix} g'(p_t) \\ b'(p_t) \\ a'(p_t) \\ \lambda'(p_t) \\ \mu'(p_t) \end{bmatrix} = - \begin{bmatrix} L_{gp_t} \\ L_{bp_t} \\ L_{ap_t} \\ L_{\lambda p_t} \\ L_{\mu p_t} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \Psi_t \\ 0 \\ 0 \end{bmatrix} \quad (7.5b)$$

Now, we can solve for $a_t'(p_t)$ by applying Cramer's rule (by substituting the vector on the RHS into the third column of the hessian matrix). Such an application gives

$$a_t'(p_t) = \frac{\Psi_t |D_3|}{|D|}. \quad (7.6b)$$

With two constraints, for a minimum to attain, the hessian and all of its border-preserving principle minors must have positive determinants. So, let D represent the hessian and D_3 represent the resulting 4-by-4 matrix after the third row and third column have been eliminated. Since both matrices are border-preserving principle minors, their determinants are both positive and failure costs are positive, $a_t'(p_t) > 0$.

Corollary 7.1: the optimal action and wage from the $(T+1)th$ period onward—after learning has stopped—is greater than the optimal action and wage in the Tth period.

Proof of Corollary 7.1: (proof is given in the first-best case because it better illustrates the issue; however, in either case it is easy to see that because of Theorem 7, when $p_T < 1$, the optimal action is less than when $p_T = 1$.) Comparing the optimality condition for period $T+1$ (Condition (1)) with the optimality condition for period T , (7.2a), yields:

$$w'(a_T) = p_T \left[\frac{[w(a_{T+1}) + c]}{a_{T+1}} \right] = p_T w'(a_{T+1}). \quad (7.4a)$$

Since $w(\cdot)$ is strictly convex and $p_T < 1$, we have that $a_T < a_{T+1}$.

Appendix 2: Theorems & Proofs

Theorem 8: The higher the future efficiency parameter, p_s , for $s > t$ —the lower the optimal action in period t , i.e., $a_t'(p_s) < 0$.

Proof of Theorem 8: (I show the proof in the unobservable action case, identical logic holds in the observable action case.) Fix all parameters after period $t + 1$, and let Ψ_t represent expected costs after period t (from period $t + 1$ onward). In period $t+1$, the principal must solve the following problem (which minimizes Ψ_t):

$$\min_{\{a_{t+1}, g_{t+1}, b_{t+1}\}} \{p_{t+1}a_{t+1}g_{t+1} + (1 - p_{t+1}a_{t+1})[b_{t+1} + \Psi_{t+1}]\} \quad (8.1)$$

subject to:

$$p_{t+1}a_{t+1}^* U(g_{t+1}^*) + (1 - p_{t+1}a_{t+1}^*) [U(b_{t+1}^*) + \Gamma_{t+1}] - V(a_{t+1}^*) \geq 0 \quad (8.2)$$

$$p_{t+1} [U(g_{t+1}^*) - U(b_{t+1}^*) - \Gamma_{t+1}] - V'(a_{t+1}^*) = 0. \quad (8.3)$$

Consider the first-order condition of the lagrangian with respect to action:

$$p_{t+1} [g_{t+1}^* - b_{t+1}^* - \Psi_{t+1}] + \mu_{t+1} V''(a_{t+1}^*) = 0. \quad (8.4)$$

Since $\mu V''(a) > 0$, the term in brackets is negative.

Now apply the envelope theorem by differentiating the lagrangian with respect to p_{t+1} :

$$a_{t+1}^* [g_{t+1}^* - b_{t+1}^* - \Psi_{t+1}] - \lambda_{t+1} [U(g_{t+1}^*) - U(b_{t+1}^*) - \Gamma_{t+1}] - \mu_{t+1} [U(g_{t+1}^*) - U(b_{t+1}^*) - \Gamma_{t+1}] \quad (8.5)$$

From the incentive compatibility constraint, (8.3), the sums inside the second and third brackets of expression (8.5) are positive. Subtracting both sums yields a negative value.

Appendix 2: Theorems & Proofs

Similarly, by the first-order condition with respect to effort, (8.4), the term in the first bracket is negative. So, $\Psi_t'(p_{t+1}) < 0$ or increases in next period's efficiency parameter reduce this period's expected failure costs.

Now, the result follows from lemma 8.1 (the proof is analogous to the proof of Thm. 7 (part b):

Lemma 8.1: $a_t'(\Psi_t) > 0$ in the interior of $[0, 1]$.

proof: because we can solve the sequence of contracts through backward induction, in the cost minimization in period t , there are five endogenous variables: g_t , b_t , a_t , λ_t , and μ_t associated with the problem (8.1 - 8.3). So, there are five first-order conditions as in (7.4b).

Differentiating this system of equations with respect to Ψ_t generates:

$$\begin{bmatrix} L_{gg} & L_{gb} & L_{ga} & L_{g\lambda} & L_{g\mu} \\ L_{gb} & L_{bb} & L_{ba} & L_{b\lambda} & L_{b\mu} \\ L_{ga} & L_{ba} & L_{aa} & L_{a\lambda} & L_{a\mu} \\ L_{g\lambda} & L_{b\lambda} & L_{a\lambda} & L_{\lambda\lambda} & L_{\lambda\mu} \\ L_{g\mu} & L_{b\mu} & L_{a\mu} & L_{\lambda\mu} & L_{\mu\mu} \end{bmatrix} \begin{bmatrix} g'(\Psi_{t+1}) \\ b'(\Psi_{t+1}) \\ a'(\Psi_{t+1}) \\ \lambda'(\Psi_{t+1}) \\ \mu'(\Psi_{t+1}) \end{bmatrix} = - \begin{bmatrix} L_{g\Psi_{t+1}} \\ L_{b\Psi_{t+1}} \\ L_{a\Psi_{t+1}} \\ L_{\lambda\Psi_{t+1}} \\ L_{\mu\Psi_{t+1}} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ p_t \\ 0 \\ 0 \end{bmatrix} \quad (8.6)$$

Now, as in the previous proof, we can solve for $a_t'(\Psi_t)$ by applying Cramer's rule to get

$$a_t'(\Psi_t) = \frac{p_t |D_3|}{|D|}. \quad (8.7)$$

For a minimum to attain, the hessian and all of its border-preserving principle minors must have positive determinants; so, $a_t'(\Psi_t) > 0$.

Appendix 2: Theorems & Proofs

Now since $a_t'(\Psi_t) > 0$ and $\Psi_t'(p_{t+1}) < 0$, we have that $a_t'(p_{t+1}) < 0$ and by induction, $a_t'(p_s) < 0$ for $s > t$.

Appendix 3: Examples of Optimal Contracts

Appendix 3.1: Assume that the principal's periodic cost, c , is 5, and the agent has logarithmic utility and exponential disutility of the form $V(a) = ae^{1/5a}$. Then if effort is observable, an infinite sequence of contracts, which require the agent to take an action = .56 and pays the agent a wage of 3.64 is optimal.

First-Best Contract

action	wage	disutility	utility	expected costs
.56	3.64	1.29	1.29	15.42


Under the same assumptions, if the agent's effort is unobservable, it is optimal for the principal to impose a deadline at the end of four periods. Thus in the second-best contract, we observe that actions are increasing over time and wages are decreasing. Because the agent commits to complete the project and the principal imposes a deadline, the agency costs are limited to 3.02. This table corresponds to Graph 6.A in the text.

Second-Best Contract

period:	actions:	good wage	bad wage	disutilities	good utils	bad utils	Agent's EU:	Princ's EC:	Cum Net Utility	Change in Util	ex ante pd.t utility	Periodic expected utility
ex ante							0.00	18.44				
1	0.32	8.62	3.24	0.52	2.15	1.17	-1.43	12.49	1.63	0.53		0.97
2	0.51	4.56	1.87	1.09	1.52	0.63	-2.88	8.64	1.08	0.37	-0.97	-0.01
3	0.69	2.35	0.80	1.97	0.85	-0.23	-4.71	5.80	-0.92	-0.21	-0.96	-1.45
4	1.00	0.80	0.00	4.48	-0.23				-6.71	-0.68	-0.48	-4.71

Appendix 3: Examples of Optimal Contracts

The following graph shows the wage-cost advantage of imposing the deadline at four periods rather than inducing a zero-defect policy which imposes a deadline in the first period. A ZDP minimizes failure-related input costs, but it does this at an extremely high wage cost; thus, the principal is willing to trade-off input costs—in expectation—for wage costs. With an input cost of 5, the principal is willing to fail (up to) three times to reduce expected wage costs over the length of the contract.



Appendix 3: Examples of Optimal Contracts

With unobservable effort and no commitment by the agent, the optimal contract is again an infinite sequence of identical contracts. Here, the induced level of effort is .37 per period instead of .56 per period in the first-best case and expected wages are 4.3408 rather than 3.64 as in the first-best. This means that the agency costs are 8.78 of the total expected costs of 25.2.

Second-Best Contract without Commitment

action	success wage	failure wage	disutility	success utility	failure utility	expected costs
.37	10.54	.70	.65	2.35	-.36	25.2

Without commitment, a ZDP is optimal at costs above 135.2. As Graph 4 shows, with commitment ZDPs are optimal at costs above 224.5.

Appendix 3: Examples of Optimal Contracts

Appendix 3.2: Another example of an optimal second-best contract with commitment. Assume that the input cost is 6.5 per period and the agent has square root utility and disutility function of $4.5a^2$ where a is the level of effort. In this case, the principal imposes the deadline at three periods:

period:	actions:	good wage	bad wage	disutilities	good utils	bad utils	Agent's EU:	Princ's EC:
ex ante							0.00	18.75
1	0.27	5.85	1.46	0.32	2.42	1.21	-1.65	13.10
2	0.33	2.95	0.93	0.48	1.72	0.96	-3.54	7.43
3	1.00	0.93	0.00	4.50	0.96	0.00	0.00	0.00

Appendix 4: An Example of a Costly Learning Problem

There are many ways to introduce costly learning into the type of problem studied in this paper; the following presentation is probably the simplest: I consider a two-iteration problem where the principal can observe the product's quality after each iteration.¹ I assume that if a failure occurred in the first period, then with probability ℓ , the agent is presented with the opportunity to improve his second period's efficiency—the probability of success for a fixed second period action—by taking an additional amount of effort, L . The opportunity to learn cannot be communicated to the principal (nor can the principal infer when the learning opportunity was accepted). Such a setting can arise if the agent can learn, by say, investigating the cause of failure. When more than one unit is produced per period, such an investigation could, in fact, be costly to the principal if the investigation resulted in a lower first-period production rate.

If the learning action does not affect reported results for the period, the problem can be solved using point-wise optimization. However, if learning reduces current-period results, by say requiring the destructive testing of certain units or a shut-down of production, then the agent's learning decision requires him to consider his utility at two points along the optimal contract curve: at the realized level and at the lower level which results from the learning action. In this case, the problem does not degenerate to pointwise optimization: one must use optimal control with "time" lags.

The agent's decision to learn, denoted d , is endogenous: of course, the only interesting case is when the principal prefers that the agent learn: when $d = 1$. With this type of costly learning, the principal's problem becomes:

¹In this general setting, a finite-period problem arises if the principal's inventory of inputs is finite.

Appendix 4: An Example of a Costly Learning Problem

$$\max_{\{g_1\}, \{b_1\}, \{a_1\}} \left\{ pa_1(r - g_1) - c - (1 - pa_1)(b_1 + c) + \right. \\ \left. (1 - pa_1) \left[\ell \left(d[a_2(L)(r - g_2) - (1 - a_2(L))b_2] + (1 - d)[pa_2(N)(r - g_2) - (1 - pa_2(N))b_2] \right) + \right. \right. \\ \left. \left. (1 - \ell) \left[pa_2(N)(r - g_2) - (1 - pa_2(N))b_2 \right] \right] \right\}$$

subject to:

IR:

$$pa_1^* U(g_1^*) - V(a_1^*) + \\ (1 - pa_1^*) \left[U(b_1^*) + \ell \left(d^* \left[-V(L) + a_2^*(L)U(g_2^*) + (1 - a_2^*(L))U(b_2^*) - V(a_2^*(L)) \right] + \right. \right. \\ \left. \left. (1 - d^*) \left[pa_2^*(N)U(g_2^*) + (1 - pa_2^*(N))U(b_2^*) - V(a_2^*(N)) \right] \right) + \right. \\ \left. (1 - \ell) \left(pa_2^*(N)U(g_2^*) + (1 - pa_2^*(N))U(b_2^*) - V(a_2^*(N)) \right) \right] \geq 0$$

Learning IC:

$$d^* \in \arg \max_{d \in [0, 1]} \left\{ d \left[-V(L) + \left[a_2^*(L)U(g_2^*) + (1 - a_2^*(L))U(b_2^*) - V(a_2^*(L)) \right] \right] + \right. \\ \left. (1 - d) \left[pa_2^*(N)U(g_2^*) + (1 - pa_2^*(N))U(b_2^*) - V(a_2^*(N)) \right] \right\}$$

it is easy to see that $d^* \in \{0, 1\}$; so, this condition can be rewritten as:

$$(2d^* - 1) \left[-V(L) + \left[a_2^*(L)U(g_2^*) + (1 - a_2^*(L))U(b_2^*) - V(a_2^*(L)) \right] \right] \\ (1 - 2d^*) \left[pa_2^*(N)U(g_2^*) + (1 - pa_2^*(N))U(b_2^*) - V(a_2^*(N)) \right] \geq 0$$

1. $a_1^* \in \arg \max_{a_1} \left\{ pa_1 U(g_1^*) - V(a_1) + (1 - pa_1) \left[U(b_1^*) + EU(\text{period 2}) \right] \right\}$
- 2L. $a_2^*(L) \in \arg \max_{a_2} \left\{ a_2 U(g_2^*) + (1 - a_2)U(b_2^*) - V(a_2) \right\}$
- 2N. $a_2^*(N) \in \arg \max_{a_2} \left\{ pa_2 U(g_2^*) + (1 - pa_2)U(b_2^*) - V(a_2) \right\}$

I replace the IC conditions by their associated first-order conditions:

Appendix 4: An Example of a Costly Learning Problem

$$1. \quad pU(g_1^*) - V'(a_1^*) - p[U(b_1^*) + EU(\text{period 2})] = 0$$

$$2L. \quad U(g_2^*) - U(b_2^*) - V'(a_2^*(L)) = 0$$

$$2N. \quad p[U(g_2^*) - U(b_2^*)] - V'(a_2^*(N)) = 0$$

Lemma 1: The second period action with learning is greater than the second period action without learning, or $a_2^*(L) > a_2^*(N)$.

Proof: From the second-period IC constraints, since (i) $V(\cdot)$ is convex (and increasing), (ii) $p < 1$, and (iii)

$$U(g_2^*) - U(b_2^*) > 0,$$

we have that:

$$U(g_2^*) - U(b_2^*) = V'(a_2^*(L)) > V'(a_2^*(N)) = p[U(g_2^*) - U(b_2^*)].$$

In other words, an "educated" agent takes a higher level of action than an uneducated one; not only does learning, by itself, increase the probability of success, but it also induces the agent to work harder.

The optimal wage schedule will satisfy the following conditions:

1st-period:

success:

$$\frac{1}{U'(g_1^*)} = \lambda + \frac{\mu_1}{a_1^*}$$

failure:

Appendix 4: An Example of a Costly Learning Problem

$$\frac{1}{U'(b_1^*)} = \lambda - \frac{p\mu_1}{(1 - pa_1^*)}$$

2nd-period

success:

$$\frac{1}{U'(g_2^*)} = \lambda - \frac{p\mu_1}{(1 - pa_1^*)} + \frac{\gamma[(2d^* - 1)a_2^*(L) + (1 - 2d^*)pa_2^*(N)] + \mu_{2L} + p\mu_{2N}}{(1 - pa_1^*)[\ell[d^*a_2^*(L) + (1 - d^*)pa_2^*(N)] + (1 - \ell)pa_2^*(N)]}$$

failure:

$$\frac{1}{U'(b_2^*)} = \lambda - \frac{p\mu_1}{(1 - pa_1^*)} + \frac{\gamma[(2d^* - 1)(1 - a_2^*(L)) + (1 - 2d^*)(1 - pa_2^*(N))] - \mu_{2L} - p\mu_{2N}}{(1 - pa_1^*)[\ell[d^*(1 - a_2^*(L)) + (1 - d^*)(1 - pa_2^*(N))] + (1 - \ell)(1 - pa_2^*(N))]}$$

if $d^* = 1$ then the second period wages become:

$$\frac{1}{U'(g_2^*)} = \lambda - \frac{p\mu_1}{(1 - pa_1^*)} + \frac{\gamma[a_2^*(L) - pa_2^*(N)] + \mu_{2L} + p\mu_{2N}}{(1 - pa_1^*)[\ell a_2^*(L) + (1 - \ell)pa_2^*(N)]}$$

$$\frac{1}{U'(b_2^*)} = \lambda - \frac{p\mu_1}{(1 - pa_1^*)} + \frac{\gamma[pa_2^*(N) - a_2^*(L)] - \mu_{2L} - p\mu_{2N}}{(1 - pa_1^*)[\ell(1 - a_2^*(L)) + (1 - \ell)(1 - pa_2^*(N))]}$$

References

1. Amershi, A. & S. Datar, "Incomplete Contracts, Production Expertise and Incentive Effects of Modern Manufacturing Practices," working paper (1991).
2. Anthony, R., J. Dearden, & V. Govindarajan, Management Control Systems, Seventh edition, Irwin (1992).
3. Arrow, K., "The Economic Implications of Learning by Doing," *Review of Economic Studies*, (1962) pp. 155 - 173.
4. Baiman, S. "Agency Research in Managerial Accounting: A Second Look," *Accounting, Organizations and Society*, (1990) pp. 341 - 371.
5. Baiman, S. and J. Demski, "Economically Optimal Performance Evaluation and Control Systems," *Journal of Accounting Research*, (1980 Supplement) p. 184-220.
6. Bowles, J., "The Riddle of German Quality," *Across the Board*, published by The Conference Board (January/February 1993) pp. 22 - 24.
7. Fine, C., "Quality Improvement and Learning in Productive Systems," *Management Science* (1986) pp. 1301 - 1315.
8. Fine, C., "A Quality Control Model Learning Effectss," *Operations Research* (1988) pp. 437 - 444.
9. Fudenberg, D., B. Holmstrom, and P. Milgrom, "Short-Term Contracts and Long-term Relationships," *Journal of Economic Theory*, (June, 1990) pp. 1 - 31.
10. Fudenberg, D., and J. Tirole, "Learning-by-Doing and Market Performance," *Bell Journal of Economics* (198x) pp. 520 - 530.
11. Hirao, Y., "Learning and Inentive Problems in Repeated Partnerships," *International Economic Review* (1993) pp. 101 - 119.

References

12. Holmstrom, B., "Moral Hazard and Observability," *Bell Journal of Economics* (1979) pp. 74 - 91.
13. Horngren, C., G. Foster and S. Datar, *Cost Accounting: A Managerial Emphasis*, Eighth edition, (1994).
14. Joskow, P., "Contract Duration and Relationship-Specific Investments: Empirical Evidence from Coal Markets," *The American Economic Review* (1987) pp. 168 - 185.
15. Kreps, D., *A Course in Microeconomic Theory*, Princeton University Press, First Edition, (1990) pp. 791 - 814.
16. Lambert, R., "Long-Term Contracting and Moral Hazard," *Bell Journal of Economics* (1983) pp. 441 - 452.
17. Manzoni, J.F. and K. Merchant, "The Achievability of Budget Targets in Profit Centers: A Field Study," *The Accounting Review*, (1989) pp. 539 - 558.
18. Radner, R., "Monitoring Cooperative Agreements in a Principal-Agent Relationship," *Econometrica*, (1981) pp. 1127 - 1148.
19. Radner, R., "Repeated Principal-Agent Games with Discounting," *Econometrica*, (1985) pp. 1173 - 1197.
20. Ramakrishnan, R., "Income Smoothing and Income Acceleration as Rational Equilibrium Behavior," Columbia University working paper, 1988.
21. Reichelstein, S., "Constructing Incentive Schemes for Government Contracts: An Application of Agency Theory," *The Accounting Review*, (1992) pp. 712 - 731.
22. Rogerson, W., "The First-Order Approach to Principal-Agent Problems," *Econometrica* (1985) pp. 1357 - 1367.

References

23. Rogerson, W., "Repeated Moral Hazard," *Econometrica*, (1981) pp. 69 - 76.
24. Rubinstein, A., and M. Yaari, "Repeated Insurance Contracts and Moral Hazard," *Journal of Economic Theory*, (1983) pp. 74 - 97.
25. Silberberg, E., *The Structure of Economics*, McGraw-Hill Book Company, (1978).
26. Spence, A., "The Learning Curve and Competition," *Bell Journal of Economics* (1981) pp. 49 - 70.
27. Varian, H., *Microeconomic Analysis*, W.W. Norton & Company, 2nd edition, (1984).
28. Whittle, P., *Optimization Over Time*, volumes I and II, John Wiley and Sons, (1982).

Copyright © 1994 by Andrew E. Spero Jr.