

Research articles

© Springer-Verlag 1994

Optimal principal agent contracts for a class of incentive schemes: a characterization and the rate of approach to efficiency*

Prajit K. Dutta1 and Roy Radner2

- Department of Economics, University of Wisconsin, Madison, WI 53706, USA and Department of Economics, Columbia University, New York, NY 10027, USA
- Mathematical Sciences Research Center, AT&T Bell Laboratories, Murray Hill, NJ 07974, USA and Department of Economics, New York University, New York, NY 10003, USA

Received: November 14, 1991; revised version August 15, 1993

Summary. In this paper we study a repeated principal-agent situation with moral hazard. We focus on a class of incentive schemes, called bankruptcy contracts. The agent is "scored" in each period, and is paid a fixed wage per period until the current score falls to zero, at which time the agent is terminated and the principal hires a new agent. The agent's current score at any time equals an initial score, plus the total output up to that time, minus an amount that is proportional to the total time. With standard assumptions about the utility functions of the principal and agent, we characterize the second-best bankruptcy contracts and show that in such a contract, the principal pays the agent an efficiency wage. We also demonstrate that such contracts lead to approximately first-best (Pareto efficient) outcomes if the principal and agent are sufficiently patient (have small discount rates). Most importantly, if the two players have a common discount rate δ , then the loss of efficiency under the second-best bankruptcy contract goes to zero at least as fast as $O(\delta^{1/2} \ln \delta)$. In order to obtain increased precision, the analysis is carried out in a continuous-time framework.

1 Introduction

A notable feature of many real life principal-agent contracts is that they specify simple compensation rules; i.e., they identify only a small set of contingencies on

^{*} This is a considerably revised version of a manuscript titled "Principal Agent Games in Continuous Time". We have benefited from the comments of seminar participants (on the earlier version) at Columbia, Northwestern, Illinois and Stony Brook. The revised version benefited from a presentation to the Game Theory Conference at Ohio State, July 1990. We would like to acknowledge helpful comments from Dilip Abreu, Tatsuro Ichiishi, Eric Maskin, William Rogerson and Aloysius Siow as well as an anonymous referee. This paper was begun while the first author was at AT&T Bell Laboratories and completed while visiting the Department of Economics at the University of Rochester; the two organizations are thanked for all help rendered.

The views expressed here are those of the authors, and not necessarily those of AT&T Bell Laboratories.

484 P.K. Dutta and R. Radner

which the agent's rewards are conditioned. This is difficult to reconcile with the theoretical work on second-best contracts, i.e., the optimal contract when no restrictions have been placed on the types of incentive schemes that are admissible (see [16], [21] and [25]). These papers demonstrate that the second-best contracts should subtly condition on various elements of an agent's past performance.¹

Perhaps this seeming paradox could be resolved by modelling the costs of contracting explicitly. In this paper we start instead by restricting ourselves to a set of simple contracts; the contracts we study pay the agent a constant wage and use the threat of dismissal as an incentive device. Furthermore, a simple statistic of past performance is employed to determine whether or not the agent is fired. The schemes studied have some of the stylized features of observable contracts such as managerial contracts which track past profits to determine tenure. Insurance contracts in which full indemnity converge is provided only if the number of past claims is no larger than a prespecified number will be seen to be a second example of the class of incentive schemes studied in this paper. Sales or franschise contracts which are renewed only if the volume of past business is sufficiently large, is a third example.²

1

The simple version of contingent-dismissal schemes that we study here was introduced by Radner (1986), who called them bankruptcy schemes. He showed that such incentive schemes generate almost efficient outcomes if the principal and the agent are sufficiently patient. We add to the analysis of that paper by explicitly characterizing the optimal contract within the class of bankruptcy schemes. Such a characterization demonstrates the optimal mix of incentive and insurance consideratations in such contracts.

It is well-known that if the principal and the agent are not infinitely patient, i.e., if they do not have discount rates of zero, exact optimality cannot be sustained even if the agency relationship is infinitely repeated. Consider the case of a risk-neutral principal and risk-averse agent. Optimal risk-sharing requires that in any Pareto optimal outcome the agent's compensation must be constant; hence, no "punishment" is feasible since compensation is independent of the only observable variable, the output generated by the agent's action. Since the first-best or efficient outcome is not sustainable, any result on asymptotic efficiency is necessarily an approximate one. We know that there are long-term contracts that are approximately efficient, as the players' common discount rate goes to zero (for instance, see Fundenberg-Maskin [9], Radner [18], [19], [20], Rubinstein [22] and Rubinstein-Yaari [23]). If we believe, however, that the "true" model involves

¹ [8], [10] and [12] have however established that under some specifications of preferences for principal and agent, history-independent short-term contracts are (constrained) optimal. See the discussion in Section 7.

² Some observable contracts add bonus provisions (contingent on immediate performance) if the agent meets his target; managerial compensations typically are of this form. On the other hand, the performance standard is simplified in certain observable contracts; adequate performance is determined by the number of periods in which the target was met (with no regard to the amount by which the standard was exceeded, if it was); an insurance contract based on the number of claims is of this form. Some of the results that follow can be adapted to these cases – see the discussion in Section 7.

³ With no discounting and an infinite horizon, Rubinstein [22] showed that exact optima could be sustained, whereas Radner [18] established the sustainability of approximate optimality in sufficiently

*

discounting, then the natural question is, how good are these approximations? For any $\varepsilon > 0$, how patient do principal and agent need to be in order to sustain ε -optimality, i.e., what is the rate of convergence to optimality in a repeated moral hazard situation? One purpose of this paper is to give a first answer to this question, within the restricted class of bankruptcy contracts. To the best of our knowledge, this is the first such estimate that is available in the literature on optimal contracts with asymmetric information.

We turn now to a summary of the model and results. We study a repeated principal-agent situation with moral hazard, in which in each period the observable "output" is a random variable that is a function of both the agent's action in that period and a random factor, neither of which is observed by the principal. We focus on a particular class of incentive contracts (called bankruptcy contracts), in which the agent is "scored" in each period as follows: (1) the agent is given an initial (positive) score, say y; (2) in each period, the current output is added to the score, and then a fixed "return", say k, is subtracted from it. The agent is paid a fixed wage per period, say w, until such time, if ever, as the current score reaches or falls below zero. At that time the agent's tenure is terminated, and the principal hires a new agent. The contract thus has three parameters, the initial score (y), the required "return" per period (k), and the wage per period (w). The agent's utility per period depends on both his action and the wage; whereas the principal's utility equals the output minus the wage (the principal is risk-neutral). As usual, we suppose that the agent's total utility is the expected value of the sum of his discounted one-period utilities, and that the principal's total unity is similarly determined. We further assume that the principal commits himself to a contract with each agent, and the agent then acts to maximize his own utility. Note that one interpretation of the "score" is that of a cash reserve. The agent is given an initial cash reserve, y, and in each period the principal withdraws from this reserve an amount k which is a "rate of return" that the agent is expected to maintain on average. When the cash reserve runs down to zero, the agent loses his job. For concreteness, we will use this interpretation in all further discussion.⁴

Within the class of contracts just described, we characterize those that are optimal for the principal, subject to the constraint that the agent can achieve some given minimum utility, i.e., we characterize the *second-best* contracts in this class. We also examine the efficiency properties of such contracts. A detailed discussion of the results follows.

Optimal choice of an agent. The agent's optimal choice problem is an example of a more general survival problem in stochastic control (see Dutta [6]). We use results from the general formulation to give a characterization of the optimal choice when the set of feasible drift-variance choices is an arbitrary convex, compact set in \mathbb{R}^2 (Theorems 3.1–3.3). It is shown that the optimal policy conditions the current action

long but finite horizon contexts. Fundenberg-Maskin [9] and Radner [19], [20] show that there exist contracts which approximate efficiency in discounted models, for sufficiently small discount rates. Indeed in all of these papers the contracts are additionally incentive compatible for the principal as well.

⁴ Note that our bankruptcy scheme is also a particular way of using a common measure of managerial performance, namely "residual income." (See, for example, [13], pp. 254 ff.)

486 P.K. Dutta and R. Radner

on the current level of the cash reserve alone. The optimal policy progressively shirks, that is the higher is the cash reserve level the higher is the agent's instantaneous utility and (under some additional conditions) the smaller the expected current output. A risk-neutral principal would like to have the control with the maximum mean used throughout, and the coasting by the agent at "safe" output levels is precisely a measure of the inefficiency of moral hazard, from the principal's point of view. As an illustration of the results obtained in this general model, we also report, without proof, the explicit form of the optimal policy in the binary case, in which the agent has only two actions (Theorem 3.4).

Principal's contract choice. The principal's return, with a (stochastic) succession of agents, is derived in Section 4. It is shown that in an optimal contract the initial cash reserve is the smallest one consistent with individual rationality. This is not a priori obvious since a lower initial reserve (or equivalently, a smaller acceptable loss) implies quicker bankruptcies and hence associated inefficiencies for the principal. On the other hand we get an efficiency-wage result, that the optimal wage is higher than the minimum wage consistent with individual rationality.

Pareto optimality. The first-best arrangement involves a constant control exercised by the agent, no dismissal and full insurance by the risk-neutral principal, by way of paying an outcome-independent wage (Propositions 5.1 and 5.2).

Asymptotic efficiency. We conclude the investigation of bankruptcy contracts by showing that the values under the optimal bankruptcy contract converge to the first best at a rate at least as fast as $O(\delta^{1/2} \ln \delta)$, as $\delta \to 0$, where δ is the (common) discount rate of principal and agent (Proposition 6.1). As an illustration, we compute the efficiency loss explicitly for a parametric example.

The principal-agent model is described in detail in Section 2. Section 3 discusses the optimal response of an agent to a bankruptcy compensation scheme. The principal's choice-of-contract problem is analyzed in Section 4. Section 5 contains the characterization of the Pareto optimal policies, whereas the analysis leading to a derivation of the rate of convergence to Pareto optimality is in Section 6. Bibliographic notes and a discussion of possible extensions of the current analysis may be found in Section 7.

2 The model

2.1 A "discrete-time" motivation

Suppose that the time axis is divided into successive intervals (periods) of length h. At the beginning of period n the agent chooses an action, which results in a stochastic output R(nh). Conditioned on the sequence of the agent's actions, the random variables R(nh) are independently distributed; however, the (marginal) distribution of R(nh) depends on the corresponding action. Let y be given and for $n \ge 0$ define the sequence Y(nh) by: Y(0) = y and

$$Y(nh + h) = Y(nh) + R(nh) - kh.$$
 (2.1)

Recall that we interpret Y(nh) as the agent's "cash reserve" at the end of period

n, y > 0 as the agent's initial cash reserve and k as the "return" per unit time that the principal requires from the agent. The principal pays the agent a constant wage w until such time (if ever) that the cash reserve reaches or falls below zero; at that time the agent is terminated and is replaced by another one.

Note that the sequence Y(nh) is a controlled random walk. If one lets h approach zero, and suitably adjusts the marginal distributions of the random variables R(nh) in the passage to the limit, then the sequence Y(nh) approaches a controlled diffusion which is a continuous time stochastic process with continuous sample paths. (For a more detailed discussion of the passage to the limit argument see [14, pp. 66-71].)

In what follows we work with such a continuous-time formulation of the problem. The reason for doing so is that while in continuous time an agent goes bankrupt the instant Y=0, in discrete time he could go bankrupt with any nonpositive level of cash reserve. The problem then is, how should the agent's continuation value depend on the level of terminal reserve. There is no obvious way in which to assign this value and any assignment clearly affects in a funamental way the agent's optimal choices while on the job. To avoid such "overshooting the boundary" problems we choose to model the principal agent question in continuous time. We turn now to the continuous time analog of (2.1), after introducing some preliminary concepts.

2.2 The continuous time model

Let $[B(t):t\geq 0]$ be a standard Brownian motion on some probability space (Ω, \mathcal{F}, P) ; for a definition see Karatzas-Shreve [14]. Let \mathcal{F}_t be the smallest family of sub σ -fields generated by the Brownian motion, i.e. \mathcal{F}_t is the smallest σ -field with respect to which B(s), $s\in[0,t]$ is measurable. Let $[f(\cdot,t):t\geq 0]$ be a \mathcal{F}_t -adapted process⁵ which further satisfies

i) For each $t \ge 0$, $[\omega: \int_0^t f^2(\omega, s) ds < \infty] = 1$ a.s.

The stochastic integral $\int_0^t f(\cdot, s) dB(s)$ is well-defined for all $t \ge 0$ a.e. A stochastic process $[\hat{Y}(t): t \ge 0]$ is said to be a diffusion if it can be written as:

$$\hat{Y}(t) = \hat{Y}(0) + \int_{0}^{t} m(s)ds + \int_{0}^{t} v^{1/2}(s)dB(s), \tag{2.2}$$

where $[m(t):t\geq 0]$ and $[v(t):t\geq 0]$ are \mathscr{F}_t -adapted and satisfy i), and $\hat{Y}(0)$ is some constant. The functions $m(\cdot)$ and $v(\cdot)$ are, respectively, the draft and variance components of the process.

In the principal-agent model, an agent controls a diffusion (the cumulative output) process. The agent's action is the choice of a feasible instantaneous drift-variance pair [m(t), v(t)]. Let the set of feasible mean-variance choices be denoted A. A choice at t conditions on the observable history of output during [0, t). An admissible strategy, π , for the agent is a pair of \mathscr{F}_t -adapted processes $[m(t): t \ge 0]$ and $[v(t): t \ge 0]$ in which $(m(t, \omega), v(t, \omega)) \in A$ for all (t, ω) and which lead to a solution

⁵ A stochastic process $[f(\cdot,t):t\geq 0]$ on (Ω,\mathcal{F}) is said to be \mathcal{F}_t -adapted if $f(\omega,t)$ is jointly measurable in ω and t, and ii) $f(\cdot,t)$ is \mathcal{F}_t -measurable, for each $t\geq 0$.

of the following stochastic differential equation (recall that k is the average return the agent is required to maintain):

$$Y(t) = Y(0) + \int_0^t m(Y(s))ds + \int_0^t v^{1/2}(Y(s))dB(s) - kt, \quad t \ge 0.$$

Note that there are several interpretations possible for the formulation in which the agent directly picks an instantaneous drift and variance. One interpretation is that the agent chooses from a menu of available projects or techniques, each involving different levels of supervision or skill or effort and having a mean and a variance.

To define a termination date, for any strategy π and initial cash reserve y > 0, let

$$T_{\pi}(y) \equiv \inf\{t \ge 0: Y(t) = 0 | Y(0) = y, \pi\}.$$

Let the constant wage paid be denoted w and let the agent's instantaneous utility function be called U. Furthermore, once fired (terminated), the agent receives a severance pay or a reassignment to a different position. We will normalize the value of such an option to zero. For any given (w, k, y) the discounted utility over an agent's uncertain lifetime, for a strategy π , is

$$g_{\pi}(y) = E\delta \int_{0}^{T_{\pi}(y)} e^{-\delta s} U(w, m(s), v(s)) ds.$$

To complete the formulation of the moral hazard problem let the principal's discounted lifetime earnings under a compensation triple (w, k, y) and agent's strategy π be denoted $H(w, k, y; \pi)$. (In Section 4 we will explicitly derive $H(\cdot)$.) Let the agent's reservation utility be denoted \hat{U} . Then the optimal contract choice problem for the principal is

$$\max_{(w,k,y)} H(w,k,y;\pi),$$
 s.t. $g_{\pi}(y) \ge g_{\pi'}(y)$ for any admissible π' , (2.3)

$$g_{\pi}(y) \ge \hat{U}. \tag{2.4}$$

Condition (2.3) is the incentive constraint and (2.4) is the individual rationality constraint.

The Pareto-optimality or first-best problem is that of maximizing the principal's discounted lifetime earnings H subject to the individual rationality constraint, i.e., in the absence of moral hazard. For this problem we shall not restrict the set of feasible contracts. The precise formulation is discussed in Section 5.

3 Incentive constraint analysis: the agent's problem

The agent's best response problem is: given w, k and y, maximize $g_{\pi}(y)$ over the set of admissible policies. This is clearly a stationary dynamic programming problem,

⁶ Strictly speaking, since the principal hires a new agent if and when the current agent fails to meet performance requirements, his returns are derived from a succession of compensation schemes offered and a succession of strategies followed by different agents. As we shall see in Section 4, the Optimality Principle implies stationarity in the compensation schemes and strategies and allows us to write H as a function of a single compensation triple (w, k, y) and a single strategy π .

and we shall denote its value function by V(y; w, k). In much of what follows, we shall concentrate on the effect of changes in the initial cash level, Y(0) (equivalently changes in the level of losses the agent is allowed to incur). Hence the dependence of the value function on w and k will frequently be suppressed and it shall be written simply as V(y), where y = Y(0). Any solution will be called an optimal strategy or policy. If an optimal policy picks controls that depend only on the level of current cumulative output, it will be called a stationary Markov optimal policy.

We make two assumptions throughout. The first says that there is some action that the agent can take which gives him a level of utility greater than the utility he derives after dismissal. Evidently, this is a minimal necessary assumption for a bankruptcy scheme to have any incentive effects at all. (Recall that we have normalized the utility level of dismissal to be zero.) The second assumption says that the agent's actions lead to uncertain outcomes. Clearly this is necessary for the principal's inference problem to be nontrivial. The assumptions are formally:

- (A0) $\sup_{(m,v)\in A}U(w,m,v)\equiv \bar{U}(w)>0 \text{ for all } w\geq 0.$
- (A1) $\inf\{v:(m,v)\in A\}>0$.

We also assume that

- (A2) The set of feasible controls A is a convex, compact set.
- (A3) The utility function U(w, m, v) is continuous and strictly concave in the last two arguments.

The agent's best response exercise is an example of a general survival problem in stochastic control (it is in fact a version of the gambler's ruin problem). In the formulation here, an instantaneous choice is being made simultaneously along three dimensions: drift, variance and utility. In previous investigations authors have allowed a choice over drift and variance (holding utility constant) or have allowed a choice over drift and utility (holding variance constant). Three dimensional trade-offs turn out to be extremely difficult to characterize in a transparent way. Dutta [6] has investigated the general control problem and we use those results to describe some basic properties of the agent's optimal choice. To add to the intuition we then report, without proof, the optimal policy in the case where the agent has only two actions available at every instant.

Let (w, k) be fixed until further notice. The following characterization of the value function holds:

Theorem 3.1.

- i) The value function V(y) is strictly increasing in y.
- ii) (Bellman equation) V is C^2 and satisfies the optimality equation

$$\max_{(m,v)\in A} \left\{ \frac{1}{2} v V''(y) + (m-k)V'(y) - \delta V(y) + \delta U(w,m,v) \right\} = 0, \quad y \ge 0.$$
 (3.1)

⁷ [11] and [17] among others, analyse related versions of the pure survival case where all controls have the same utility, whereas [3] and [5], among others, analyze problems in which all controls have the same variance. See [6] for further references.

iii) The marginal valuation satisfies the following

$$V'(y) = V'(0)Ee^{-\delta T^*(y)},$$

where $T^*(y)$ is the termination date under the agent's optimal policy. Consequently, the value function is strictly concave.

Remark. The proof of Theorem 3.1 may be found in Dutta [6].

The following comparative statics and boundary properties of the value function, with respect to the parameters (w, k), are easy to verify:

Proposition 3.1

- i) The value function V(y; w, k) is increasing in w, provided the utility function is increasing in w, and decreasing in k. Further, it is continuous in (w, k).
- ii) $\lim_{y \to \infty} V(y; w, k) = \overline{U}(w)$, for all k, and V(0; w, k) = 0, for all (w, k).

Turning to a characterization of the optimal strategy, we first define a stationary Markov policy $\beta: \mathbb{R}_+ \to A$ to be interior if $\beta(y) \in \text{int } A$ for all $y \in \mathbb{R}_+$. Further, the utility function is said to be separable if there exist functions $\xi_w(m)$ and $\phi_w(v)$ such that $U(w, m, v) = \xi_w(m) - \phi_w(v)$.

Theorem 3.2

- i) There is a unique stationary Markov optimal policy $\beta^* \equiv (m^*, v^*)$: $\mathbb{R}_+ \to A$, and this policy is given by the maximizers from (3.1). Furthermore, β^* is a continuous function.
- ii) The optimal strategy has the property that as the cumulative cash reserve grows, the agent switches to higher variance and/or lower mean options. In other words, y' > y implies that either or both of the following conditions hold:

a)
$$v^*(y') \ge v^*(y)$$
 or

b)
$$m^*(y') \le m^*(y), \frac{m^*(y') - k}{v^*(y')} \le \frac{m^*(y) - k}{v^*(y)}.$$

iii) Suppose that U is separable and β^* is interior. Then, at high cumulative cash reserves, the agent employs high variance-low mean actions. Furthermore, if U is decreasing (respectively, increasing) in m (respectively, in v), then the agent's actions at higher cash reserves give him higher instantaneous utility. In other words, y' > y implies that $m^*(y') < m^*(y)$ and $v^*(y') > v^*(y)$ and $U(w, m^*(y'), v^*(y')) > U(w, m^*(y), v^*(y))$.

Proof. That any selection from the maximizers correspondence of the optimality equation defines a stationary Markov optimal policy follows from a standard argument via Ito's lemma (e.g. see Krylov [15, 1.1 and 1.4]). By the Maximum theorem of Berge [2] and the fact that the value function is C^2 , this correspondence is upper hemi-continuous. From the strict concavity of the utility function, the set of maximizers is actually single valued for every y. Hence this function, β^* , is continuous.

Suppose we denote the optimal choice at y' by (m', v') (respectively the optimal choice at y by (m, v)). Then it follows from the optimality equation that

$$\frac{1}{2}(v-v')[V''(y)-V''(y')]+(m-m')[V'(y)-V'(y')] \ge 0$$
 (3.2)

$$\left(\frac{m-k}{v} - \frac{m'-k}{v'}\right) [V'(y) - V'(y')] - \delta\left(\frac{1}{v} - \frac{1}{v'}\right) [V(y) - V(y')] \ge 0.$$
 (3.3)

Dutta [6], Theorem 3.1, establishes that V'' increases in y. That combined with Theorem 3.1 and (3.2)–(3.3) yields the second part of the theorem. In the separable-utility case, first-order conditions yield

$$V'(y) = -\delta \xi'_{w}(m^{*}(y)), \quad V''(y) = \delta \phi'(v^{*}(y)).$$

The third part of the theorem follows from the strict concavity of V (Theorem 3.1(iii)). \square

The order of usage of the drifts points directly to the inefficiency, from the principal's point of view, that persists under a bankruptcy incentive scheme. At low cumulative output levels, with the threat of dismissal near, the agent does in fact forego instant gratification to boost immediate returns for the principal. However at higher and safer levels, after a run of good luck or "hard work", the agent rests on his laurels. Of course, if the principal could renege on his commitment to the bankruptcy contract, this is precisely when he would like to do so, and dismiss an agent in order to hire a new one for whom the threat of dismissal is more effective. 8

In order to sharpen our intuition, we actually computed the agent's optimal policy in the specific case of a binary choice problem; the set of feasible actions contains two elements (m_1, v_1) and (m_2, v_2) . The results are reported below and the relevant computations may be found in [7]. Denote $U_i \equiv U(w, m_i, v_i)$, i = 1, 2 and suppose that $U_2 \ge U_1$.

The principal result states that faced with a bankruptcy scheme the agent finds it optimal to employ a switchpoint strategy of the following kind: above a critical aggregate cash reserve \hat{y} the agent uses control 2 while below \hat{y} the agent switches to the other control in order to improve tenure prospects. As long as the preference between the two controls is strict, i.e. $U_2 > U_1$, the agent must eventually shirk, i.e. $\hat{y} < \infty$. Typically the higher utility action will also be the control with a lower mean, and hence that which the principal does not want employed.

Consider the quadratic function $\frac{1}{2}v_ix^2 + (m_i - k)x - \delta = 0$, and denote by θ_i (resp. λ_i) the positive (resp. absolute value of the negative) root.

Theorem 3.3

ii) There is a unique stationary Markov optimal policy for the agent's problem, and this policy is a switchpoint strategy.

⁸ Such breach of contract brings into the picture further considerations of reputation effects for the principal. Further, a rational agent foreseeing such a possibility would also adjust his behavior. At some cost of complexity the present analysis could be extended to generate the commitment of the principal as a self-enforcing outcome. Given our focus we prefer just to assume that such a breach of contract is not possible.

⁹ A version of this problem was first studied by Sheng [24]. Since our formulation turned out to be not immediately covered by her analysis, we directly computed the optimal policy.

2

ii) Suppose that $U_2 > U_1$. Then the optimal switchpoint \hat{y} is finite. It is zero if and only if

$$\xi_1(\lambda_2) \ge \frac{\delta(U_1 - U_2)}{U_2},\tag{3.4}$$

where $\xi_1(\lambda_2) = \frac{1}{2}v_1(-\lambda_2)^2 + (m_1 - k)(-\lambda_2) - \delta$.

- iii) Suppose $U_2 = U_1$. The optimal policy is: exclusive use of control 1 (i.e. $\hat{y} = \infty$) if $\lambda_1 > \lambda_2$ or exclusive use of control 2 (i.e. y = 0) if $\lambda_2 > \lambda_1$. If $\lambda_1 = \lambda_2$, the agent is indifferent between the two controls at all cash reserve levels.
- iv) The value function satisfies all of the properties that hold for the general case (Theorem 3.1).

Note that, using Theorem 3.3(iv), it can be shown that all of the subsequent analysis, which is proved for the convex case, will also hold for the binary case.

4 Optimal principal-agent contracts

4.1 The principal's problem

Given the agent's best response and the individual rationality constraint, the principal picks a contract triple (w, k, y) to maximize net receipts. The specific interpretation of the "score", Y(t), will determine the exact form of the principal's return function. We discuss this issue now.

Consider the cash reserve interpretation of the agent's score. Recall that in this interpretation, the specified rate of return k is an actual outflow. The principal pays the agent compensation w and receives a dividend k-w every period. Hence, the cumulative index Y(t) is a cumulative cash reserve and when it runs down to zero the agent is dismissed and the principal hires a second agent. So in this interpretation, the principal incurs in every period an interest payment on the initial cash y. Let us suppose the principal's discount rate is also δ (this is unnecessary for the analysis in the current section but will be required in Section 5). The principal's net receipts, denoted $H(w, k, y; \pi)$, when the agent follows a strategy π and so do successive agents, is

$$H(w,k,y;\pi) = E\delta \int_0^{T_\pi} e^{-\delta s} [k-w-\delta y] ds - Ee^{-\delta T_\pi} [\delta y - H(w,k,y;\pi)].$$

Note that there is, of course, no loss in generality in restricting successive agents to the same best response strategy. Collecting terms,

$$H = k - w - \frac{\delta y}{1 - Ee^{-\delta T}},\tag{4.1}$$

where

$$\frac{1 - Ee^{-\delta T}}{\delta} = E \int_0^{T_x} e^{-\delta s}$$

is the expected (discounted) time to failure by the agent. No matter which generation

of agent is currently employed, the principal always gets per period returns of k-w. However, the expected sum of discounted initial cash outlays, $\delta y/1 - Ee^{-\delta T}$, depends on the agent's best respone. The principal's problem is:

$$\max_{(w,k,y)\in\mathbb{R}^3_+} k - w - \frac{\delta y}{1 - Ee^{-\delta T^*(y)}}$$
s.t.
$$V(y; w, k) \ge \hat{U} \ge 0,$$

where

$$T^*(y) = \min \{t: Y(t) = 0; Y(0) = y, \pi = \beta^*(k, w)\}.$$

A second interpretation of the score is one in which the rate of return k is not an actual outflow but is simply used to track the agent's performance against this standard. In this interpretation, all of the incremental returns, dY + k accrues to the principal out of which he pays the agent the constant compensation w. The conclusions under these two interpretations are similar (indeed as $\delta \to 0$, the two returns converge to the same limit) and so in this paper we pursue only the cash reserve interpretation.¹⁰

4.2 Individual rationality

The requirement that the agent be able to make at least the reservation utility \hat{U} restricts the set of feasible contracts. Recall that $\bar{U}(w)$ is defined as the highest instantaneous utility when the prevailing wage is w. Define the minimum wage \underline{w} as:

$$\bar{U}(w) = \hat{U}$$
.

Note that the minimum wage is independent of the rate of return k. We know from Proposition 3.1 that the agent's value in a bankruptcy scheme is bounded above by $\overline{U}(w)$. Clearly, any compensation scheme offered by the principal must pay a wage at least as large as \underline{w} . Further, define $y^*(w, k)$ as the minimum security level for fixed (w, k):

$$V(y^*(w, k); w, k) = \hat{U}, \quad w \ge \underline{w}.$$

$$H = E\delta \int_0^T e^{-\delta s} [dY + (k - w)ds] + Ee^{-\delta T}H.$$
 (i)

Using Ito's lemma, (i) can be rewritten as

$$H = E\delta \int_0^T e^{-\delta s} [\delta Y(s) + k - w] ds - \delta y + Ee^{-\delta T} H$$

$$= k - w - \frac{\delta y}{1 - Ee^{-\delta T}} + \frac{E\delta \int_0^T e^{-\delta s} \delta Y(s) ds}{1 - Ee^{-\delta T}}.$$
(ii)

The difference in the net receipts (4.1) and (ii), is the last term in (ii) which reflects the fact that in the "score" approach the principal gets, on average, the dividend k - w plus the excess profits that would go into the cash reserve.

¹⁰ In the second formulation

Given Theorem 3.1 and Proposition 3.1, $y^*(w, k)$ is well-defined, and indeed is decreasing in w and increasing in k. The set of feasible compensation schemes then is

$$B = \{ w, k, y \in \mathbb{R}^3 : w \in [w, k], y \ge y^*(w, k) \}.$$

4.3 Initial cash level choice

For fixed (w, k), the principal picks an initial cash level $y \ge y^*$, to minimize the expected discounted per period setup costs $\delta y/1 - Ee^{-\delta T}$. The initial cash level (or equivalently, the tolerable loss level) is one mechanism by which the principal transfers risk to the agent. Note the a priori conflicts: a lower initial cash level implies smaller interest payments for the principal but also quicker failure on the part of the agent and hence a more frequent outlay of initial capital by the principal.

Proposition 4.1. For fixed (w, k), the optimal loss level choice is $y^*(w, k)$.

Proof. From Theorem 3.1, $V'(y) = V'(0)Ee^{-\delta T^*(y)}$. Since V is C^2 , it follows that $V''(y) = V'(0)\frac{d}{dy}Ee^{-\delta T^*(y)}$. Dutta [6] (Theorem 3.1) shows that V'' increases in y. It then follows that $Ee^{-\delta T^*(y)}$ is a decreasing, convex function, or equivalently that $1 - Ee^{-\delta T^*(y)}$ is an increasing, concave function. Hence $\delta y/1 - Ee^{-\delta T^*(y)}$ is minimized over $[y^*, \infty)$, at y^* . \square

It is somewhat surprising to find that in the model under study, under reasonable general conditions, the principal finds it optimal to transfer all the risk that can be feasibly transferred through the choice of initial cash level. The result may not hold when there are costs to new hires, e.g., when there are training costs for new agents. However, it is still the case that an optimal level of initial cash will exist in general. This is so since the principal's returns tend to $-\infty$, as $y \uparrow \infty$.

4.4 Compensation level choice

Given the results of the previous subsection, the optimal choice of a compensation level involves the maximization of

$$(k-w)-\frac{\delta y^*}{1-Ee^{-\delta T}}.$$

A lower compensation w increases the net dividend to the principal, k-w (for k fixed). Since the agent's value increases in w (Proposition 3.1), in order to guarantee the agent expected utility \hat{U} , the tolerable loss level y^* has to increase, thereby raising interest payments for the principal. This is the direct cost of lowering w. There is a further indirect cost, in that the agent's best response is affected, and he may be moved to take actions which lead to (stochastically) more frequent failure.

Proposition 4.2. For fixed k, there is an optimal choice of compensation level w^* , with $w < w^* \le k$.

Proof. The optimization problem is: Minimize

$$w + \frac{\delta y^*(w)}{1 - Ee^{-\delta T} w^{(y)}},$$

over w in $[\underline{w}, k]$, where we write $T_w(y)$ to denote the termination date when an agent uses his optimal response for wage w and an initial cash level y. It is easy to see that as w is lowered to $\underline{w}, y^*(w) \to \infty$. Since $1 - Ee^{-\delta T}w^{(y^*)}$ is bounded between 0 and 1, the minimand goes to ∞ , as $w \downarrow \underline{w}$. The agent's value function V is a continuous function of w and hence so is y^* (and $1 - Ee^{-\delta T}$). So a minimum is achieved over $(\underline{w}, k]$. \square

Lowering the agent's compensation increases the principal's dividend linearly but also increases the expected debt and the latter increases "infinitely" fast as the compensation is lowered to the minimum wage. The result, that wages are strictly higher than minimum wage, looks like an efficiency wage conclusion although the explanation here is a combination of incentive and individual rationality arguments and therefore different from the standard purely incentive-based argument.

4.5 Rate of return choice and the optimal contract

The final component of the principal's choice problem is to pick a required average rate of return k to maximize

$$k - w^* - \frac{\delta y^*}{1 - Ee^{-\delta T}}.$$

The incentive and individual rationality considerations are similar to those involved in the choice of w. An increase in k makes the agent more receptive to tenure considerations (which the principal prefers). On the other hand the (binding) individual rationality constraint implies that the allowable loss has to be larger. It is our conjecture that the optimal choice of k lies between the highest and lowest drifts. We have however not been able to prove this. As is clear from the principal's objective function, a determination of the optimal k involves both the expected discounted time to failure, through $\delta/[1-Ee^{-\delta T}(y^*(k))]$, and the minimum initial cash reserve, $y^*(k)$. The first function in particular depends in an extremely complicated way on k, since its evaluation involves the entire optimal policy function $\beta^*(y, k)$. Furthermore, unlike the choice of w, the initial cash level does not become unbounded as k becomes larger than the highest drift. It appears therefore that this question may not have a resolution in the full generality of our framework.

5 First-best analysis

The first-best or Pareto-optimality problem is one of maximizing the principal's net receipts subject to the individual raionality constraint, but in the absence of moral hazard. There are consequently two differences: firstly, since the agent's actions are observable (and agents are identical) there is no need for dismissal as an incentive device. Secondly, actions are taken so as to maximize a (weighted) sum of principal and agent utilities. The principal result of this section shows that if the agent's utility is separable in compensation and action and the agent is averse to risk, then the Pareto-optimal policy is to choose always the control that maximizes an appropriate

*

weighted sum of instantaneous returns. Formally, the first best problem is,

Maximize
$$E\delta \int_0^\infty e^{-\delta s} \{ [dY + kds] - w(s)ds \}$$

s.t. $E\delta \int_0^\infty e^{-\delta s} U(w(s), (m(s), v(s)))ds \ge \hat{U},$

where the compensation scheme $[w(t): t \ge 0]$ is some \mathcal{F}_t -adapted process. In the remaining sections we shall strengthen our assumptions on the agent's instantaneous utility.

(A4) U is separable in compensation and action; U(w, m, v) = u(w) + q(m, v). Further u is increasing and strictly concave, and $\lim_{w \to \infty} u'(w) = 0$.

Recall that the principal and agent's discount rates are the same. Given the agent's risk-aversion and identical discount rates, standard arguments show that the principal should completely insure the agent in the first-best strategy.

Proposition 5.1. In the first-best solution, the agent's compensation is completely independent of outcomes, i.e. $w(\omega, t) = w$, for all (ω, t) in $\Omega \times [0, \infty)$.

Proof. See Appendix.

Define the weight first-best problem as

$$\max(1-\lambda)\bigg\{E\delta\int_0^\infty e^{-\delta s}[dY+kds]-w\bigg\}+\lambda\bigg\{E\delta\int_0^\infty e^{-\delta s}q(m(s),v(s))ds+u(w)\bigg\},$$
(5.1)

where λ is in [0, 1].

It is well-known that the principal-agent values generated by the weighted first-best problem as λ varies, are exactly the Pareto optimal values. Further we have

Proposition 5.2. For any λ [0, 1), a solution to (5.1) is a strategy using control (\tilde{m}, \tilde{v}) exclusively, and a compensation $\tilde{w}(\lambda)$ where,

- i) $(\tilde{m}, \tilde{v}) \in argmax[(1 \lambda)m + \lambda q(m, v)],$
- ii) $\tilde{w}(\lambda)$ is the (unique) maximizer of $\lambda u(w) (1 \lambda)w, w \ge 0$.

Proof. See Appendix.

The reason that constant use of a single control is optimal is clear. The principal, in the formulation of the first-best problem above, is assumed to have an infinite pocket. Hence principal (and agent) at every instant face an infinite horizon problem which is invariant over the cumulative profits to date. So myopic optimization, i.e., maximization of (weighted) one-period utilities, is dynamically optimal. With firm bankruptcy possible, the simple results here would no longer hold; 11 note however

¹¹ Note that the infinite-pocket assumption is also implicit in the moral hazard formulation of Section 4 and hence in order to compare moral hazard and first-best values, as we shall do shortly, we need to maintain this consistency in assumption.

that in an optimal solution the probability of the principal and agent accumulating negative infinite wealth is zero.

6 Convergence-to first-best utilities

The following result bounds the rate at which principal-agent values under bankruptcy contracts approach first-best efficiency as the discount rate approaches zero. The analysis (whose details may be found in the Appendix) proceeds as follows: denote the (discount independent) first-best action (respectively, agent compensation) as \tilde{m} , \tilde{v} (respectively, \tilde{w}). Consider a bankruptcy contract that requires the agent to maintain a rate of return $k = \tilde{m}$, pays him \tilde{w} if he does so and allows a discount dependent initial cash level $y(\delta)$. This initial cash level has to be chosen as an appropriate balance between reducing the risk of firing a "good" agent and maintaining incentives. In the Appendix we show that one choice of initial cash level that achieves this balance is $\frac{-n \ln \delta}{\sqrt{\delta}}$, where n is an appropriate constant. Clearly

the rate of asymptotic efficiency of this particular bankruptcy contract is a lower bound for the rate at which second-best values approach efficiency. It is an open question as to how tight these bounds are.

Proposition 6.1 For any first-best values (G, H), there exist principal-agent contracts $(w(\delta), k(\delta), y(\delta))$ such that

$$1 - \frac{V_{\delta}[y(\delta); w(\delta), k(\delta)]}{G} = O(\delta^{1/2} \ln \delta)$$

$$1 - \frac{H_{\delta}[y(\delta); w(\delta), k(\delta)]}{H} = O(\delta^{1/2} \ln \delta)$$

Proof. See Appendix.

Remark. The arguments in the proof of Proposition 6.1 are completely independent of the particular Pareto-optimal point that is being approximated. So, the result is really a statement on the rate of uniform convergence of the principal-agent value frontier to the Pareto optimal first-best frontier.

We next report exact computations on the rate of convergence for a parametric example. Suppose that

$$U(w, m, v) = \frac{25}{16} \frac{w}{w+1} + z[v(1-m)]^{1/4},$$

where z > 0, and suppose further that the set of agent actions is

$$A = \left\{ (m, v) \in \mathbb{R}^2_+ : m + v \le 1, \frac{z^2}{8} \le v \right\}.$$

It is straightforward to check that all of the assumptions made above are indeed satisfied by U and A. Set the principal and agent weights in the first-best exercise

to be equal; i.e. $\lambda = 1/2$. It is then possible to show (see the Appendix for details) that

$$\begin{split} &1 - \frac{V_{\delta}[y(\delta); w(\delta), k(\delta)]}{G} \leq \sqrt{\delta} \\ &1 - \frac{H_{\delta}[y(\delta); w(\delta), k(\delta)]}{H} \leq - \left[\frac{1}{4}f(z)z\right] \sqrt{\delta} \ln \delta \end{split}$$

where

*

$$f(z) = \frac{25 + 8z^2}{(3 - z^2)(5 + 8z^2)}.$$

For example, when z=1/16, f(z)<2 and hence that rate of convergence is bounded by $1/8\sqrt{\delta}\ln\delta$. This particular parameter value is represented in Figures 1-2. Figure 1 plots the efficiency loss associated with discount rates between 1%

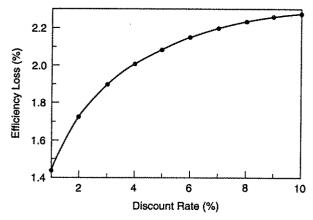


Figure 1. Efficiency loss: moderate discount rates, z = 1/16.

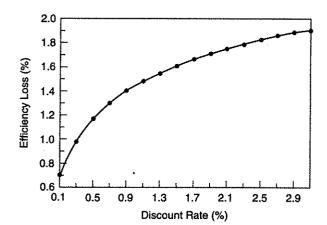


Figure 2. Efficiency loss: very low discount rates, z = 1/16.

and 10%. Note that even at a discount rate of 10%, the size of efficiency loss is less than 2.3%. Figure 2 reports the efficiency loss for this parameter specification between discount rates of 0.1% and 3.1%. It can be seen that the size of efficiency loss is less than 1% for discount rates of 0.3% and less.

7 Discussion and extensions

Second-best contracts have been characterized by [16], [21], [12] and [25]. The work of Rogerson [21] and Lambert [16] has shown that second-best incentive schemes will, in general, have "memory"; compensations in any period will depend in a subtle manner on previous compensations and/or outcomes. The theoretical reason for this is the fact that, although optimal contracts will depend in the expected manner on the information revealed by observed outcomes, this information can be linked quite arbitrarily to the outcomes themselves. Spear–Srivastava [25] provide some reduction in the dimension of contingent variables. They show that the second-best scheme conditions on current output and the agent's expected continuation value. Unfortunately this last statistic is not easy to relate to any aggregate of outcomes. As mentioned in the introduction, these results are difficult to reconcile with the simplicity of observed contracts.

Holmstrom-Milgrom [12] show in a T-period diffusion model that if principal and agent utilities are multiplicatively separable and exponential then the second-best contract has the attractive feature of being a succession of short-term contracts (and indeed is linearly related to observed outcomes). In a two-period model, Fellingham-Newman-Suh [8] isolated a couple of other configurations of principal and agent preferences for which the same conclusion holds. Unfortunately, the linear short-term characterization is very delicately predicted on the constant absolute-risk-aversion specification of preferences.

As described in the introduction, a line of research has indeed looked at some simple schemes, and shown that any single-period efficient utility level can be attained arbitrarily closely by such schemes in the limit (Radner [18], [19], [20], Rubinstein [22] and Rubinstein—Yaari [23]). This is the literature that motivated us directly. In particular, we have tried in this paper to complement the findings of this line of inquiry by providing a direct analysis of the optimal principal-agent contracts (within a class of simple schemes) and by providing an estimate of the rate of approach to efficiency. Note also that Fudenberg—Maskin [9] employ ideas used in the oligopoly context by Abreu—Pearce—Stachetti [1] to study the entire set of sustainable payoffs in settings of imperfect information more general than the repeated moral hazard problem. They establish the asymptotic sustainability of all individually rational payoffs (and hence first-best payoffs) under some conditions.

In an interesting paper, Fudenberg-Holmstrom-Milgrom [10] argue the general point that if the agent is allowed to insure himself, then some of the insurance that the principal has to provide in standard contracts without this feature becomes unnecessary. In particular they show that when the preferences of principal and agent are additively separable and of the constant absolute risk aversion class, then long-term contracts can be replaced with a succession of second-best short-term contracts. In this context it is worth noting that Yaari [26] has shown that, for

500 P.K. Dutta and R. Radner

some specifications of bankruptcy, a patient risk-averse agent subject to income fluctuations finds it optimal to consume every period the expected income; i.e. to behave as a risk-neutral agent. This suggests the conjecture that if the agent is allowed to self-insure and is made the residual claimant in our model, then the resulting outcomes would again approximate the first-best, provided principal and agent are sufficiently patient.¹²

In a recent paper, Brock-Evans [4] have pointed out that in our model the principal can gather additional information on the agent's variance choice. They discuss results from statistical estimation theory which have the following implication: if the principal samples the (stochastic) cash-reserve process sufficiently frequently over a time-interval [t, t+g], and if the agent's actions over that time-interval change sufficiently slowly, then the principal can get precise estimates of the true variance choices made by the agent. This additional information could be utilized by the principal in the design of optimal contracts. Furthermore, it may be possible to also utilize this additional information to improve the rate of convergence to Pareto efficiency.

Two possible generalizations of the model can be attempted. First one can study compensation schemes in which the agent's compensation is linked directly to immediate performance (as well as indirectly through the possibility of being fired), i.e., salary plus bonus schemes. Secondly, as discussed above, the agent can be allowed to insure himself, allowing for the smoothing of consumption across periods even when income is erratic. In a model incorporating these features, it is thus far possible to derive some general results. ¹³ but not enough to allow an explicit characterization of the optimal contract choice.

References

1

- Abreu, D., Pearce, D., Stachetti, E.: Optimal cartel equilibria with imperfect monitoring. J. Econ. Theory 39, 251-269 (1986)
- 2. Berge, C.: Topological spaces. New York: MacMillan 1963
- 3. Benes, V.: Existence of optimal stochastic control laws. SIAM J. Control 9, 446-472 (1971)
- Brock, W.A., Evans, L.T.: The space of contracts in continuous time asymmetric information models. University of Wisconsin, Madison, unpublished (1992)
- Davis, M.H.A.: On the existence of optimal policies in stochastic control. SIAM J. Control 11, 587-594 (1973)
- Dutta, P.K.: Bankruptcy and expected utility maximization. J. Econ. Dynam. Control, forthcoming (1993)

¹² There are two important differences between our formulation and Yaari's on account of which we cannot immediately infer that the Yaari result holds. First, Yaari allows only a consumption choice for the agent; stochastic returns are generated every period with mean μ and the agent picks a consumption rate c, thereby determining an effective mean of m. In our framework, the agent picks from a given menu a particular project or effort level (corresponding to a mean μ) and additionally picks consumption c (and therefore a net mean m). Second, Yaari formulates bankruptcy in his (finite-horizon) model in a way that has no analog in the infinite horizon problem.

¹³ Dutta, unpublished notes. The special class of performance standards in which the principal only tracks the number of successes and failures (rather than the amounts by which the agent succeeded or failed) can however be easily accommodated into our analysis, as is easy to check.

- Dutta, P.K., Radner, R.: Optimal principal agent contracts for a class of incentive schemes: a characterization and the rate of approach to efficiency. University of Rochester Working Paper #300, 1991
- Fellingham, J., Newman, D., Suh, Y.: Contracts without memory in multiperiod agency models. J. Econ. Theory 37, 340-355 (1985)
- 9. Fudenberg, D., Maskin, E.: Discounted repeated games with unobservable actions, I: one sided moral hazard. Harvard University Working Paper, #1280, 1986
- Fudenberg, D., Holmstrom, B., Milgrom, P.: Short term contracts and long term relationships. J. Econ. Theory 51, 1-31 (1990)
- Heath, D., Orey, S., Prestien, V., Sudderth, W.: Maximizing or minimizing the expected time to reach zero. SIAM J. Control Optimiz. 25, 195-205 (1987)
- Holmstrom, B., Milgrom, P.: Aggregation and linearity in the provision of intertemporal incentives. Econometrica 55, 303-329 (1987)
- 13. Kaplan, R.S.: Advanced management accounting. New Jersey: Prentice-Hall 1982
- 14. Karatzas, I., Shreve, S.: Brownian motion and stochastic calculus. New York: Springer 1987
- 15. Krylov, N.: Controlled diffusion processes. New York: Springer 1981
- 16. Lambert, R.A.: Long-term contracts and moral hazard. Bell J. Economics 14, 441-452 (1983)
- Majumdar, M., Radner, R.: Linear models of economic survival under uncertainty. Econ. Theory 1, 13-31 (1990)
- Radner, R.: Monitoring cooperative agreements in a repeated principal-agent relationship. Econometrica 49, 1127-1148 (1981)
- 19. Radner, R.: Repeated principal-agent games with discounting. Econometrica 53, 1173-1198 (1985)
- Radner, R.: Repeated moral hazard with low discount rates. In: Heller, W., Starr, R., Starrett, D. (eds.) Essays in honor of Kenneth Arrow. Cambridge: Cambridge University Press 1986
- 21. Rogerson, W.: Repeated moral hazard. Econometrica 53, 69-76 (1985)
- Rubinstein, A.: Offenses that may have been committed by accident an optimal policy of retribution. In: Brams, S.J., Schotter, A.S., Schrodiauer, G.S. (eds). Applied game theory. Würzburg: Physica 1979
- Rubinstein, A., Yaari, M.: Repeated insurance contracts and moral hazard. J. Econ. Theory 30, 74-97 (1983)
- Sheng, D.: Two-mode control of absorbing Brownian motion. Bell Laboratories Technical Memorandum, Murray Hill, NJ, 1980
- 25. Spear, S., Srivastava, S.: On repeated moral hazard with discounting. Rev. Econ. Studies 55 (1988)
- Yaari, M.: A law of large numbers in the theory of consumption choice under uncertainty. J. Econ. Theory 12, 202-217 (1976)

Appendix

In this Appendix we prove the results of Sections 5 and 6.

Proof of Proposition 5.1. Consider any sample path with agent compensation $w(\omega, \cdot)$. Write $\bar{w}(\omega) = \delta \int_0^\infty e^{-\delta s} w(\omega, s) ds$, the "mean wage" for the measure induced by the discount rate $\delta e^{-\delta s}$. By Jensen's inequality,

$$u[\bar{w}(\omega)] > \delta \int_{0}^{\infty} e^{-\delta s} u[w(\omega, s)] ds.$$

Denote $\bar{w} = E\bar{w}(\omega)$, taking the expectation now with respect to the measure induced by the given control strategy. Since u is concave, again, by Jensen's inequality,

$$U(\bar{w}) \geq Eu[\bar{w}(\omega)] \geq E\delta \int_0^\infty e^{-\delta s} u[w(\omega, s)] ds.$$

Since the principal discounts the future at the same rate as the agent and is risk-neutral, along any sample path the principal's returns are identical under time varying compensation $w(\omega, \cdot)$ or mean wage $\bar{w}(\omega)$. From risk-neutrality it follows that the principal is indifferent between environment-varying compensation $\bar{w}(\omega)$ or a constant compensation \bar{w} .

Proof of Proposition 5.2. The weighted first-best maximand is

$$(1-\lambda)\left\{E\delta\int_0^\infty e^{-\delta s}[dY+kds]-w\right\}+\lambda\left\{\delta\int_0^\infty q(s)ds+u(w)\right\} \tag{A.1}$$

and a strategy is of course the choice of instantaneous controls (m, v) for every time instant and environment, and a constant compensation level w. From (A.1) it is clear that the two choices can be made independently. Further, the maximand for w is

$$\lambda u(w) - (1 - \lambda)w$$
 and this is clearly maximized for w s.t. $u'(\tilde{w}) = \frac{1 - \lambda}{\lambda}$ when $\tilde{w} > 0$ or

at $\tilde{w} = 0$, e.g., when $\lambda = 0$. It is further clear from (A.1), that $E\delta\{\int_0^\infty e^{-\delta s} [dY(1-\lambda) + q(m,v)\lambda]ds\}$ is maximized by the constant use of the control which maximizes $m(1-\lambda) + q(m,v)\lambda$.

Proof of Proposition 6.1. One way in which one could estimate the rate of approach to efficiency would be to directly analyze the asymptotic behavior of the optimal bankruptcy scheme $(w^*(\delta), k^*(\delta), y^*(\delta))$, as $\delta \downarrow 0$. Since explicit expressions for these parameters cannot be obtained, such a direct line of attack is not very fruitful. Instead, we concentrate on finding a particular set of schemes $(w(\delta), k(\delta), y(\delta))$ for which the rate of convergence stated in Proposition 6.1 is valid. Clearly, such a rate is therefore a lower bound for the rate implied by $(w^*(\delta), k^*(\delta), y^*(\delta))$, which in turn is a lower bound for the general class of all admissible compensation schemes.

Suppose the first-best constant control is (\tilde{m}, \tilde{v}) and the associated wage is \tilde{w} . Consider, $k(\delta) = \tilde{m}$, $w(\delta) = \tilde{w}$, for all δ . We will specify $y(\delta)$ shortly. For any $y \ge 0$, let $\beta^*(\tilde{m}, \tilde{w}, \delta)$ denote the stationary Markovian optimal best response policy of the agent. Define

$$\begin{split} \tilde{U} &\equiv U(\tilde{w}, \tilde{m}, \tilde{v}) \\ T_{\delta}(y) &= \min \left\{ t > 0 \colon Y(t) = 0 \mid Y(0) = y, \ \beta^*(\tilde{m}, \tilde{w}, \delta) \right\} \\ \tilde{T} &= \min \left\{ t > 0 \colon Y(t) = 0 \mid Y(0) = y, \ \pi \equiv (\tilde{m}, \tilde{v}) \right\}. \end{split}$$

Clearly,

$$V_{\delta}(y) \le \bar{U}(\tilde{w})(1 - E^{-\delta T_{\delta}(y)}) \tag{A.2}$$

$$V_{\delta}(y) \ge \tilde{U}(1 - E^{-\delta \tilde{T}(y)}) \tag{A.3}$$

(A.2) and (A.3) imply that

$$1 - Ee^{-\delta T_{\delta}(y)} \ge \frac{\tilde{U}}{\bar{U}(\tilde{w})} (1 - Ee^{-\delta \tilde{T}(y)})$$
$$\equiv b(1 - Ee^{-\delta \tilde{T}(y)}).$$

Now, by standard arguments (e.g. Dutta [6]) it follows that

$$1 - Ee^{-\delta \tilde{T}(y)} = 1 - e^{-\tilde{\lambda}y} \tag{A.4}$$

where

1

$$\tilde{\lambda} = \biggl(\sqrt{\frac{2}{\tilde{v}}}\biggr)\sqrt{\delta} \equiv a\sqrt{\delta}.$$

From (A.4) it follows that

$$H_{\delta}(y) \ge (\tilde{m} - \tilde{w}) - \frac{\delta y(\delta)}{b(1 - e^{-\tilde{\lambda}y(\delta)})}.$$
 (A.5)

Collecting (A.4) and (A.5) together we have

$$\frac{H_{\delta}(y)}{\tilde{m} - \tilde{w}} \ge 1 - c \frac{\delta y(\delta)}{1 - e^{-\tilde{\lambda}y(\delta)}}$$

$$\frac{V_{\delta}(y)}{\tilde{U}} \ge 1 - e^{-\tilde{\lambda}y(\delta)}$$

where $c^{-1} \equiv (\tilde{m} - \tilde{w})b$.

The remainder of the proof will be as follows. We shall demonstrate the existence of $y(\delta)$ such that i) $\delta y(\delta) = O(\sqrt{\delta} \ln \delta)$, ii) $e^{-\frac{\tilde{\lambda}}{\tilde{\lambda}}y(\delta)} = O(\sqrt{\delta} \ln \delta)$. Clearly, the proof will then be complete.

For any n > 0 define

$$y(\delta) \equiv \frac{-n\ln\delta}{\sqrt{\delta}}.$$

It follows that $\delta y(\delta) = -n\sqrt{\delta} \ln \delta = O(\sqrt{\delta} \ln \delta)$. Further, $-\tilde{\lambda}y(\delta) = an \ln \delta$ and so $e^{-\lambda y(\delta)} = \delta^{an}$. If n is chosen such that $an \ge \frac{1}{2}$, then δ^{an} goes to zero faster than $\sqrt{\delta} \ln \delta$. Hence, for $n \ge \frac{1}{2a}$, $e^{-\tilde{\lambda}y(\delta)} = O(\sqrt{\delta} \ln \delta)$, and the proof of Proposition 6.1 is complete.

Computations for the parametric example

Since the agent's utility is increasing in the variance, it is immediate that for any mean, the largest associated variance is chosen in a first-best solution: i.e. v = 1 - m. The first-best \tilde{m} then maximizes $m + z(1 - m)^{1/2}$. It follows that $\tilde{m} = 1 - z^2/4 = 1 - \bar{v}$. Furthermore, it is easy to check that the first-best wage $\tilde{w} = 1/4$. In the terminalogy of the proof of Proposition 6.1, $a = \left(\frac{2}{z^2}\right)^{1/2}$ and consequently, $n \ge 1/2\left(\frac{2}{z^2}\right)^{-1/2}$. Finally, substituting for \tilde{m} , \tilde{v} and \tilde{w} , it follows that $c = (25 + 8z^2)/(3 - z^2)(5 + 8z^2)$. The computations are complete. \square