

Short-Term Contracts and Long-Term Agency Relationships*

DREW FUDENBERG

*Massachusetts Institute of Technology,
Cambridge, Massachusetts 02139*

BENGT HOLMSTROM

Yale University, New Haven, Connecticut 06520

AND

PAUL MILGROM

Stanford University Stanford, California 94305

Received August 24, 1988; revised June 20, 1989

Long-term contracts are valuable only if optimal contracting requires commitment to a plan today that would not otherwise be adopted tomorrow. We show that commitments are unnecessary, and hence short-term contracts are sufficient if (1) all public information can be used in contracting, (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. *Journal of Economic Literature* Classification Numbers: 022, 026.

© 1990 Academic Press, Inc.

1. INTRODUCTION

The simplest incentive schemes are probably the piece rate scheme once commonly used for laborers and the commission contract used to compensate salesmen; in both, the worker's pay is proportional to the number of

* This research was supported by funds from the National Science Foundation, the Alfred P. Sloan Foundation, and the John Simon Guggenheim Foundation. Much of the research reported here was done while Drew Fudenberg and Paul Milgrom were in residence at UC Berkeley and Bengt Holmstrom at Stanford University. We are especially grateful to John Roberts for his detailed review of an earlier version. We also benefited from comments by Oliver Hart, Michael Jensen, and participants in seminars at MIT, University of Rochester, University of Western Ontario, Stanford University, Universidad Autonoma de Barcelona, and University of Chicago.

products of various kinds produced. These contracts involve no element of deferred compensation—each period's pay depends on that period's product—and neither party pays a penalty when a quit or layoff occurs or when the contract is renegotiated. For these reasons, piece rate and commission contracts can be viewed as a series of very short-term contracts, with the worker's pay for any accounting period being the sum of his earnings over each hour or day in the period. Our purpose is to investigate when a series of short-term contracts similar to commission contracts can constitute an efficient incentive scheme, and when, to the contrary, the parties would do better to sign a long-term contract that cannot be renegotiated.¹

Long-term contracts enjoy an obvious advantage if they expand the agent's ability to smooth consumption over time, that is, if the firm acts as a banker for a worker who would otherwise have limited access to banking services. But such a rationale can hardly explain the observed variation in the length of employment contracts. Since higher income employees are likely to have easier access to capital markets, the smoothing argument would suggest that lower income employees in particular would be on long-term contracts and this appears empirically false. We are therefore led to reject the hypothesis that consumption smoothing is a major reason for long-term contracts. To focus exclusively on alternative reasons we will make the unconventional assumption that the employee can borrow and save on the same terms as the firm.²

To assess the advantages of long-term contracts we construct a general principal-agent model in which the agent can act, consume, and receive and communicate information over time. The agent is always at least as well informed as the principal. Our main finding is that the timing of the agent's information advantage is central for determining the value of long-term contracts. If at all times of potential recontracting, the principal and the agent share the same beliefs about the payoff-relevant future, then there are no gains to long-term contracts. In our model, long-term contracts only serve to avoid recontracting under asymmetric information.

More precisely, we show that the following conditions are sufficient for an efficient long-term contract to be implementable as a sequence of short-

¹ Commitments to non-renegotiable long-term contracts may be difficult to enforce. It has been suggested that long-term contracts must therefore be renegotiation-proof. We are not concerned with this issue, because our main result will provide conditions under which non-renegotiable long-term contracts are equivalent to a sequence of short-term contracts and hence, a fortiori, to renegotiation-proof long-term contracts.

² This banking assumption is necessary to model the important case where the employer cannot observe the employee's consumption or wealth. We sketch in Section 7 how to eliminate third-party banking from the model for the case where the employer can observe consumption and wealth.

term contracts: At the start of each period, (1) the preferences of the principal and the agent over future contingent outcomes (action–payment sequences) are common knowledge, (2) future technological opportunities are common knowledge, (3) payments in the period can be made contingent on all information shared by the principal and the agent, and (4) the utility frontier is downward sloping. By the utility frontier at the beginning of a period, we mean the set of expected utility pairs (given the agent's information) that can be achieved using feasible, incentive-compatible long-term contracts written at that time.

Condition 3 requires that payments can be indexed on joint information without delay. This prevents insurance opportunities from being lost in a short-term contracting mode. Condition 4 ensures that given any past history, the agent can be offered the same level of future expected utility using an efficient long-term contract as using any feasible, incentive compatible long-term contract. This condition is met if the agent's preferences are additively or multiplicatively separable.

Conditions 1 and 2 are the information restrictions. Together they rule out any form of adverse selection at the time contracts may be renegotiated. We emphasize that information asymmetries leading to adverse selection can arise either from exogenous signals that the agent receives or from unobserved actions that the agent takes. Both are problematic for contracting. It is well-known that private information of the former kind leads to inconsistencies in long-term plans, because *ex ante* incentives are efficiently provided by making *ex post* allocations inefficient (see for instance papers by Baron and Besanko [2] and Laffont and Tirole [17]). Less recognized is the fact that private information about actions will lead to similar adverse selection problems (however, see Milgrom [19]). To see why, suppose the agent has private information about past actions that affect future outcomes. As this means that future outcomes must be informative about what the agent did in the past, the sufficient statistic results of Holmstrom [13] and Shavell [27] imply that a long-term contract, which awaits the arrival of additional performance information, improves on a short-term contract. Interestingly, the intuition about allocational distortions and the intuition about sufficient statistics, while seemingly different, are really part of the same general adverse selection logic.

In view of this discussion it is not surprising that Conditions 1 and 2 are needed for our result. The surprise is that the conditions suffice. It is instructive to sketch the simple logic behind this.

The main result is derived in two steps. The first step shows that under Conditions 1–4 an efficient long-term contract is sequentially efficient: given any history, the contract is immune to renegotiations. The key insight is that Conditions 1–3 imply that the principal and the agent can

trade in a “complete” set of contingencies, namely all contingencies that are relevant for determining future expected utilities. Therefore, ex post efficiency is a necessary condition for ex ante efficiency, assuming the full range of agent incentives can be provided within the set of efficient contracts. The latter is guaranteed by Condition 4.

For the second step, we need a bargaining solution that determines which sequentially efficient agreements would be reached if new contracts had to be negotiated midstream. We adopt as a solution any contract on the efficient frontier that gives the firm a future payoff (conditional expected present value of future profits) of zero. (Such a contract can be identified by both parties when the common knowledge conditions hold.) The proof proceeds by modifying the timing of payments under the original efficient long-term contract so that the firm’s future payoff is indeed zero at every date. Since the agent has access to banking services, a mere change in the timing of the agent’s compensation affects neither his welfare nor his incentives.

Combining the two steps, any efficient contract which in the beginning gives the firm zero expected profits can be replaced by a sequentially efficient contract with a zero future payoff for the firm at each date and in each event. Since the continuations under this contract always coincide with what the parties would have agreed to anyway given our bargaining convention, there is no gain to long-term contracting.

An outline follows. Section 2 introduces the model. Section 3 discusses two examples of adverse selection caused by private action and how commitments against renegotiation are of value in this case. Section 4 introduces our assumptions on common knowledge and a decreasing utility frontier. Section 5 presents the main result. Section 6 explores a more specialized model in which the agent’s preferences are additively separable over time and exponential in each period’s consumption and his consumption and savings decisions are unobservable. With unobserved savings, Condition 1 requires that the agent’s wealth not influence his preferences, which is the reason we must assume exponential utility. We investigate the situation when the technology is of a kind that seems most relevant for factory laborers and some salesmen: Each period’s efforts are assumed to affect only that period’s production. We show that if the environment is stationary, the optimal long-term contract prescribes piece rates and commission schemes, which can be achieved with short-term contracts. Thus, in response to the inquiry in the opening paragraph, our model suggests an answer based on the information conditions and nature of work.

The properties of our model when the interest rate is small are studied in Section 7, where our conclusions are analyzed in relation to the literature on “folk theorems” in game-theoretic models with moral hazard. We show that with little discounting the first-best utility frontier can be

approached. Since this is achieved by a sequence of short-term contracts, this result cannot be interpreted to mean that long-term contracts and repeat dealings alleviate the single period incentive problem. Rather, the rationale for the near first-best performance in our model is based on self-insurance. Section 8 discusses extensions and is followed by concluding remarks in Section 9.

2. THE MODEL

We consider a multi-period principal-agent model with time indexed $t = 0, 1, \dots$. Periods are to be construed as the minimum length of time of commitment to a contract or alternatively as dates of potential renegotiation; of course, long-term contracts may permit commitments beyond one period. Within each period, events take place in the following order: (i) a contract is negotiated (not relevant if a long-term contract is in effect), (ii) the agent takes a productive action, (iii) he consumes, (iv) he observes a private signal, (v) the principal and the agent observe a public outcome, and (vi) the principal pays the agent.³ We denote the agent's productive action in period t by e_t ("effort"), his consumption by c_t , his private signal by σ_t , the public outcome by x_t , and the principal's payment by s_t .

The action e_t is a vector that includes any choice by the agent which affects his utility or the outcome of the firm. We assume consumption is one-dimensional, measuring c_t in dollars expended.⁴ The agent's private signal, σ_t , is a vector that may contain information about the agent's preferences or the technology. Note that e_t , as an arbitrary vector, can include strategies that map signals σ_t into productive decisions. Thus, the model can handle reverse timing of actions and signals within a period.

The public outcome x_t is a vector that records everything new the principal and the agent jointly observe in period t before payment of s_t . It may include information about the agent's past or present actions, consumption decisions, and private signals, as well as any period- t messages from the agent to the principal.⁵ The principal's period- t profit π_t is assumed public information by specifying that $\pi_t = \pi_t(x_t)$.

Histories are denoted by superscripts. Thus, $e^t = (e_0, \dots, e_t)$, $c^t = (c_0, \dots, c_t)$, and so on. We let the vector $h^t \equiv (x^t, s^t)$ record the public

³ Assuming a different order of events within a period is equally acceptable. Some convention is necessary for consistent notation.

⁴ A more general formulation incorporates consumption as part of the action vector. This formulation allows consumption to be multi-dimensional. In addition one must specify a mapping (deterministic or stochastic) that determines the agent's expenditures in each period as a function of his action. The analytics of this case is only notationally different.

⁵ Period- t messages, if present, are part of the action vector e_t .

history of events through time t and the vector $z^t \equiv (e^t, c^t, \sigma^t, x^t, s^t)$ the corresponding full history. Occasionally, we will refer to z^t and h^t as *nodes* in the unfolding tree of events. At the end of period t , the agent knows z^t and the principal knows h^t . Thus, at no time does the principal know more than the agent. At the beginning, both sides have identical information, so we let the initial history be $z^{-1} = \emptyset$.

Regarding payments we make the following important assumption.

Assumption 1 (Verifiability). Payments to the agent at time t , s_t , can be conditioned on all publicly available information up to that time, x^t .

Thus, we make no distinction between observable and verifiable information. Third parties, responsible for enforcing contractual promises, can also observe the vector of public outcomes, x_t . Note that Assumption 1 makes no reference to contracts. The distinction between long-term and short-term contracts is one of commitment, not one of contingencies.

We assume there is a *terminal date* T such that after period T , profits will no longer depend on the agent's actions (past or present), there will be no additional information arriving, and the principal will make no further payments to the agent. Formally:

Assumption 2 (Finite Contract Term). For $t \geq T+1$, $\pi_t = 0$, $x_t = 0$, $\sigma_t = 0$, and $s_t = 0$.

The agent may (but need not) continue to consume forever (in one of our applications it is essential that the agent consume over an infinite horizon; see Section 6). However, we will not model this explicitly in the main analysis. Instead we will express post- T consumption preferences in the form of an indirect utility over terminal wealth w_{T+1} . The agent's wealth is determined by his savings and borrowing options to be specified below.

The von Neumann–Morgenstern utility of the agent is initially specified as a general function:

$$U(e_0, \dots, e_T, c_0, \dots, c_T, \sigma_0, \dots, \sigma_T, w_{T+1}). \quad (2.1)$$

The agent cares about his consumption, his actions, and his final wealth. Potentially, his preferences are affected by the signals σ_t . Special cases of (2.1) will come up later.

The principal is risk neutral and evaluates profit and payment streams through the net present value:

$$\sum_{t=0}^T \delta^t [\pi_t - s_t]. \quad (2.2)$$

The stochastic *technology* is described by a set of probability distributions $\{F_t; t=0, \dots, T\}$ over public outcomes x_t and signals σ_t . Each

distribution $F_t(x_t, \sigma_t | e^{t-1}, \sigma^{t-1}, x^{t-1}, e_t)$ depends in general on the action, signal, and outcome histories up to time t as well as on the agent's action in period t . This formulation permits outcomes in different periods to be time dependent as well as stochastically dependent, either through the agent's actions or through exogenous events. Also, observe that our specification permits the agent to be inactive in any period (by assuming that his action in that period is without consequence); in particular, he may retire before the contract termination date T .

Long-Term Contracts

Since the agent consumes and acts in period t contingent on the information z^{t-1} , the agent's *plans of action and consumption* take the forms $\{e_t(z^{t-1})\}$ and $\{c_t(z^{t-1})\}$. We will often write e for the agent's action plan and c for his consumption plan.

We adopt the convention that plans of action and consumption specify what the agent will do at all conceivable nodes z^{t-1} , even nodes that can be reached only by following some alternative plan. For instance, if the action plan \hat{e} specifies \hat{e}_0 as its initial action, then of course \hat{e} can never lead to a node z^{t-1} for which $e_0 \neq \hat{e}_0$. Yet, we will insist that \hat{e} specify what the agent will do at node z^{t-1} . This convention will facilitate the discussion of recontracting later on. (Until then, the convention can be ignored.) Note that since $z^{-1} = \emptyset$, the date 0 action and consumption plans e_0 and c_0 are always singletons.

A *long-term contract* is a triple $\Delta = (e, c, s)$, where (e, c) are to be construed as the instructions (or suggestions) for the agent's effort and consumption plans above and $s \equiv \{s_t(x^t)\}$ is a *payment plan*, which specifies what the principal promises to pay the agent as a function of the publicly available information x^t . As with action and consumption plans, a payment plan specifies payments for all conceivable outcome contingencies x^t .

DEFINITION. A long-term contract $\Delta = (e, c, s)$ is *incentive compatible* if the agent finds it optimal to follow the instructions (e, c) , that is, if

$$(e, c) \underset{(\hat{e}, \hat{c})}{\text{Maximizes}} E[U(\hat{e}, \hat{c}, \sigma^T, w_{T+1}(\hat{c}, s)) | \hat{e}]. \quad (2.3)$$

The expectation in (2.3) is taken with respect to the distribution of the stochastic process $\{z_t\}$ generated by the agent's plan (\hat{e}) as indicated by the conditioning. If a contract is incentive compatible, the agent is willing to follow instructions. Note that (2.3) does not require (e, c) to be optimal conditional on nodes that cannot be reached, since it is a matter of indifference what the agent plans to do in events that cannot happen along the optimal path. A stronger notion, requiring incentive compatibility conditional on every history, will be introduced later.

The expected utility and the expected profit from a long-term contract Δ , evaluated at date 0, are

$$U_0(\Delta) \equiv E[U(e, c, \sigma^T, w_{T+1}(c, s)) | e], \quad (2.4)$$

and

$$\Pi_0(\Delta) \equiv \sum_{t=0}^T \delta^t E[\pi_t - s_t(x^t) | e]. \quad (2.5)$$

DEFINITION. An incentive compatible long-term contract Δ is *efficient* if there is no other incentive compatible long-term contract $\tilde{\Delta}$ that both parties prefer, i.e., $(U_0(\tilde{\Delta}), \Pi_0(\tilde{\Delta})) > (U_0(\Delta), \Pi_0(\Delta))$, where the inequality allows one, but not both components of the vectors to be equal.

An efficient long-term contract that guarantees the principal expected profits Π_0 solves the program

$$\text{Maximize } U_0(\Delta), \text{ subject to} \quad (2.6)$$

- (i) program (2.3)
- (ii) $\Pi_0(\Delta) \geq \Pi_0$.

DEFINITION. An efficient long-term contract that guarantees the principal zero expected profits at time zero, i.e., solves (2.6) with $\Pi_0 = 0$, is called *optimal*.

Our focus on optimal contracts is motivated by the idea that competition in the market for agents will force the principal to offer the agent the best zero profit contract. Later on, in our discussion of sequential contracting, a similar zero profit constraint will be imposed at each recontracting date to describe the outcome of short-term contracting.

We postpone a formal definition of short-term contracting until after the introduction of common knowledge, since we wish to avoid a general discussion of sequential contracting under adverse selection, which still is an imperfectly understood topic.

Banking

In order to isolate the incentive concerns from the issue of intertemporal smoothing of consumption we will assume that the agent has free access to a bank.

Assumption 3 (Equal Access to Banking). The agent can borrow and save at a secure bank, with a savings or loan balance of 1 at the end of today growing to $1/\delta$ at the beginning of tomorrow. The bank allows the agent

to borrow and save any amount up to time T . Post- T transactions are embodied in the agent's preferences over terminal wealth.

With this banking assumption, the wealth of the agent at the beginning of period $t + 1$, w_{t+1} , is determined by

$$w_{t+1}(z^t) = \delta^{-(t+1)} \left[w_0 + \sum_{\tau=0}^t \delta^\tau (s_\tau(x^\tau) - c_\tau(z^{\tau-1})) \right], \quad (2.7)$$

where w_0 is the wealth that the agent starts out with. To indicate explicitly that w_{T+1} depends on c and s we often write $w_{T+1}(c, s)$.

Remarks. 1. Assumption 2 implies that the agent and the principal can transact in the capital market on equal terms (the principal also discounts payments at the rate δ).

2. Banking may lead to a negative balance at the end of time T , and hence negative consumption afterwards. If desired, this possibility can be ruled out by specifying the utility function over terminal wealth w_{T+1} so that the agent voluntarily keeps w_{T+1} nonnegative. What is important is that the bank can enforce repayment of debts at least as economically as the employer could.

3. Two payment plans $s = \{s_t(x^t)\}$ and $\tilde{s} = \{\tilde{s}_t(x^t)\}$, which have the same net present values along every complete path x^T , offer the agent the same consumption opportunities and hence the same expected utility under any given action plan. Therefore, if the contract $\Delta = (e, c, s)$ is incentive compatible, so is the contract $\tilde{\Delta} = (e, c, \tilde{s})$, implying that the agent and the principal will both be indifferent between Δ and $\tilde{\Delta}$.

By assuming that the timing of payments is unimportant, we have removed all potential advantages from deferred payment plans such as pension plans.

3. EXAMPLES

Before proceeding to the general theory, we pause to discuss two illustrative examples. They help motivate the common knowledge conditions we will be introducing shortly. The point of the examples is to indicate how adverse selection, caused by the agent's choice of action (Example 1) or by his choice of consumption (Example 2), prevents short-term contracts from emulating optimal long-term contracts. Related examples by Baron and Besanko [2] and Laffont and Tirole [17] show that adverse selection caused by exogenous information asymmetries leads to similar short-term contracting problems.

EXAMPLE 1. There are two periods: $t=0, 1$. The agent works only in the initial period. Work involves choosing either a low level of effort or a high level of effort: $e_0=0$ or 1. Outputs in the two periods, x_0 and x_1 , are distributed according to the probability mass function $f(x_0, x_1|e_0)$.

The agent is risk and work averse. His preferences are represented by a utility function of the form $U(e_0, c_0, c_1, w_2) = u(w_2) - e_0$; that is, the agent does not consume until after period 1. The interest rate is zero, c_0 and c_1 are restricted to be non-negative, and the function u is assumed increasing, strictly concave, and unbounded from below. Clearly, the agent will set $c_0 = c_1 = 0$ so $w_2 = s_0 + s_1$.

Because the agent works only once, this is essentially a standard "single-period" agency model with the twist that the outcome is revealed over time. Let $(e_0^*, s_0^*(x_0), s_1^*(x_0, x_1))$ be an optimal long-term contract. We assume that $e_0^* = 1$, that is, it pays to implement a high level of effort.

A familiar result from single-period analyses (Holmstrom [13], Shavell [27]) tells that the agent's optimal period-1 compensation, s_1^* , must generally depend on x_1 unless x_0 is a sufficient statistic for e_0 , that is, unless $f(x_1|e_0=1, x_0) = f(x_1|e_0=0, x_0)$ for all (x_0, x_1) . The intuition is simple. If x_0 is not a sufficient statistic for e_0 , x_1 contains additional information (about the agent's action) worth incorporating into the contract.

However, consider what will happen if the parties can renegotiate the contract at date 1. At that point, the agent has no further actions, so there is no longer an incentive reason to let the final payment s_1 be contingent on either x_0 or x_1 . In fact, efficient risk sharing at date 1 dictates that the agent bear no risk at all. Therefore, we can expect the initial contract to be renegotiated to make the payment s_1 a constant. But if that is the case, the optimal long-term contract cannot be implemented as intended.⁶

Commitments have value here because the optimal long-term contract is not sequentially efficient, i.e., it is not immune to renegotiation. Sequential inefficiency can be traced to the presence of adverse selection. In our example, adverse selection emerges after the initial period because the agent is better informed than the principal about the distribution of future outcomes.

Contrast this with the case in which x_0 is a sufficient statistic for e_0 . Now a one-period contract will do as well as a long-term contract: Because s_1^* will not depend on the (uninformative) outcome x_1 , all payments conditional on x_0 can be made in period 0 (see earlier Remark 3). Put

⁶ Fudenberg and Tirole [8] provide an analysis of renegotiation-proof contracts in situations like this. Note that the agent, foreseeing a renegotiation to full insurance, would shirk if the optimal long-term contract were attempted. Consequently, an equilibrium must involve randomization unless the agent is instructed to choose low effort. Milgrom [19] makes a similar observation in the case where no contracting is allowed until after the agent's effort decision.

differently, without any information worth waiting for, the contract can be closed out after period 0. Note well the simultaneous absence of adverse selection: if the sufficient statistic condition holds, the principal, after observing x_0 , is as informed as the agent about the distribution of x_1 .

The next example shows that adverse selection problems can also arise from unobserved consumption decisions.

EXAMPLE 2. Suppose that the agent's preferences are described by the utility function $U(e_0, e_1, c_0, c_1, w_2) = u(c_0) + u(c_1) - e_1$, with the additional restrictions c_0, c_1 , and $w_2 \geq 0$. Note that we assume that the agent consumes in both periods, but works only in the last one. No outcomes are observed in period 0 ($x_0 \equiv 0$). The productive outcome in period 1 may be either a success ($x_1 = S$) or a failure ($x_1 = F$). The probability of success is p if the worker supplies high effort and q if he supplies low effort, where $q < p$. Consumption c_0 is not observed by the principal, so contracts cannot be indexed on it. The interest rate is zero. Again, we shall argue that the optimal long-term contract is vulnerable to renegotiation.

To verify this, let us first examine the two-period efficient contract. The compensation scheme takes the form $(s_0, s_1(x_1))$. By Remark 3, following the banking assumption, we can set $s_0 = 0$ and have all payments take place at the end of the horizon. Write $s_1(S) = w_S$, $s_1(F) = w_F$, and $c_0 = c$. An optimal compensation scheme is a solution to

$$\text{Max}_{\hat{c}, w_F, w_S} \{u(\hat{c}) + pu(w_S - \hat{c}) + (1-p)u(w_F - \hat{c}) - 1\}, \quad (3.1)$$

subject to

$$\begin{aligned} &u(\hat{c}) + pu(w_S - \hat{c}) + (1-p)u(w_F - \hat{c}) - 1 \\ &= \text{Max}_c \{u(c) + pu(w_S - c) + (1-p)u(w_F - c) - 1\} \\ &\geq \text{Max}_c \{u(c) + qu(w_S - c) + (1-q)u(w_F - c)\}, \end{aligned} \quad (3.2)$$

and

$$pw_S + (1-p)w_F \leq \bar{w}. \quad (3.3)$$

This optimization problem incorporates an incentive constraint (3.2) and a participation constraint (3.3). The incentive constraint requires that the worker would rather work diligently than shirk, given that in each case he can adapt his consumption decision to his effort choice, and that the consumption choice contemplated in the contract is the preferred choice of a

diligent worker. The participation constraint specifies a maximum expected wage (and hence a minimum expected profit) for the firm. This formulation is equivalent to (2.6).

At an optimal solution (\hat{w}_S, \hat{w}_F) to the problem (3.1)–(3.3), $\hat{w}_F < \hat{w}_S$; otherwise the no-shirking constraint (3.2) cannot be satisfied. Let c_H be the optimal choice of initial consumption for a worker who plans to work hard and c_L the optimal choice for a worker who plans to be lazy. These are unique since u is strictly concave. The first-order conditions determining the consumption choices are

$$u'(c_H) = pu'(\hat{w}_S - c_H) + (1 - p)u'(\hat{w}_F - c_H),$$

and

$$u'(c_L) = qu'(\hat{w}_S - c_L) + (1 - q)u'(\hat{w}_F - c_L),$$

From these equations and the facts that $\hat{w}_F < \hat{w}_S$, $p > q$, and u' is decreasing, it follows that $c_H > c_L$. The worker consumes more in the first period when he plans to be diligent because his income is (stochastically) greater then and first-period consumption is a normal good.

Now suppose that the parties sign the optimal long-term contract and that the employee, planning to be diligent, consumes c_H . Once he has made the initial consumption decision, the employee strictly prefers to work hard in the second period, as seen from

$$\begin{aligned} u(c_H) + pu(\hat{w}_S - c_H) + (1 - p)u(\hat{w}_F - c_H) - 1 \\ \geq u(c_L) + qu(\hat{w}_S - c_L) + (1 - q)u(\hat{w}_F - c_L) \\ > u(c_H) + qu(\hat{w}_S - c_H) + (1 - q)u(\hat{w}_F - c_H). \end{aligned} \quad (3.4)$$

The first inequality in (3.4) is just a restatement of the incentive-compatibility condition (3.2); the second follows because c_L is the unique optimal choice for an agent who plans to shirk.

Once we see that the incentive constraint does not hold with equality along the equilibrium path, it is clear that the optimal long-term contract is vulnerable to renegotiation: Given that the agent has consumed c_H , an efficient continuation contract must make the agent just indifferent about his choice of effort. More precisely, the following equality will hold:

$$pu(w_S - c_H) + (1 - p)u(w_F - c_H) - 1 = qu(w_S - c_H) + (1 - q)u(w_F - c_H). \quad (3.5)$$

However, according to (3.4), the wages (\hat{w}_S, \hat{w}_F) specified by the unique optimal two period contract do not satisfy (3.5).

We conclude that private information about preferences at recontracting dates (adverse selection) can make optimal long-term contracts vulnerable to renegotiation. Commitments to a non-renegotiable long-term contract can therefore benefit both parties.

4. COMMON KNOWLEDGE AND DECREASING UTILITY FRONTIER

We proceed to the major assumptions of the model, beginning with common knowledge conditions on technology and preferences that will eliminate adverse selection problems of the kind discussed above.

Common Knowledge

Since the agent has more detailed information than the principal at all times, the principal's information is always common knowledge. Thus, the assumptions on common knowledge only involve limitations on the agent's information advantage. Our conditions will imply that, given any history up to and including period t , the principal knows the agent's preference ordering over all potential contracts that may be offered at the beginning of period $t + 1$.

Assumption 4 (Common Knowledge of Technology). At the beginning of each period t and for all possible histories z^{t-1} ,

$$F_t(x_t, \sigma_t | e^{t-1}, \sigma^{t-1}, x^{t-1}, e_t) = F_t(x_t, \sigma_t | x^{t-1}, e_t). \quad (4.1)$$

Assumption 4 says that the information provided by the public observations x^{t-1} is sufficient to determine how period t 's actions will affect future outcomes and signals. Condition (4.1) is obviously satisfied if periods are independent of each other as in repeated principal-agent models. But we stress that (4.1) is much weaker than independence. Past actions and signals can affect current outcomes and signals as long as these dependencies are publicly revealed. To illustrate, suppose $\sigma_t \equiv 0$ for all t . Then the Markov technologies (4.2) satisfy (4.1),

$$x_t = h(x_{t-1}, e_t, \varepsilon_t), \quad (4.2)$$

where the ε_t 's are independent stochastic disturbances unobserved by both parties and the function h is arbitrary. Using (4.2), we can include in the model transitory or non-transitory components of the economy, the industry, or the agent's technology, provided information regarding these components is symmetric at recontracting dates.

Importantly, Assumption 4 requires that x_t convey no *new* information about past actions e^{t-1} . Information about the agent's activities cannot

arrive with a lag as in Example 1 of the previous section. Other technologies which violate (4.1) on these grounds include

$$x_t = e_t + e_{t-1} + \varepsilon_t, \quad (4.3)$$

$$x_t = e_t + \varepsilon_{t-1} + \varepsilon_t. \quad (4.4)$$

Condition (4.1) also rules out traditional adverse selection in which the agent receives payoff relevant information (through σ_t) without taking any actions (e.g., dynamic insurance). The models by Baron and Besanko [2] and Laffont and Tirole [17] violate (4.1), because in these models the agent privately learns a productivity parameter, which remains unchanged over time ($\sigma_0 = \sigma_1$).⁷

Our common knowledge assumption regarding agent preferences will be stated in terms of preferences over contingent action–payment streams. Let $\tilde{e}^T = (\tilde{e}_1, \dots, \tilde{e}_T)$ and $\tilde{s}^T = (s_1, \dots, s_T)$ denote a random action and a random income vector, respectively. Let P be a probability measure over $(\tilde{e}^T, \tilde{s}^T)$ and let E_P denote expectations with respect to the measure P . We write $\tilde{e}^T \setminus e^t$ for the random action vector $(e_0, \dots, e_t, \tilde{e}_{t+1}, \dots, \tilde{e}_T)$ in which the actions in the first t periods have been fixed at e^t . We use similar notation for random payment streams. The notation $\hat{c} \setminus c^t$ refers to a *continuation plan of consumption*, defined as a plan in which the first t periods of consumption are fixed at $c^t = (c_1, \dots, c_t)$ and the succeeding periods are governed by the terms of the long-term plan \hat{c} . Define

$$V_{t+1}(P|z^t) \equiv \max_{\hat{c}} E_P[U(\tilde{e}^T \setminus e^t, \hat{c} \setminus c^t, \sigma^T, w_{T+1}(\hat{c} \setminus c^t, \tilde{s}^T \setminus s^t)) | z^t]. \quad (4.5)$$

This represents the maximal expected utility that the agent can obtain by choosing an optimal consumption plan for the future given the history z^t of past consumptions, signals, and outcomes and a probability measure P over the random stream $(\tilde{e}^T, \tilde{s}^T)$. We emphasize that the value V_{t+1} refers to preferences over probability distributions of future action–*payment* streams (in which consumption is optimal), rather than arbitrary action–*consumption* streams. The following assumption asserts that these preferences are common knowledge.

Assumption 5 (Common Knowledge of Preferences over Action–Payment Streams). For all t and any two distributions P and R ,

$$\{z^t: V_{t+1}(P|z^t) > V_{t+1}(R|z^t)\} \in \sigma(h^t),$$

where $\sigma(h^t)$ is the σ -algebra generated by the public history $h^t = (x^t, s^t)$.

⁷ With an appropriate change of variables, the productivity parameter becomes a cost parameter, in which case the models can alternatively be viewed as violating common knowledge of preferences (Assumption 5).

Assumption 5 can be satisfied in several ways, all involving some form of intertemporal separability of preferences. Suppose the signals σ_t are independently distributed preference parameters. Then if the agent's utility is either additively separable,

$$U(e^T, c^T, w_{T+1}) = \sum_{t=0}^T \delta^t u_t(e_t, c_t, \sigma_t) + g(w_{T+1}), \quad (4.6)$$

or multiplicatively separable,

$$U(e^T, c^T, w_{T+1}) = \prod_{t=0}^T u_t(e_t, c_t, \sigma_t) g(w_{T+1}), \quad (4.7)$$

and if the agent's consumption (hence wealth) is observable, Assumption 5 holds. Here the agent's private information only relates to past u_t values, while separability assures that these values do not affect preferences over future action-income streams. Various special cases of (4.6), with consumption publicly observed, have been studied frequently. Allen [1] and Green [9] adopted the "repeated insurance problem" version of this model, which is defined by the additional restrictions that $u_t = u_t(c_t, \sigma_t)$ (with the σ_t 's independent). Malcomson and Spinnewyn [18] and Rogerson [24] treated models satisfying (4.6) with the restriction $u_t = u_t(c_t, e_t)$.

Even when consumption is not observed, Assumption 5 may be satisfied provided there are no wealth effects on preferences. This happens in (4.7) if in addition

$$u_t \equiv -\exp[-r(c_t - v(e_t, \sigma_t))] \quad \text{and} \quad g(w_{T+1}) = -\exp[-rw_{T+1}]. \quad (4.8)$$

The preferences resulting from (4.7) and (4.8) with $v(e_t, \sigma_t) = v(e_t)$ were used in the principal-agent models of Fellingham, Newman, and Suh [6] and Holmstrom and Milgrom [15]. Our main theorem below applies to each of these models as well as the ones mentioned earlier, thus helping to clarify the relationships among them. In Section 6 we shall study another model without wealth effects to which our theorem applies—one in which the agent's consumption is unobserved and preferences are additively exponential ((4.6) and (4.8) hold with $v(e_t, \sigma_t) = v(e_t)$).

Recontracting

Given our convention that actions, consumptions, and payments in a long-term contract are defined for all conceivable histories, we can view recontracting at time t simply as the adoption of a new long-term contract. So let $\Delta = (e, c, s)$ be a long-term contract adopted at time $t + 1$, after the history $\hat{z}^t = (\hat{e}^t, \hat{c}^t, \hat{\sigma}^t, \hat{h}^t)$ (Δ may, of course, have been in effect also before

$t + 1$). By a *continuation* of Δ at \hat{z}^t , denoted $\Delta \setminus \hat{z}^t$, we mean the collection of future plans in Δ , which remain relevant given the history \hat{z}^t (i.e., plans that prescribe actions, consumptions, and payments for contingencies that follow \hat{z}^t). The agent's expected utility from the long-term contract Δ given history \hat{z}^t , i.e., his valuation of the continuation $\Delta \setminus \hat{z}^t$, is

$$U_{t+1}(\Delta \setminus \hat{z}^t) \equiv E[U(e \setminus \hat{e}^t, c \setminus \hat{c}^t, \sigma^T \setminus \hat{\sigma}^t, w_{T+1}(c \setminus \hat{c}^t, s \setminus \hat{s}^t) | e \setminus \hat{e}^t, \hat{z}^t)], \quad (4.9)$$

where, as before, $c \setminus \hat{c}^t$ denotes the agent's continuation plan of consumption and $e \setminus \hat{e}^t, s \setminus \hat{s}^t$ are analogously defined continuation plans of action and payment. If the principal knew \hat{z}^t , his corresponding valuation of the continuation $\Delta \setminus \hat{z}^t$ would be given by the period- $(t + 1)$ expected present value of profits,

$$\Pi_{t+1}(\Delta \setminus \hat{z}^t) \equiv \sum_{\tau=t+1}^T \delta^{\tau-t-1} E[\pi_{\tau} - s_{\tau}(x^{\tau}) | e \setminus \hat{e}^t, \hat{z}^t]. \quad (4.10)$$

Note that $\Pi_{t+1}(\Delta \setminus \hat{z}^t)$ excludes profits realized prior to period $t + 1$.

DEFINITION. A continuation $\Delta \setminus \hat{z}^t$ is *incentive compatible* if, given the payment plan $s \setminus \hat{s}^t$, the agent prefers the action-consumption plan $(e \setminus \hat{e}^t, c \setminus \hat{c}^t)$ to any other plan $(\tilde{e} \setminus \hat{e}^t, \tilde{c} \setminus \hat{c}^t)$ (as measured by (4.9)). A long-term contract Δ is *sequentially incentive compatible* if for every t and every \hat{z}^t , the continuation $\Delta \setminus \hat{z}^t$ is incentive compatible. The set of sequentially incentive compatible long-term contracts is denoted SIC.

Sequential incentive compatibility means that the agent is willing to follow the instructions in Δ conditional on any history. This is the proper consistency condition given that Δ may be started at any node z^t , following some other long-term contract.

Since the principal cannot observe \hat{z}^t , he generally cannot tell how the agent values a contract Δ , nor how he himself would value Δ if he had the agent's information. Therefore, negotiations over which contract to adopt at node z^t will typically have to take place under asymmetric information about the value of alternative options. Such negotiations (and their analysis) can be expected to be complicated, because of adverse selection. However, when technology and preferences are common knowledge (Assumptions 4 and 5), these problems vanish: for every history \hat{z}^t , and every pair of sequentially incentive compatible contracts Δ and Δ' we have

$$\{z^t: U_{t+1}(\Delta \setminus z^t) > U_{t+1}(\Delta' \setminus z^t)\} \in \sigma(\hat{h}^t), \quad (4.11)$$

$$\{z^t: \Pi_{t+1}(\Delta \setminus z^t) > \Pi_{t+1}(\Delta' \setminus z^t)\} \in \sigma(\hat{h}^t). \quad (4.12)$$

According to (4.11) the principal knows how the agent ranks incentive compatible continuation contracts at any node z' .⁸ This is weaker than assuming that $\{z': U_{t+1}(\Delta \setminus z') = \bar{u}\} \in \sigma(\hat{h}')$ for all \bar{u} , i.e., that the principal knows the actual value of the agent's conditional expected utility. The distinction is important when the agent's consumption cannot be observed (Section 6). By contrast, (4.12) is equivalent to assuming that the principal knows the expected profit of each incentive compatible continuation contract, since contracts that give all future profits to the agent in exchange for a fixed rental fee will provide the requisite calibration.

If we let $Z(\hat{h}') = \{z' | h' = \hat{h}'\}$ denote the set of histories consistent with the principal's observation \hat{h}' , conditions (4.11) and (4.12) can be rephrased as saying that the agent's and the principal's preference orderings over contracts never change within a set $Z(\hat{h}')$.

Decreasing Utility Frontier

We turn to the assumption that the utility frontier is downward sloping, along with conditions on the agent's preferences sufficient to ensure that the assumption is satisfied.

The *utility possibility set* conditional on history z' is the set of feasible payoff pairs $\text{UPS}(z') \equiv \{(\pi, u) | \exists \Delta \in \text{SIC} \text{ such that } \pi = \Pi_{t+1}(\Delta \setminus z'), u = U_{t+1}(\Delta \setminus z')\}$. The *utility frontier* conditional on z' is the function $\pi = \text{UPF}(u | z')$, obtained by maximizing the principal's expected profit among all payoff pairs in $\text{UPS}(z')$, which give the agent a *fixed* utility u . The *efficient frontier* conditional on history z' is the set of undominated feasible payoff pairs $\text{EF}(z') \equiv \{(\pi, u) \in \text{UPS}(z') | \nexists (\pi', u') \in \text{UPS}(z') \text{ such that } (\pi', u') \geq (\pi, u) \text{ and } (\pi', u') \neq (\pi, u)\}$. Points in $\text{EF}(z')$ are obtained by maximizing the principal's expected profit among payoff pairs in $\text{UPS}(z')$ which give the agent *at least* a utility level u , say.

DEFINITION. A continuation $\Delta \setminus z'$ is *efficient* if $\Delta \in \text{SIC}$ and if the payoffs of $\Delta \setminus z'$ are on the efficient frontier $\text{EF}(z')$.

Assumption 6 (Decreasing Utility Frontier). For every history z' the function $\text{UPF}(u | z')$ is strictly decreasing in u .

Restated, Assumption 6 says that the efficient frontier coincides with the utility frontier. The significance for our analysis is that under Assumption 6 one can replace any contract with an efficient contract without altering the agent's payoff. Thus, the full range of agent incentives can be provided within the set of efficient contracts. Figure 1 depicts a utility frontier that

⁸ Note that incentive compatibility is essential. The principal may not be able to rank contracts in which the agent is assumed to consume in a suboptimal way, as is the case in the example of Section 6.




FIGURE 1

violates Assumption 6. Even though point A is inefficient, one cannot move to a Pareto preferred point without raising the agent's utility.

In one-period models, requirements on minimum payments ($s_t \geq \bar{s}$) or limited liability may cause Assumption 6 to fail; the efficiency wage model of Shapiro and Stiglitz [26] and the quality-assuring price model of Klein and Lefler [16] are well-known examples.⁹ In our model no minimum payments are imposed. Provided preferences are separable and the agent's consumption can be observed, the frontier of the utility possibility set will be downward sloping:

THEOREM 1. *If consumption is observable, then Assumption 6 is satisfied when either of the following two conditions holds:*

- (i) *Preferences satisfy (4.6) (additive separability over time) and the function g is increasing, continuous, and unbounded below.*
- (ii) *Preferences satisfy (4.7) (multiplicative separability over time), each u_t is positive, the function g is increasing and continuous, and either g is negative and unbounded below or it is positive and has greatest lower bound zero.*

Proof. We give the argument for case (i); case (ii) is analogous.

Let z^t be an arbitrary history, Δ any sequentially incentive compatible long-term contract, and k an arbitrary positive number. Let the continuation $\Delta \setminus z^t$ have payoffs (π, u) . We shall construct another sequentially incentive compatible contract $\hat{\Delta}$ such that $\hat{\Delta} \setminus z^t$ has payoffs $(\hat{\pi}, \hat{u})$ with $\hat{\pi} > \pi$

⁹ For a dynamic model in which bonds are built up over time, see Borland [3].

and $\hat{u} = u - k$, i.e., the agent's utility is reduced by k and the principal's utility is strictly higher. This will establish the claim.

The contract \hat{A} calls for the same actions, consumptions, and payments as A , except that the final payment, \hat{s}_T , is determined by the equation $g(w_T + \hat{s}_T - c_T) = g(w_T + s_T - c_T) - k$, for each pair (w_T, c_T) . A solution (\hat{s}_T) to the equation can be found, since g is increasing continuous, and unbounded from below and the solution will be a function of x^T as required, since (w_T, c_T) is a function of x^T when consumption is observable. Compared to A , the contract \hat{A} reduces the agent's realized utility by the amount k along every complete history z^T . Consequently, the agent is willing to act the same way under either contract, proving that the new contract \hat{A} is also sequentially incentive compatible. The principal's payoff is higher, since the agent's behavior stays the same and the new payments are smaller. Q.E.D.

In Section 6 we will see that even if consumption is not observable, Assumption 6 will hold for the additively exponential case.

5. MAIN RESULT

As the examples in Section 3 suggest, a long-term contract must be immune to renegotiation if one hopes to emulate it by a sequence of short-term contracts. The following notion of sequential efficiency is a sufficient condition for a long-term contract to be renegotiation-proof.¹⁰

DEFINITION. A long-term contract A is *sequentially efficient* if for every history z' , $A \setminus z'$ is an efficient continuation.

Sequential efficiency is a strong requirement. Since the payments of a long-term contract only depend on the public history x' , this does not provide enough flexibility in general to maintain payoffs on the efficient frontier in all contingencies z' . Indeed, neither of the examples in Section 3 admits sequentially efficient contracts.¹¹ The following key result (which,

¹⁰ That sequential efficiency implies renegotiation-proofness is obvious given common knowledge. But even without common knowledge it is irrational for the principal to agree to replace a sequentially efficient contract, because if the agent accepts the change, the principal must infer that the new contract gives him something less, since there is no contract that makes both better off.

¹¹ Note that even when there are no sequentially efficient contracts, there typically will exist renegotiation-proof contracts. Renegotiation-proof contracts are determined subject to the extra constraint that future renegotiations cannot be desirable for both parties. (For analyses of renegotiation-proof contracts under asymmetric information, see for instance Dewatripont [5] and Hart and Tirole [12].)

we stress, does not require the banking assumption), shows that sequentially efficient contracts always exist under the conditions we have postulated:

THEOREM 2. *Assume verifiability of public outcomes (Assumption 1), a finite contracting horizon (Assumption 2), common knowledge of technology and preferences (Assumption 4 and 5), and a decreasing utility frontier (Assumption 6). Then, for any efficient long-term contract, there is a corresponding sequentially efficient long-term contract providing the same initial expected utility and profit levels.*

Proof. Let Δ be an efficient long-term contract. Suppose it is not sequentially efficient. Then, given some history \hat{z}' , there exists an efficient continuation $\Delta' \setminus \hat{z}'$, which Pareto dominates the continuation $\Delta \setminus \hat{z}'$. By Assumption 6, we can assume $\Delta' \setminus \hat{z}'$ provides the same expected utility to the agent and a strictly higher expected profit to the principal. By Assumptions 4 and 5, it is common knowledge at \hat{z}' that $\Delta' \setminus \hat{z}'$ has this property, i.e., for all $z' \in Z(\hat{h}')$, the agent is indifferent between $\Delta \setminus z'$ and $\Delta' \setminus z'$ and the principal strictly prefers $\Delta' \setminus z'$.

Now construct a new contract that differs from the original contract only in that its terms, once the event \hat{h}' occurs, are those specified by Δ' . We claim the revised contract is incentive compatible. As noted above, for all $z' \in Z(\hat{h}')$, the agent's conditional expected utility is unchanged. Also, his expected utilities conditional on any $z' \notin Z(\hat{h}')$ are unchanged, since the subsequent terms are unchanged. By the Optimality Principle of Dynamic Programming, the agent's incentives up to time t are therefore unchanged. Consequently, the agent is precisely as well off and the principal weakly better off with the revisions. By doing a similar substitution in all contingencies z' for which $\Delta \setminus z'$ is not an efficient continuation, we arrive at a sequentially efficient contract. Q.E.D.

The logic of our proof is the same as the one used to argue that ex ante optimality implies ex post optimality in complete markets. This is because common knowledge together with verifiability of public outcomes admits complete contracting in the limited sense that all contingencies that are relevant for forecasting the (payoff-relevant) future can be contracted on at time t .

Finally, we are ready to state what we mean by short-term contracting. When the assumptions of Theorem 2 hold, it is reasonable to postulate that any continuation contract, agreed to at time t conditional on some history z' , will have payoffs on the efficient frontier $EF(z')$. Furthermore, if we envision principals competing for the agent's services, the contract will be one which gives the agent the highest expected utility subject to the principal's making non-negative profits. This motivates the following:

DEFINITION. A sequentially efficient contract which gives the principal zero expected profits conditional on any history z^t is called *sequentially optimal*.¹²

A sequentially optimal contract has the feature that if the agent and the principal were to terminate their relationship (cancel their contract) at any time and start negotiating for a new long-term contract immediately afterwards, the old contract would be accepted anew. Working backwards from date T , it is then clear that a sequentially optimal long-term contract signed at date 0 can be decomposed into a sequence of short-term contracts negotiated at the beginning of each period and only specifying payments and plans for that period.

By adding the banking assumption to the conditions of Theorem 2, sequential optimality follows from sequential efficiency by a simple rearrangement of payments. This will establish the main result of the paper:

THEOREM 3. *Assume verifiability of public outcomes (Assumption 1), a finite contracting horizon (Assumption 2), equal access to banking (Assumption 3), common knowledge of technology and preferences (Assumptions 4 and 5), and decreasing utility frontier (Assumption 6). If there is an optimal long-term contract, then there is a sequentially optimal contract, which can be implemented via a sequence of short-term contracts.*

Proof. Let Δ be an optimal long-term contract. By definition, $\Pi_0(\Delta) = 0$. By Theorem 2, there is a sequentially efficient contract Δ' which provides the same initial payoffs as Δ , in particular $\Pi_0(\Delta') = 0$. By Assumption 2, $\Pi_{T+1}(\Delta' \setminus z^T) \equiv 0$, for every z^T .

We will modify the timing of payments to make expected profits zero from each node z^t onwards. Let

$$\hat{s}(x^t) = s'(x^t) - \delta \Pi_{t+1}(\Delta' \setminus z^t) + \Pi_t(\Delta' \setminus z^{t-1}), \quad \text{for } t = 0, 1, \dots, T.$$

By common knowledge the right hand side varies with x^t only, so this construction is possible. For every complete history z^T , the present value of the agent's compensation is the same under \hat{s} as under s' (by telescoping series), so the contract $\hat{\Delta} \equiv (e', c', \hat{s})$, which substitutes \hat{s} for s' in Δ' , is incentive compatible (cf. Remark 3 following Assumption 3). By construction expected profits in $\hat{\Delta}$ are zero from each node onwards so the contract is sequentially optimal.

The identification of sequential optimality with short-term contracting was explained before the theorem. Q.E.D.

¹² We assume that there is an efficient contract in which the principal makes zero profits. This follows from Assumption 6, if there is some contract that gives non-positive profits and if the set $\text{UPS}(z^t)$ is convex. The latter condition can be assured by randomizing among contracts.

6. ADDITIVELY SEPARABLE EXPONENTIAL UTILITY

In general, if consumption cannot be observed, it will not be true that preferences are common knowledge, as we saw in Example 2, and therefore a commitment to a long-term contract will be of value. However, we do not think that unobserved consumption is an empirically significant reason for long-term contracts. Firms seem to make little effort to monitor employee wealth, which would be of value if consumption effects were important. It seems reasonable therefore to look for an explanation of observed contract characteristics using models in which Assumption 5 (Common Knowledge of Preferences) is satisfied even when consumption is not observed. One such model is specified below.

The agent observes no periodic private signals σ_t . Preferences display constant absolute risk aversion and additive separability over time, as follows:

Assumption 7 (Exponential Utility). The agent's utility function is

$$-\sum_{t=0}^{\infty} \delta^t \exp\{-r(c_t - v(e_t))\}, \quad \text{with } v(e_t) \equiv 0 \text{ for } t > T. \quad (6.1)$$

Assuming that the agent has access to a bank in all periods using the same discount factor δ as that in (6.1), he will find it optimal after period T to consume just the interest on his terminal wealth w_{T+1} :

$$c_t = (1 - \delta) w_{T+1} \quad \text{for } t > T. \quad (6.2)$$

Consequently, the specification (6.1) corresponds to (4.9) with

$$g(w_{T+1}) \equiv -\delta^{T+1}(1 - \delta)^{-1} \exp\{-r(1 - \delta) w_{T+1}\}. \quad (6.3)$$

THEOREM 4. *Suppose Assumptions 1–4 and 7 hold. Then the agent's preferences will be common knowledge (Assumption 5) and the utility frontier will be decreasing (Assumption 6).*

Proof. Fix τ , $0 < \tau < T$, and a history z^τ . Let $w_{\tau+1} = w_{\tau+1}(z^\tau)$ be the agent's wealth at the beginning of period $\tau + 1$ as defined earlier in (2.3).

Define $d_t = c_t - (1 - \delta) w_{\tau+1}$, for $T \geq t \geq \tau + 1$. This represents consumption in excess of interest on the wealth the agent has at the beginning of time $\tau + 1$. Simple algebra shows that

$$w_{T+1} = w_{\tau+1} + \sum_{t=\tau+1}^T \delta^{t-T-1} (s_t - d_t). \quad (6.4)$$

Consequently, the agent's preferences at time τ over future action-consumption streams can be represented by

$$-u((1-\delta)w_{\tau+1}) \left\{ \sum_{t=\tau+1}^T \delta^t u[d_t - v(e_t)] + \delta^{T+1}(1-\delta)^{-1} u \left[(1-\delta) \sum_{t=\tau+1}^{\infty} \delta^{t-T-1} s_t - d_t \right] \right\}, \quad (6.5)$$

where $u(w) = -\exp(-rw)$. To establish (6.5) we have used the fact that $u(ab) = -u(a)u(b)$ as well as the representation (4.6) and (6.3) of the agent's utility function.

Given a probability distribution P over action-payment streams, we compute $V_{\tau+1}(P|z')$ (see (4.5)) by maximizing (6.5) over consumption plans, or equivalently over d . From (6.4) and (6.5), the agent's expected utility takes the form $u((1-\delta)w_{\tau+1}) E_P[f(d)]$ for some function $f(\cdot)$. So his maximal expected utility takes the following form for some function $H(\cdot)$:

$$V_{\tau+1}(P|z') = u[(1-\delta)w_{\tau+1}] H(P). \quad (6.6)$$

Assumption 5 follows immediately from (6.6). In intuitive language, the agent's preferences over distributions P are fully determined by $H(P)$, which does not depend on z' . With additively separable utility, history affects current preferences only through wealth effects. In deriving Eq. (6.6) we have established that, with exponential utility of periodic consumption, wealth effects are absent, too, so preferences are common knowledge.

Assumption 6 also follows from (6.6). A reduction of a fixed amount in the agent's compensation in any period (say, the first) does not affect his preferences over lotteries. Hence, such a change preserves incentive-compatibility, lowers the agent's utility, and raises the principal's expected profits. This shows that the initial utility frontier is downward sloping, and a similar argument applies to renewal dates. Q.E.D.

In view of Theorem 4, the conclusion of Theorem 3 applies when Assumptions 1-4 and 7 hold. We can obtain stronger conclusions by strengthening Assumption 4 as follows.

Assumption 8 (History-Independent Technology). For all $0 \leq t \leq T$, $F_t(x_t | e^{t-1}, x^{t-1}, e_t) = F_t(x_t | e_t)$. There are no signals.

Assumption 9 (Stationary, History-Independent Technology). For all $0 \leq t \leq T$, $F_t(x_t | e^{t-1}, x^{t-1}, e_t) = F(x_t | e_t)$. There are no signals.

According to Assumption 8, the outcome at date t depends on date t actions alone, and not on past outcomes or actions. We take this assump-

tion to be a tolerable approximation for the work of a laborer engaged in a repetitive task, where the outcome of each successive operation affects the quality of one particular item, or the effort exerted over each item affects the time to completion for that item. The assumption might also apply to one who sells consumer goods, abstracting from any unobserved investments the salesman may have to make in such things as his reputation, knowledge of the stock and current styles, etc. For such situations, the optimal incentive compensation schemes are modified piece-rate or commission rules, in which the commission rate or piece-rate may vary over time. It is obvious that when such rules are optimal, there is no gain from having a long-term employment relationship. With the additional Assumption 9, requiring that the environment be stationary, a standard commission or piece-rate scheme emerges as optimal.

THEOREM 5. *Suppose that Assumptions 1–3, 7, and 8 hold and that there exists an optimal long-term contract. Then there is an optimal contract for which*

- (i) *current instructions and payments do not depend on past performance: $e_t(z^{t-1}) = e_t$ and $s_t(x^t) = s_t(x_t)$,*
- (ii) *the principal's expected profit in every period is zero, and*
- (iii) *action and payment plans are identical to those in the optimal contract that would be offered in the "one-period problem" in which the agent retires at the end of the initial period ($T=0$) and the available technology is that of period t .*

COROLLARY. *Suppose, in addition to the hypotheses of Theorem 4, that Assumption 9 holds. Then there is an optimal contract in which, for all t , $e_t = e_0$, and $s_t(x^t) \equiv s(x_t)$. Thus the net present value of the agent's total compensation when he retires with history z^T is*

$$\sum_{t=0}^T \delta^t s(x_t).$$

Proof of Theorem 5. Assumptions 5 and 6 are implied by Theorem 4 and Assumption 5 is implied by Assumption 7. Applying Theorem 3, there is a sequentially optimal contract Δ . By Assumption 7, the efficient frontier $EF(z^t)$ is (up to an affine rescaling) independent of the past, and hence Δ can be chosen to be history-independent in the sense (i). Part (ii) follows since Δ is sequentially optimal.

In view of conclusion (i) and (6.6), the agent's maximal expected future utility at the beginning of period $\tau + 1$, if his current wealth is $w_{\tau+1} = w$

and he is employed under the sequentially efficient contract Δ , is expressible as

$$\begin{aligned} \text{Max}_c u(c - v(e_{\tau+1})) + \delta E\{u[(1 - \delta) \delta^{-1}(w + h(P) - c + s_{\tau+1}(x_{\tau+1}))]\} \\ = H(P) \text{Max}_d u(d - v(e_{\tau+1})) + \delta E\{u[(1 - \delta) \delta^{-1}(-d + s_{\tau+1}(x_{\tau+1}))]\}, \end{aligned} \quad (6.7)$$

where P is the probability distribution of future incomes and efforts under Δ , $h(\cdot)$ is defined by $H(P) = -u[(1 - \delta)(w + h(P))]$, and we have made the change of variables $d = c - (1 - \delta)(w + h(P))$. In view of (6.7), the agent's preferences over action-income lotteries (or short-term contracts) for period $\tau + 1$ does not depend on T , the continuation lottery P , or the date $\tau + 1$, except through the technology specified for the date. Obviously, the same is true for the principal's preferences. Hence, the efficient action and payment plans in Δ only depend on the period t technology, and (iii) is verified. Q.E.D.

The optimal contract specifies the employee's actions each period as a function of the current technology and his compensation as a function of the current outcome (which depends only on the current action). The contract requires no "memory", and the ability to provide correct incentives in this model is not enhanced by having the employee write a long-term contract (or have a long-term relationship) with the employer. Further, each period's contract is the same as it would be if this were the only period in which the agent worked; hence repetition in no way helps to ameliorate the single period incentive problem. We should emphasize that these "one-period contracts" are not the same as those which would be optimal if the agent only *lived* for one period; even when the agent works only one, he lives (and consumes) infinitely often.¹³

The Corollary asserts that when the environment is stationary the agent's aggregate compensation depends on the total number of times each possible outcome has occurred, corrected for discounting. Thus the optimal contract is linear in suitably defined accounting aggregates. (Caution: this does not mean the contract is linear in, for example, the total dollar volume of the agent's sales, but rather that it depends on the number of sales of each possible dollar amount.) This result corresponds to the result on aggregation over time obtained by Holmstrom and Milgrom [15] for a multiplicatively separable exponential specification.

¹³ Actually, one can show that (i)-(ii) of Theorem 5 continue to hold even when the agent is finitely lived after retirement. In that case, (iii) fails because the agent's preferences over contracts depend on the length of his remaining lifetime. The optimal compensation scheme for each period will then depend both on the current technology and on the number of periods of remaining life.

7. SMALL DISCOUNT RATES

It is of interest to relate our conclusion, that short-term contracts may be optimal, to the literature on “folk theorems” in infinitely repeated principal-agent games (Radner [21, 22], Rubinstein [25] Fudenberg, Levine, and Maskin [7]). These papers show that when there is little or no discounting, efficient (first-best) payoffs of the one shot game can be approximated by the normalized discounted payoffs of a perfect equilibrium of the repeated game. In other words, in a long-term relationship agency costs are negligible if players are sufficiently patient. For this to be consistent with our results, it must be that if our banking assumption is added to the repeated game models, asymptotic efficiency can be attained by a series of short-term contracts. We will show that this is indeed the case.

Intuitively, the agent is quite tolerant of single-period risks when δ is close to 1, because by smoothing he can translate large periodic income variations into small changes in his periodic consumption. To illustrate, consider our exponential model. Let $T=0$ so that the agent only works in period 0. His consumption in all future periods will be $(1-\delta)w_1$ and his lifetime utility $u(c_0 - e_0) + (\delta/(1-\delta))u((1-\delta)w_1)$. His indirect utility for period 1 wealth is thus $u((1-\delta)w_1)$, which has a coefficient of absolute risk aversion equal to $(1-\delta)r$ (cf. (6.2) and (6.3)). For δ close to one, the agent is almost risk neutral with respect to income at date 0. As one might expect, the first-best outcome can then be approximated by a contract $s_0(x_0) = \pi_0$, in which the agent is fully responsible for the consequences of his actions. The fact that the agent acts only once and yet first-best can be approximated arbitrarily closely, serves to underscore that, in our model repetition does not alleviate incentive problems because of improved monitoring.

In a finite horizon model without productive actions, Yaari [29] proved that, as the number of periods grows, a completely patient consumer can approach the full insurance outcome (in average utility) by smoothing his consumption. Following Yaari’s basic logic, we can generalize the exponential case discussed above. Consider an infinitely repeated game with stationary, history-independent technology $F(x_t | e_t)$. Let $\pi_t(x_t) = x_t$ for all t and assume that $F(\cdot | \cdot)$ has compact support. As is standard in the repeated game literature, we normalize the agent’s utility function U in the repeated game (and throughout this section) so that payoffs in the repeated game are measured in per period terms:

$$U = (1-\delta) \sum_{t=0}^{\infty} \delta^t u(c_t, e_t).$$

Let e^* be the agent’s level of effort in the first-best optimal contract, let

$\underline{x}^* = \min\{x \mid x \in \text{support } F(x \mid e^*)\}$, and let $m^* = E[x \mid e^*]$. We assume that it is *feasible* for the agent to consume any $c_t \geq \delta \underline{x}^*$.

Now we introduce the bank. In previous sections we assumed that the agent could not obtain new loans once he retired. The formalization of banking constraints in the infinitely repeated model is more delicate. The obvious formulation is that the agent's wealth be non-negative "at infinity," but it is hard to say how this type of constraint could be enforced. Instead we use a *more restrictive* constraint: we require that the agent is *never* allowed to borrow, i.e., that he is restricted to consumption plans along which his wealth w_t is *always* non-negative. We assume $w_0 = 0$. With these model specifications we can prove the following:

THEOREM 6. *Let the principal pay the agent $s_t(x_t) = x_t$ in every period. Then for every $\varepsilon > 0$, there exists a $\delta(\varepsilon) < 1$, such that the agent can ensure himself a utility level $u(m^*, e^*) - \varepsilon$ for all $\delta > \delta(\varepsilon)$.*

The proof is in the Appendix. The idea of the proof is to construct a strategy which guarantees with high probability that the agent is able to consume approximately the mean output in every period after a finite number of periods. The strategy specifies that the agent chooses an efficient effort level and consumes close to the mean output unless wealth falls below a critical level in which case he consumes the minimum output. Wealth under this strategy will follow a submartingale with bounded increments. Therefore with a probability arbitrarily close to one, the agent will eventually stay above the critical wealth level for all periods to come.

Fudenberg, Levine, and Maskin [7] prove folk theorems for general classes of repeated games with moral hazard which include principal-agent games as a special case. As they explain, their results imply that when players are patient, there is no conflict between the need to provide incentives and the desire to insulate players from risk. This is not to say that risk aversion is unimportant when players are patient: If two players must share a risky endowment stream and do not have access to a storage technology, then even in the limit as $\delta \rightarrow 1$ they cannot do as well as with a deterministic endowment stream of the same expected value. Note also that while a risk neutral principal could serve as a substitute for a bank, the principal's promise of future rewards is only credible if equilibrium strategies make honoring his promises a best response.

8. EXTENSIONS AND APPLICATIONS

Elimination of Third-Party Banking

In a recent paper, Rey and Salanie [23] have demonstrated how, in a certain class of environments, long-term commitments can be supported by

a series of short-term contracts. Their essential idea is that long-term banking arrangements can be supported by a series of shorter term contracts in which the deposit (or loan) is rolled into a new deposit (or loan) as desired at the start of each period. Thus, in our model, if we allow the agent to make a deposit (or loan) with the principal from the end of period t to the beginning of period $t+1$ for each t , our conclusions about short-term contracting could be extended to the case where there is no third-party banker. Note that in this case consumption is effectively observed, since the principal can observe the agent's deposits and withdrawals.

We wish to stress that a sequence of short-term contracts need by no means be "simpler" than the corresponding long-term contract. The labor contract model in Harris and Holmstrom [11] provides a good illustration. The optimal long-term contract is a simple wage guarantee. In principle, the same wage guarantee could be achieved by a sequence of one-period contracts in which the firm acts as a short-term insurer and banker, but evidently such an arrangement would be more complex and cumbersome.

Multiple Agents

Joint production, in which several agents contribute to the success of an enterprise, is a natural extension. Under conditions similar to those specified in the preceding sections, efficient contracts can always be replaced by equivalent sequentially efficient ones.

The analysis is somewhat different when the "principal" works, and so is in effect another agent. It is well-known that in one-period models, joint production of this sort can lead to free-rider problems (even with all parties risk neutral), if one requires that all returns be distributed among the productive agents (see Holmstrom [14]). It can be shown that if compensation rules are subject only to the weaker constraint that the total compensation in any period not exceed that period's gross revenues (with the balance being effectively disposed of, for example by a "charitable" contribution) and if natural analogues of the Common Knowledge and Decreasing Utility Frontier conditions hold, then for every efficient contract there is an equivalent sequentially efficient one. Consequently, long-term contracts and relationships do not help to solve the free rider problem.

Incomplete Contracts

Our analysis is of relevance for the emerging literature on incomplete contracts (Williamson [28], Grossman and Hart [10]), because it identifies a range of cases in which short-term contracts are sufficient to support efficient arrangements, even when imperfectly observed long-lived investments are required. Hence, our analysis does not support the

common argument that relationship-specific investments must always be protected by long-term contracts in order for proper investments to be made. This point is developed more fully by Crawford [4] and Milgrom and Roberts [20].

9. CONCLUSION

In the introduction we motivated our study in terms of the observed variety in incentive schemes used for different types of workers. Why are managerial workers paid differently than salesmen or factory workers? We have found an answer in the information conditions of their work. Many more of the activities of managers than of factory workers or salesmen contribute directly to future 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.¹⁴

Our formal analysis has considered lifelong and single period contracts, but most actual employment contracts are of more moderate terms. What can be said about the relationship between the length of contracts and the extent of information lags? Our results suggest an obvious conjecture: The benefits of extending contract length are positively related to the length and extent of the information lag. Consequently, one would expect contracts to be designed to balance the gains from incorporating all the information relevant to the current contract period against the costs of lengthening the contract term. Similarly, our analysis suggests the conjecture that employee turnover in jobs that do not exhibit substantial information lags is higher than in jobs that do. Further development of our model will be needed to generate a set of testable hypotheses of this sort.

APPENDIX

Proof of Theorem 6. Fix $\bar{\delta} < 1$. For any $\delta > \bar{\delta}$, one feasible strategy for the agent is to choose the first-best action $e_t = e^*$ and, for any fixed $\gamma = (0, m^* - x^*)$, to consume according to

$$\hat{c}_t = \begin{cases} m^* - \gamma\delta + (1 - \delta)w_t, & \text{if } w_t \geq (m^* - x^*)/\delta, \\ \delta x^* + (1 - \delta)w_t + (1 - \delta)m^*, & \text{if } w_t < (m^* - x^*)/\delta. \end{cases} \quad (\text{A.1})$$

¹⁴ Note that when long-term contracts are desirable, the question about the parties' ability to commit not to renegotiate them becomes important. For instance, if a manager's contract were conditioned on information that arrives after he retires, then as in Example 1 of Section 3, there would be a desire to renegotiate the deal upon retirement. Fudenberg and Tirole [8] note that renegotiation threats may explain why managers, unlike workers, frequently are offered a choice from a menu of compensation schemes (because of adverse selection).

In (A.1), when γ is small, the agent consumes close to his mean income plus the interest on his wealth, unless his wealth level falls precariously low. Note that with this plan of consumption the agent's wealth will follow a stochastic process $\{w_t\}$. Mean wealth will progress according to

$$E[w_{t+1}|e^*, w_t, h^t] = \begin{cases} w_t + \gamma & \text{if } w_t \geq (m^* - \underline{x}^*)/\delta \\ w_t + (m^* - \underline{x}^*)/\delta & \text{if } w_t < (m^* - \underline{x}^*)/\delta. \end{cases} \quad (\text{A.2})$$

So $E[w_{t+1}|e^*, w_t, h^t] \geq w_t + \min\{\gamma, m^* - \underline{x}^*\}$, that is, $\{w_t\}$ is a submartingale with drift bounded away from zero. Since the increments of w_t are uniformly bounded for $\delta \in (\bar{\delta}, 1)$ (because F has compact support), there is for all $\rho \in (0, 1)$ a finite $T(\rho)$ such that for all $\delta \in (\bar{\delta}, 1)$, $\text{Prob}\{w_t \geq (m^* - \underline{x}^*)/\delta \text{ for all } t \geq T(\rho)\} \geq 1 - \rho$. Hence, the agent's utility from working efficiently and following the consumption plan in (A.1) is at least

$$(1 - \rho) \delta^{T(\rho)} u(m^* - \gamma\delta, e^*) + [1 - (1 - \rho) \delta^{T(\rho)}] u(\underline{x}^*, e^*). \quad (\text{A.3})$$

By inspection of (A.3), for all $\varepsilon > 0$ there exist ρ, γ and $\delta(\varepsilon)$ such that the bound (A.3) exceeds $u(m^*, e^*) - \varepsilon$ for all $\delta > \delta(\varepsilon)$. Q.E.D.

REFERENCES

1. F. ALLEN, Repeated principal-agent relationships with lending and borrowing, *Econ. Lett.* **17** (1985), 27-31.
2. D. BARON AND D. BESANKO, "Commitment and Fairness in a Dynamic Regulatory Relationship," unpublished working paper, 1987.
3. J. BORLAND, "Optimal Labor Contracts with Moral Hazard and Limited Liability," unpublished dissertation, Yale University, 1988.
4. V. CRAWFORD, "Long-Term Relationships Governed by Short-Term Contracts," Economics Discussion Paper 86-22, University of California, San Diego, 1986.
5. M. DEWATRIPONT, Renegotiation and information revelation over time: The case of optimal labor contracts, mimeo, Harvard University, 1985.
6. J. FELLINGHAM, D. NEWMAN, AND Y. SUH, Contracts without memory in multiperiod agency models, *J. Econ. Theory* **37** (1985), 340-355.
7. D. FUDENBERG, D. LEVINE, AND E. MASKIN, *The folk theorem in repeated games with imperfect public information*, mimeo, MIT, 1989.
8. D. FUDENBERG AND J. TIROLE, "Moral Hazard and Renegotiation in Agency Contracts," unpublished working paper, MIT, 1989.
9. E. GREEN, Lending and the smoothing of uninsurable income, in "Contractual Arrangements for Intertemporal Trade" (E. Prescott and N. Wallace, Eds.), pp. 3-25, Minneapolis, University of Minnesota, 1987.
10. S. GROSSMAN AND O. HART, The costs and benefits of ownership: A theory of vertical and lateral integration, *J. Polit. Econ.* **94** (1986), 691-719.
11. M. HARRIS AND B. HOLMSTROM, A theory of wage dynamics, *Rev. Econ. Stud.* **49** (1982), 315-33.

12. O. HART AND J. TIROLE, "Contract Renegotiation and Coasion Dynamics," MIT Working Paper, 1987.
13. B. HOLMSTROM, Moral hazard and observability, *Bell J. Econ.* **10** (1979), 74–91.
14. B. HOLMSTROM, Moral hazard in teams, *Bell J. Econ.* **13** (1982), 324–40.
15. B. HOLMSTROM AND P. MILGROM, Aggregation and linearity in the provision of intertemporal incentives, *Econometrica* **55** (1987), 303–28.
16. B. KLEIN AND K. LEFFLER, The role of market forces in assuring contractual performance, *J. Polit. Econ.* **89**(4) (1981), 615–41.
17. J. LAFFONT AND J. TIROLE, The dynamics of incentive contracts, *Econometrica* **56** (1986), 1153–1175.
18. J. MALCOMSON AND F. SPINNEWYN, The multiperiod principal-agent problem, *Rev. Econ. Stud.* **55** (1988), 391–407.
19. P. MILGROM, "Adverse Selection without Hidden Information." University of California, Berkeley, IBER Working Paper No. 8742. 1987.
20. P. MILGROM AND J. ROBERTS, Bargaining costs, influence costs, and the organization of economic activity, in "Perspectives on Positive Political Economy" (J. Alt and K. Shepsle, Eds.), Cambridge, Cambridge University Press, 1990.
21. R. RADNER, Monitoring cooperative agreements in a repeated principal-agent relationship, *Econometrica* **49** (1981), 1127–1148.
22. R. RADNER, Repeated principal-agent games with discounting, *Econometrica* **53** (1985), 1173–1198.
23. P. REY AND B. SALANIE, "Long-Term, Short-Term and Renegotiation," INSEE Working Paper No. 8715, 1987.
24. W. ROGERSON, Repeated moral hazard, *Econometrica* **53** (1985), 69–76.
25. A. RUBINSTEIN, Offenses that may have been committed by accident—An optimal policy of retribution, in "Applied Game Theory" (S. J. Brams, A. Schotter, and G. Schrodiauer, Eds.), Wurzburg, Physica-Verlag, 1979.
26. C. SHAPIRO AND J. STIGLITZ, Equilibrium unemployment as a worker discipline device, *Amer. Econ. Rev.* **74**(3) (1984), 433–444.
27. S. SHAVELL, Risk sharing and incentives in the principal and agent relationship, *Bell J. Econ.* **10** (1979), 55–73.
28. O. WILLIAMSON, "The Economic Institutions of Capitalism," The Free Press, New York, 1985.
29. M. YAARI, A law of large numbers in the theory of consumer's choice under uncertainty, *J. Econ. Theory* **12** (1976), 202–217.