



Aggregation and Linearity in the Provision of Intertemporal Incentives

Bengt Holmstrom; Paul Milgrom

Econometrica, Vol. 55, No. 2 (Mar., 1987), 303-328.

Stable URL:

<http://links.jstor.org/sici?sici=0012-9682%28198703%2955%3A2%3C303%3AAALITP%3E2.0.CO%3B2-K>

Econometrica is currently published by The Econometric Society.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/econosoc.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

AGGREGATION AND LINEARITY IN THE PROVISION OF INTERTEMPORAL INCENTIVES

BY BENGT HOLMSTROM AND PAUL MILGROM¹

We consider the problem of providing incentives over time for an agent with constant absolute risk aversion. The optimal compensation scheme is found to be a linear function of a vector of N accounts which count the number of times that each of the N kinds of observable events occurs. The number N is independent of the number of time periods, so the accounts may entail substantial aggregation.

In a continuous time version of the problem, the agent controls the drift rate of a vector of accounts that is subject to frequent, small random fluctuations. The solution is as if the problem were the static one in which the agent controls only the mean of a multivariate normal distribution and the principal is constrained to use a linear compensation rule. If the principal can observe only coarser linear aggregates, such as revenues, costs, or profits, the optimal compensation scheme is then a linear function of those aggregates. The combination of exponential utility, normal distributions, and linear compensation schemes makes computations and comparative statics easy to do, as we illustrate.

We interpret our linearity results as deriving in part from the richness of the agent's strategy space, which makes it possible for the agent to undermine and exploit complicated, nonlinear functions of the accounting aggregates.

KEYWORDS: Principal-agent problems, moral hazard, incentives, piece-rates, Brownian motion, aggregation, linear incentive schemes.

"Then I thought a minute, and says to myself, hold on,—s'pose you'd a done right and give Jim up; would you felt better than what you do now? No, says I, I'd feel bad—I'd feel just the same way I do now. Well, then, says I, what's the use you learning to do right when it's troublesome to do right and ain't no trouble to do wrong, and the wages is just the same? I was stuck. I couldn't answer that. So I reckoned I wouldn't bother no more about it, but after this always do whichever come handiest at the time." *Huckleberry Finn*

1. INTRODUCTION

THERE IS A LONG TRADITION in economic theory of analyzing the conditions under which market-mediated outcomes are efficient, but only recently has there been a systematic effort to investigate the kinds of nonmarket arrangements that might prevail when those conditions fail. Much of this recent work has centered on the simplest nonmarket paradigm: the principal-agent (or multiagent) relationship. Principal-agent models, in which the agent enjoys some informational advantage over the principal, offer a natural framework for investigating the thesis that informational constraints are economically significant in many trading relationships and should therefore be treated on par with standard resource constraints. These models have been successful in suggesting explanations for nonmarket institutions and contracting practices that had been left unexplained by received micro theory.

Yet, those interested in incorporating the rudimentary agency models into more lively and realistic economic settings have encountered a problem: Optimal incentive contracts tend to be complicated even in the simplest situations, making

¹ This research was supported by NSF Grants IST-8411595 and SES-8411732. For their helpful comments, we thank Jacques Dreze, Darrell Duffie, David Kreps, Michael Harrison, Edward Sherry, two anonymous referees, and the participants in seminars at IMSSS (Stanford), MIT, Hebrew University, Oxford, and CORE (Louvain).

the models hard to work with in extended settings. Indeed, very few restrictions can be placed on the shape of optimal contracts; even such a basic property as monotonicity may be hard to assure (Milgrom, 1981; Grossman and Hart, 1983). This is explained by the fact that outcomes (output, profits, etc.) provide information about the agent's action. Optimal contracts respond directly to the strength of this information, which may vary in a rather arbitrary way with the outcome. For the same reason, when agents are risk averse, optimal contracts will generally depend on all available information about the agent's action (Holmstrom, 1979; Shavell 1979). This admits aggregation through sufficient statistics, but such statistics rarely correspond to common accounting measures, which aggregate information to a substantially greater degree.²

Real world incentive schemes appear to take less extreme forms than the finely tuned rules predicted by the basic theory. One could explain the simplicity of actual incentive schemes partly by the costs of writing intricate contracts, but that is hardly the whole story. Agents in the real world typically face a wider range of alternatives and principals a more diffuse picture of circumstances than is assumed in the usual models. Optimal schemes derived from a spare and approximate model of reality may perform quite poorly in the richer real environment. Indeed, it is our purpose to show that a linear compensation scheme based on aggregates—whose performance is quite robust to the specification of the environment³—emerges as optimal in a class of intertemporal problems for which the agent has great freedom of action.⁴

Our main idea can best be described in the context of a nonexistence result discovered by Mirrlees (1974). Consider a moral hazard problem in which the principal is risk neutral and the strictly risk-averse agent privately (and at a cost) controls the mean of a normal distribution with fixed variance. If the agent's utility is additively (or multiplicatively) separable in consumption and action and unbounded from below in consumption, there is no optimal solution, because the first-best (informationally unconstrained) solution can be approximated arbitrarily closely using a compensation scheme of the following sort: Pay a fixed

² There is a trivial and uninteresting sense in which it is always possible to make the optimal compensation of the agent a linear function of a single numerical "aggregate," namely, his optimal compensation rule can be expressed as a linear function of any multiple of that rule, and the latter is a single numerical aggregate. When we say that compensation can be based on aggregates, we have something more sensible in mind. First, the aggregate must be determined as a *linear* function of some separately observed variables, such as the profits earned in two different periods of time. Second, its definition must not depend on such parameters of the problem as the risk aversion of the principal or agent or the costs of various actions the agent may take.

Idealized accounting systems are based on aggregates of this sort. Account balances are accumulated sums over time, and they are defined in a way that depends only to a limited degree on the tastes of the owners and managers.

³ Bull, Schotter, and Weigelt (1985) have found in experimental studies that piece-rate compensation schemes are more consistent than tournament schemes for eliciting productive behavior.

⁴ Recently, Laffont and Tirole (1985) and McAfee and McMillan (1986) have proposed another explanation of linear contracts. In their analyses, the principal and agent are both risk neutral and therefore indifferent across a range of incentive contracts which sometimes includes linear contracts. That idea is complementary to our contention that linear contracts enjoy a robustness that makes them effective in a wide range of situations.

wage unless output is very low, and pay a very low wage for very low output. This "two-wage" scheme is effective because normal distributions have the property that very low outputs are much more likely when the agent shirks than when he does not. Consequently, the two-wage scheme can be designed to impose virtually no risk on an agent who follows instructions but a large risk on an agent who shirks.

This nonexistence result and the near-optimal rules that go with it are disturbing. In practice, one feels, schemes that adjust compensation only when rare events occur are not likely to provide correct incentives for the agent in ordinary, high probability circumstances. Moreover, to construct the scheme, the principal requires very precise knowledge about the agent's preferences and beliefs, and about the technology that he controls. The two-wage scheme performs ideally if the model's assumptions are precisely met, but can be made to perform quite poorly if small deviations in the assumptions about the variance or (especially) about the agent's ability to control the probability of rare events are introduced. Later, we shall consider a variant of the Mirrlees model in which the agent chooses his labor input over time in response to observations of how well he is doing. In our model, a two-wage scheme paid as a function of profits over some period of time leads the agent to work hard only when that appears necessary to avoid a disaster. Intuitively, a linear scheme, because (apart from wealth effects) it applies the same incentive pressure on the agent no matter what his past performance has been, would lead to a more appropriate choice of effort over time.

We show in this paper that this intuition can be made precise. We consider a continuous time model in which the agent controls the drift rate of a Brownian motion over the unit time interval. Assuming that the agent has an exponential utility function (to abstract from wealth effects) and that the cost of control is monetary, the optimal incentive scheme will indeed be linear in output. The agent will choose a constant drift rate independently of the path of output. This means that the optimal incentive scheme can be computed *as if* the agent were choosing the mean of a normal distribution only once and the principal were restricted *a priori* to using a linear rule. Thus, the dynamic perspective leads not only to a natural resolution of the Mirrlees nonexistence problem, but also to a remarkably simple and easily computed optimal incentive scheme. Moreover, the resulting linear rule is quite robust; the agent's behavior and both parties' payoffs hardly vary when one makes small changes in the specifications of the problem.

We turn to a brief outline of the analysis we will pursue. Section 2 presents a single-period model of moral hazard with the agent controlling a multinomial distribution. The purpose of this section is two-fold: to isolate the key implications of the agent's exponential utility function and to illustrate how "richness" in the agent's opportunity set limits the principal's choice of incentive schemes.

In Section 3 we look at a finitely repeated version of the single-period multinomial model. The main result is that even though the principal can observe the outcomes of each period separately, the optimal scheme depends only on aggregates, and it depends linearly on them. The aggregates are enumeration aggregates,

that is, they count the numbers of times that each particular outcome occurs during the pay period. For instance, in a three-period problem it is enough to know that the agent made \$200 twice and \$100 once; the principal gains nothing from further information about the precise sequence of these events. However, it is not generally the case that an optimal scheme is a function of profits only, nor is it true that if the scheme were restricted to be a function of profits only, the constrained optimal scheme would be a *linear* function of profits.

In Section 4 we go on to study a continuous time approximation of the discrete time model in which optimal compensation schemes can be shown to be linear functions of more comprehensive aggregates than the simple enumeration accounts of Section 3. In the model, the agent controls the instantaneous drift (but not the covariance) of a multi-dimensional Brownian motion. This model serves as a good approximation to any repeated multinomial model satisfying three conditions: (i) the multinomial model involves a large number of periods, (ii) the costs and profits in each individual period are small relative to the principal's and the agent's risk tolerances, and (iii) the number of outcomes that the principal can distinguish in a single period is smaller than the dimension of the agent's control set. The optimal scheme for the multidimensional Brownian model is a linear function of the end-of-period levels of the different dimensions of the process (which we interpret as different account balances). In analogy with the discrete time model, the optimal scheme uses only account balances aggregated over time, but it generally requires more information than just the accumulated profits. Only for the one-dimensional case is the scheme generally linear in end-of-period profits.

This may seem to make the one-dimensional Brownian "linearity in profits" result very special. However, it has an important corollary implication for the multidimensional Brownian model, in which the manager controls a more interesting and complex process. If, in the latter model, the compensation paid must be a function of profits alone (perhaps because reliable detailed accounts are unavailable) or if the manager has sufficient discretion in how to account for revenues and expenses, then the optimal compensation scheme will be a linear function of profits. This is a central result, because it explains the use of schemes which are linear in profits even when the agent controls a complex multi-dimensional process.

We devote Section 5 to illustrations of the ease with which one can apply the continuous time model to obtain simple and explicit solutions to various kinds of agency problems. Our examples should be viewed as merely suggestive; we do not attempt any systematic analysis of possible applications.

2. A SINGLE-PERIOD MODEL

We begin by studying a fairly general single-period model of moral hazard.

An agent controls a stochastic technology owned by a principal. Uncertainty is described by a state of nature θ , which can take a finite number of values

$\theta_0, \dots, \theta_N$. The agent's actions affect the probability distribution over these states: $p = (p_0, \dots, p_N)$. With little loss of generality we can view p itself as the agent's action. The feasible set of actions is then a subset P in the N -dimensional probability simplex. P is convex by randomization and compact by assumption.

Each state θ_i is associated with a monetary payoff π_i that belongs to the principal and a publicly observed information outcome x_i . We will write π for the random payoff and refer to it as "profits." The random public information is denoted x and called the "public outcome" or just the "outcome" of the agent's action. Note that if $x_i = x_j$ implies $\pi_i = \pi_j$, then the outcome x contains at least as much information as π (i.e. profits are observed). Notice, too, that it is possible that some of the information states are never observed.

Both the principal's and the agent's preferences over wealth can be described by exponential utility functions. The constant coefficient of absolute risk aversion is R for the principal and r for the agent; hence, the utility function is $v(y) = -\exp(-Ry)$ for the principal and $u(y) = -\exp(-ry)$ for the agent. The case $R = 0$ ($r = 0$) represents a risk neutral principal (agent) with utility function $v(y) = y$ ($u(y) = y$).

The principal rewards the agent based on the observable outcome x using a sharing rule $s(x)$. The agent's final income is $s(x)$ minus the cost of taking the action p . Thus, we assume that the agent's cost of action is a financial expenditure or an opportunity cost incurred by foregoing some other income generating activity. The cost is allowed to be stochastic and we write it $c(p; \theta)$. The cost function is assumed to be continuously differentiable on P .

The principal's problem is to select a sharing rule s and instructions p for the agent under the two standard constraints that (i) the agent can maximize his expected utility by following instructions and (ii) the agent can attain a certain minimum level of expected utility from the contract. We will measure this minimum expected utility level in terms of the agent's certain equivalent w ; thus, his expected utility has to be at least $u(w)$. The *Principal's Problem* can then be formally stated as:

- (1) $\max_{p, s} \sum_i v(\pi_i - s(x_i)) p_i, \quad \text{subject to:}$
- (2) $p \text{ maximizes } \sum_i u(s(x_i) - c(p'; \theta_i)) p'_i \text{ on } P.$
- (3) $\sum_i u(s(x_i) - c(p; \theta_i)) p_i \geq u(w).$

Before going on to analyze this problem it may be helpful to suggest some interpretations that indicate the scope of the formulation.

In the simplest situation and the one most often studied, the agent controls a one-dimensional action variable, usually interpreted as effort. The cost of effort is deterministic and the agent chooses his effort with no more information about the production possibilities than the principal has. Commonly, output is the only observable variable. In our formulation this standard set-up would correspond to letting P be a one-dimensional manifold (a curve) in the (N) -simplex, letting $c(p; \theta_i) = c(p)$ for all θ_i and letting $x_i = \pi_i$ for all i .

The standard model can be enriched by having the agent observe a signal about the production technology or the cost function before choosing his effort but after entering into a binding contract with the principal. Since the agent is assumed to have no private information at the time of contracting, this is not a model of adverse selection in the usual sense; we refer to it as a Hidden Information Model.⁵ In a Hidden Information Model, the agent's action can be thought of as a strategy that maps his observed signal into an effort level. Note that any such strategy, when paired with a prior distribution over signals, induces a distribution over the final outcome. Therefore, we can again think of the agent as choosing among final distributions at the outset, before any information has been revealed. The set of distributional strategies available to the agent is, of course, restricted. For a distribution to be feasible, the agent must be able to induce via some contingent effort strategy. Such restrictions can be incorporated into the feasible set P , implying that the Hidden Information Model is indeed subsumed in our general framework.⁶

Our formulation includes many more complex models than the Hidden Information Model discussed above. The agent could be choosing a sequence of actions over time. Information of relevance for future decisions could be entering along the way. The agent's actions could influence this information stream as well as future costs, payoffs, and opportunity sets. At each stage actions could be multi-dimensional (e.g. include effort choice, production decisions, project selections, etc.). The cost of action could be stochastic and the observable information (x) essentially anything. In short, we could permit rather arbitrary production and information technologies and still have the reduced form map into the conceptually simple structure (1)–(3), where the agent is viewed as choosing at the outset a final distribution over states θ from some feasible set P . Our principal restrictions are that (i) the agent evaluates wealth at a single point in time, after all actions have been taken, (ii) the cost of actions can be expressed in monetary units, (iii) the utility functions are exponential, and (iv) neither party has any relevant private information at the time the contract is signed.

In view of our subsequent interests it is worth noting that even if the agent's action is a one-dimensional variable (e.g. effort), a contingent strategy, arising in the extended settings just discussed, often permits control of p in more than one dimension. For instance, in the Hidden Information Model, the set P that the agent can choose from will generally be of higher dimension than one, even though the action itself is one-dimensional. Thus, a natural way of increasing the dimensionality of P and hence enriching the agent's action space, is to let the agent act on the basis of private information.

We now proceed to analyze the Principal's Problem. As is well known, the existence of a solution to the program (1)–(3) cannot be taken for granted. The

⁵ This descriptive language was suggested by Arrow (1985), who emphasized the distinction between hidden knowledge (about a state variable) and hidden action. For a contract theory such as ours, a more important distinction is between informational asymmetries that arise before the contract is signed and those which arise after.

⁶ The attentive reader will notice that in the Hidden Information Model, the agent's cost function will generally depend nontrivially on θ .

following result gives sufficient conditions for existence; it is a variant of an existence result of Grossman and Hart (1983).

THEOREM 1: *Suppose that P is compact and that $c(\cdot, \theta_i)$ is continuous for all i . Let p' maximize $\sum p_i u(-c(p, \theta_i))$ and suppose that for all p assigning zero probability to some public outcome, $c(p, \theta_i) > c(p', \theta_i) + \max_{j,k} (\pi_j - \pi_k)$ for all i . Then a solution (p^*, s^*) to the Principal's Problem exists and p^* assigns positive probability to every outcome.*

PROOF SKETCH: Let p' be as specified in the Theorem and let $-\gamma$ be the agent's certain equivalent for the random variable $-c(p, \theta)$. Then, since u is exponential, the pair (p', s') , in which $s'(x_i) \equiv w + \gamma$ for all i , satisfies (2) and (3). One can show that the set of pairs (p, s) , such that (2) and (3) are satisfied and the principal's payoff is at least as high as for (p', s') , is compact and excludes any p that assigns zero probability to any outcome. Also, the principal's payoff is continuous in (p, s) on this set. Hence there is some optimum (p^*, s^*) with the required property. *Q.E.D.*

Given any sharing rule s , the agent's problem (2) has a solution p because P is compact and the objective function is continuous. If the sharing rule s and the optimal response p results in an expected utility level with certain equivalent w , we say that s implements p with certain equivalent w . The set of sharing rules that implement p with certain equivalent w is denoted $S(p, w)$. This set may be empty for some (p, w) . Therefore, define $P^0(w) = \{p \mid S(p, w) \text{ is not empty}\}$ and $P^*(w) = \{p \mid \text{for some } s, (s, p) \text{ solves the Principal's Problem}\}$. The key implications of assuming that the agent's utility is exponential can now be stated as follows:

THEOREM 2: *For any s , w and $p \in P^0(w)$: (i) $s \in S(p, w)$ if and only if $s - w \in S(p, 0)$, (ii) $P^0(w) = P^0$ for all w , (iii) $P^*(w) = P^*$ for all w .*

PROOF: Because utility is exponential,

$$(4) \quad \sum_i u(s(x_i) - w - c(p; \theta_i)) p_i = -u(-w) \sum_i u(s(x_i) - c(p; \theta_i)) p_i.$$

Since $-u(-w) > 0$, any p that is best for the agent against $s(x) - w$ is also best against $s(x)$, and conversely. Also, $-u(-w)u(w) = u(0)$. This proves (i) and (ii). Part (iii) then follows from the exponential form of the principal's objective function in (1). *Q.E.D.*

Theorem 2, part (iii) asserts that the optimal choice of an instruction p^* given to the agent does not depend on the required minimum certain equivalent w . Also, the optimal incentive scheme s^* adjusts to changes in w by a simple shift, that is, $s^* - w$ does not depend on w . Computationally, this means that the principal can deal with the two constraints (2) and (3) separately. He governs the agent's incentives by the choice of the differences $s(x_i) - s(x_0)$, $i = 1, \dots, N$, and he assures sufficient expected utility by adjusting $s(x_0)$. This separation result will play a key simplifying role in the subsequent multi-period analysis.

For notational convenience, we will henceforth write $S(p)$ for $S(p, 0)$. From $S(p)$ we can recover schemes in $S(p, w)$ by adding w . Note that Theorem 1 says nothing about the shape of an optimal scheme $s^*(\cdot)$. Indeed, our approach is not to characterize the solution to the one-period problem directly, but rather to say something about the relationship between the solution to the one-period problem and its multi-period extension (which could, of course, itself be viewed as an extended one-period problem in the way described earlier).

Before moving on to the multi-period case, we wish to make a brief digression into the relation between the dimensionality of the agent's feasible set of actions P and the principal's freedom to choose a sharing rule to implement any particular $p \in P$. This relationship will play a central role in the Brownian model treated in Section 4.

Notice that the principal chooses a sharing rule, which is a point in an $(N+1)$ -dimensional space, to control the agent's choice of action and to provide a particular certain equivalent. When the agent's choice is a distribution from a one-dimensional manifold in P , the mapping from sharing rules to action-certain equivalent pairs is a mapping from \mathbb{R}^{N+1} into \mathbb{R}^2 , so there is typically an $N-1$ dimensional continuum of sharing rules that lead to any implementable (p, w) pair. As the agent's action space grows in dimension (for instance, because the agent acts based on private information), the principal's options for inducing a given behavior become correspondingly more limited. We show below that in a simple version of our model where the agent's action space is of full dimension, the rule that implements any particular (p, w) pair is in fact unique. This requires an additional assumption.

ASSUMPTION A: (i) P has a nonempty interior in the N -dimensional simplex; (ii) $c(p; \theta_i) = c(p)$ for all θ_i ; (iii) $c(p)$ is continuously differentiable on P ; (iv) $c(p) - c(p') \geq \max_{ij} (\pi_i - \pi_j)$ for p on the boundary of P , where p' minimizes $c(p)$.

THEOREM 3: For any p in the interior of P , the set $S(p)$ is either empty or a singleton under Assumption A; that is, if an interior p can be implemented, then the implementing scheme (with any certain equivalent w) is unique. In particular, the sharing rule that implements an optimal action p^* with any certain equivalent w is unique.

REMARK: It could still be the case that P^* has more than one element and therefore that there are many optimal incentive schemes.

PROOF: Assume for the moment that $x_i = \pi_i$ for all i . Fix a p in the interior of P . Let c_j be the partial derivative of $c(p)$ with respect to p_j after substituting $p_0 = 1 - \sum_{i=1}^N p_i$ into the cost function. If $S(p)$ is empty we are done, so assume there exists an $s \in S(p)$. Since p is in the interior of P which is of full dimension, the first-order conditions for the agent's optimization problem (2), imply:

$$(5) \quad - \sum_{i=0}^N u'(s(x_i) - c(p)) c_j p_i + u(s(x_j) - c(p)) - u(s(x_0) - c(p)) = 0,$$

for $j = 1, \dots, N$. If we define $z_i = u(s(x_i))/u(s(x_0)) = -u(s(x_i) - s(x_0))$, then (5) can be written (using the exponential form of utility) as:

$$(6) \quad \sum_{i=0}^n rz_i p_i c_j + z_j - 1 = 0 \quad (j = 1, \dots, N).$$

Let

$$K = \sum_{i=0}^N rz_i p_i.$$

Then

$$z_j = 1 - c_j K \quad \text{and}$$

$$K = \sum_{i=0}^N r(1 - c_i K) p_i \quad \text{or}$$

$$K = r / \left(1 + \sum_{i=0}^N rc_i p_i \right),$$

which implies unique values for the z_j 's. Note that if $\sum_{i=0}^N rc_i p_i = -1$, then (6) has no solution contradicting the assumption that $S(p)$ is nonempty. Consequently, $s(x_i) - s(x_0)$ is uniquely determined for all i ; specifying $s(x_0)$ determines the agent's certain equivalent.

If we do not have $x_i = \pi_i$ as assumed, but instead have that x provides coarser information than the state, then uniqueness is implied *a fortiori*, because system (6) will have added constraints of the form $z_i = z_j$ (in case $x_i = x_j$).

The last statement of the theorem follows from the first part and the fact that part (iv) of Assumption A implies that P^* is in the interior of P . *Q.E.D.*

Theorem 3 contrasts sharply with the conclusions of the one-dimensional moral hazard models. When P is one-dimensional (that is, a curve), $S(p, w)$ normally contains infinitely many schemes (unless $N = 1$). The analysis then centers on the characterization of the best scheme in $S(p, w)$. Under some rather restrictive assumptions (see Grossman and Hart, 1983; Rogerson, 1985) one can use variational techniques to provide an intuitive and useful characterization of the best scheme (Mirrlees, 1974; Holmstrom, 1979). However, little can in general be said about the optimal choice of p . Also, once the agent's action space expands, the corresponding characterization result becomes much less informative. It may then be both realistic and analytically tractable to go to the opposite extreme and let P be of full dimension so that $s(p)$ can be obtained uniquely from (6). This route has the potential of offering more information about the optimal p to be implemented and in addition providing a useful characterization of the best scheme. In fact, this point is illustrated by earlier Hidden Information Models such as Mirrlees' (1971) model of optimal taxation (as well as by adverse selection and other nonlinear pricing models). It will also be illustrated by our Brownian model for which we shall derive an explicit closed form for the sharing rule that implements any given implementable strategy with any given certain equivalent.

3. A MULTI-PERIOD MODEL

Consider a T -period version of the previous model. In each period $t = 1, \dots, T$, the agent picks a $p' \in P$, incurring a periodic cost $c(p'; \theta_t)$. We denote the outcome x' , the resulting state θ' , and the profit level π' . It is assumed that periods are stochastically independent. We call $X' = (x^1, \dots, x^t)$ the *history* of the stochastic outcome process up to time t . A key assumption is that the agent can observe X^{t-1} before deciding p' . Thus, a *strategy* for the agent is a stochastic process $\{p'(X^{t-1})\}$.

We assume that the principal pays the agent at the end of the last period based on the entire realized path X^T of the outcome process. The incentive scheme is denoted $s(X^T)$. The agent is assumed to be concerned about his final wealth, which will equal $s(X^T) - \sum_{t=1}^T c(p'; \theta^t)$. He values this wealth according to the exponential utility function $u(y) = -e^{-ry}$. The principal's final wealth is $\sum_{t=1}^T \pi' - s(X^T)$, which he values according to the exponential utility function $v(y) = -e^{-Ry}$.

The principal's problem is to select a sharing rule and a strategy for the agent (interpreted as a set of instructions) such that it maximizes his expected end of period utility, subject to the instructions being incentive compatible and the agent being assured a minimum certain equivalent, which we henceforth normalize to zero. Formally, the problem can be stated as:

$$(7) \quad \max_{\{p'\}, s} E \left\{ v \left(\sum_{t=1}^T \pi' - s(X^T) \right) \right\}, \quad \text{subject to}$$

$$(8) \quad E \left\{ u \left(s(X^T) - \sum_{t=1}^T c(p'(X^{t-1}); \theta^t) \right) \right\} \geq u(0),$$

$$(9) \quad \{p'\} \text{ maximizes } E \left\{ u \left(s(X^T) - \sum_{t=1}^T c(p'(X^{t-1}); \theta^t) \right) \right\}.$$

The expectations are taken with respect to the distribution over states induced by the agent's strategy $\{p'\}$.

We analyze the agent's problem using dynamic programming. Fix a compensation rule $s(X^T)$ and let $\{p'(\cdot)\}$ be an optimal strategy for the agent given that rule. Define $V_t = V_t(X')$ by

$$V_t = E \left\{ u \left[s(X^T) - \sum_{\tau=t+1}^T c(p^\tau; \theta^\tau) \right] \middle| X' \right\}.$$

Since u is exponential, V_t differs from the standard dynamic programming value function only by the multiplicative factor $-u(\sum_{\tau=1}^t c(p^\tau; \theta^\tau))$, which is a positive constant from the perspective of time t . Thus, we may use V_t for purposes of dynamic programming, interpreting it as the maximal expected utility to the agent of continuing after time t given the history up to and including the outcome at time t , but excluding the accumulated sunk costs. Let $w_t = w_t(X')$ be the corresponding certain equivalent; i.e. $u(w_t) = V_t$. We wish to examine w_t as a function of x' , holding the history X^{t-1} constant. For this purpose we will write $w_t(X') = w_t(X^{t-1}, x')$. The dynamic programming equation for the agent's program

requires that $p'(X^{t-1})$ solve

$$(10) \quad \max_p \sum_{i=0}^N u(w_i(X^{t-1}, x_i) - c(p; \theta_i)) p_i.$$

Note that the sunk cost term $\sum_{\tau=1}^{t-1} c(p^\tau; \theta^\tau)$ has been dropped as it gives rise to a positive constant that can be factored out.

The problem (10) has the same form as the single-period problem (2). Thus, Theorem 2 applies. It follows that $p'(X^{t-1})$ is optimal in (10) and makes the certain equivalent of the maximum value of (10) equal to $w_{t-1}(X^{t-1})$ if and only if $w_t(X^{t-1}, \cdot) - w_{t-1}(X^{t-1}) \in S(p'(X^{t-1}))$. Thus,

$$(11) \quad w_t(X^t) = s_t(x_t; p'(X^{t-1})) + w_{t-1}(X^{t-1}),$$

where $s_t(\cdot; p)$ denotes a scheme in $S(p)$, that is, a scheme which in the single-period problem implements p with certain equivalent zero. Summing (11) over t (from 1 to T) and noting that, by definition, $w_T(X^T) = s(X^T)$ gives the following:

THEOREM 4: *A strategy $\{p'(X^{t-1})\}$ can be implemented if and only if for every date and history, $p'(X^{t-1}) \in P^0$ (that is, if and only if each $p'(X^{t-1})$ can be implemented in the single period problem). A sharing rule $s(X^T)$ implements $\{p(X^{t-1})\}$ with certain equivalent w_0 if and only if it can be written in the form:*

$$(12) \quad s(X^T) = \sum_{t=1}^T s_t(x_t; p'(X^{t-1})) + w_0,$$

where each $s_t(\cdot; p)$ is a sharing rule that implements p with certain equivalent zero in the single-period problem.

It is instructive to think of each possible outcome as being recorded in a different account. There may be fewer than $N+1$ such outcomes, since two different states may correspond to the same outcome ($x_i = x_j$). If there are M possible outcomes, then there are M accounts. Let A'_t be the number of times in the first t periods that the i th outcome occurs and let A' be the vector (A'_1, \dots, A'_M) . Also, it is convenient to represent the sharing rule $s_t(\cdot; p)$ by the M -vector $s_t(p)$ whose i th component is the compensation payable when the i th outcome occurs. Then we can write (12) as:

$$(13) \quad s(X^T) = \sum_{t=1}^T s_t(p'(X^{t-1})) \cdot (A' - A^{t-1}) + w_0.$$

Written this way, the sharing rule can be thought of as a "stochastic integral" of the account process $\{A^t\}$; it is this form that is suitable for an extension to continuous time models.

Theorem 4 recovers a sharing rule from the strategy that is to be implemented. Note that there may be many sharing rules that implement the same strategy, because we may not have uniqueness in the single-period model. Of course, in view of Theorem 3, if each p' always lies in the interior of P and $c(\cdot)$ is differentiable, we do have a unique implementation.

We wish to stress that while (12) has an additive form over time, this does not imply that every sharing rule $s(X^T)$ is additively separable in the x^t 's. Indeed no sharing rule that implements a history-contingent strategy has this separability property because, for history-contingent strategies, x^t affects the actions and hence the summands in periods after time t .

Without assuming exponential utility, we could have derived a formula similar to (12), namely:

$$(14) \quad s(X^T) = w_0 + \sum_{t=1}^T s_t[x^t; p'(X^{t-1}), w_{t-1}(X^{t-1})] - w_{t-1}(X^{t-1})$$

where $s_t(\cdot; p, w)$ is a scheme that implements p with certain equivalent w in the single-period problem. What is special about exponential utility is that $s_t(\cdot; p, w) = w + s_t(\cdot; p)$ as we saw in Theorem 2.

Turning to the Principal's Problem, we see from (12) that the principal's expected utility is:

$$(15) \quad E \left\{ v \left[\sum_{t=1}^T \pi^t - s_t(x^t; p'(X^{t-1})) - w_0 \right] \right\}.$$

Thus, the Principal's Problem reduces to choosing $\{p', s_t\}$ and w_0 with $w_0 \geq 0$ and $s_t \in S(p')$ to maximize (15). Our next Theorem characterizes one optimal solution.

THEOREM 5: *An optimal strategy for the principal to implement is $p'(X^{t-1}) = p^*$ for all t , where p^* is any single-period optimum. An optimal compensation rule to use is:*

$$(16) \quad s(X^T) = \sum_{t=1}^T s(x^t; p^*) = s(p^*) \cdot A^T,$$

where $s(\cdot; p^*)$ is an optimal single-period scheme that implements p^* , and $s(p^*)$ is the corresponding M -vector.

PROOF: The proof is by induction. For the case when $T=1$, the Theorem merely restates the definitions of p^* and s^* .

Suppose the conclusions hold when $T=\tau$. Let v_τ^* be the optimal value for the T -problem when $w_0=0$. Then, suppressing the arguments of s_t , the principal's payoff in the $T=\tau+1$ problem is, by (15) and the exponential form of v , equal to:

$$\begin{aligned} & -\exp(Rw_0)E \left[v(\pi^1 - s_1) E \left\{ v \left[\sum_{t=2}^{\tau+1} (\pi^t - s_t) \right] \middle| X_1 \right\} \right] \\ & \leq -\exp(Rw_0)E[v(\pi^1 - s_1)v_\tau^*] \leq -\exp(Rw_0)v_1^*v_\tau^*, \end{aligned}$$

where the inequalities follow from the definition of v_τ^* for $T=1, \tau$. By the inductive hypothesis, the first inequality holds as an equality when $p_t = p^*$ and $s_t = s^*$ for

all $t \geq 2$. The second holds as an equality when $p_1 = p^*$ and $s_1 = s^*$. With those substitutions, (14) becomes (16). *Q.E.D.*

Again we note that the scheme in (16) is generally not the unique scheme that implements p^* (unless Theorem 3 applies). And even if $s(x^T)$ is unique in implementing p^* , there will be several optima if P^* has more than one element. Any string of actions from P^* with accompanying one-period optimal schemes would solve the principal's problem. However, Theorem 5 tells us that there is no need to do anything more complex than apply the same scheme in each period separately.

The sharing rule in (16) has a ready interpretation in terms of aggregation and linearity. The agent's optimal compensation is a linear function of the account balances recorded in A^T . These balances represent time-aggregated information about the outcome path. In view of previous sufficient statistics results (Holmstrom, 1979; Shavell 1979), it is noteworthy that an optimal scheme can be based on the aggregated information A^T , which is not a sufficient statistic for the agent's full strategy.⁷ At the particular optimum identified in the theorem, the only binding incentive constraints are those that prohibit the agent from switching from the specified constant, uncontingent action p^* to another constant, uncontingent action. The time-aggregated information is a sufficient statistic for deviations in that class, and so an optimal compensation scheme can be based on it.

According to Theorem 5, the principal suffers no loss by relying on enumeration aggregates, that is, aggregates that report the number of times that each observable outcome occurs. Can some further aggregation of information be done without loss to the principal? For instance, suppose each x^t is a monetary payoff, say profits, of the agent's activity in period t , so that A_i^t is the number of periods the agent has made a profit of $\$x_i$. When will the optimal incentive scheme be linear in the total profits $\sum_{t=1}^T A_i^t x_i$? Theorem 5, it should be stressed, does not tell us that the optimal scheme is linear in profits: Two periods with a profit of \$100K each are not generally compensated the same as one of \$50K and one of \$150K. Nevertheless, there is a special case for which the optimal scheme in (16) is linear in money. That is the case when there are only two outcomes yielding two different profit levels; in other words, if the agent controls a binomial process. In that case, the two account balances A_1^T and A_2^T are both linear in total profits, so the compensation rule (16) is linear in profits.

In the next section we will take a more careful look at linearity in nonenumeration accounts using a continuous time model, which will prove to be better suited for this analysis. Before that a few remarks on the robustness of our discrete time results are in order.

1. Two assumptions drive our results: exponential utility and a history- and time-independent technology. We elaborate on the roles on these assumptions

⁷ This conclusion does not contradict the results of Holmstrom (1979) and Shavell (1979). These results posit that only local incentive constraints are binding and show that the optimal compensation scheme will utilize any evidence that the employee has "shirked" (made a local deviation). In their models, as in ours, the optimal contract does not use evidence that bears solely on violations of nonbinding incentive constraints.

here. Theorem 4, stating how a given strategy for the agent can be implemented, follows directly from the form of the agent's preferences. If the agent's feasible set P and cost function $c(\cdot, \cdot)$ at time t depended on the history of public outcomes X^t , our argument would establish the decomposition (12) where $s_t(\cdot; p X^t)$ implements p in the one-period problem indexed by the available technology at time t , which might depend on the history X^t . Theorem 5—the basic aggregation result—requires the two additional assumptions that the principal has exponential utility and that the technology is stationary and history independent. These assure that the T single period problems faced by the principal are identical, and so have identical solutions.

2. Our conclusion that the optimal sharing rule is a linear function of time aggregates depends on specifying preferences for the principal and agent that are stationary over time. However, the specification used here is not the only one with that property. In Fudenberg, Holmstrom, and Milgrom (1986), similar results are obtained with a time-separable utility function and periodic consumption, when the agent has access to a bank for borrowing and saving. That model assumes that both the principal and the agent discount payoffs over time. Then, the relevant time aggregates are discounted enumeration accounts. Of course, if periods of pay are relatively close to each other, the effects of discounting on the aggregation are negligible, and the optimal compensation scheme resembles the one derived here.

3. As we stressed in the introduction, the timing of information is a crucial aspect of the formulation. It matters a great deal that the agent gets to observe his performance before proceeding to the next period. Thus, if instead of choosing sequentially the probabilities of success of T identical binomial variables, the agent were to choose them simultaneously, the optimal scheme would no longer be linear in the number of successes. The principal could do better, because he would be faced with fewer incentive constraints. This point is dramatically illustrated by the contrast between the Mirrlees "normal distribution" model, described earlier, and the Brownian model which we turn to next.

4. THE BROWNIAN MODEL

In this section we investigate a situation in which the agent takes actions very frequently in time. We will use a controlled Brownian process as the relevant abstraction. The formal analysis of the model requires some sophisticated mathematics, but it repays the investment with some strikingly simple results. We will show that if the agent controls the drift rate of a multi-dimensional Brownian process over the unit time interval then, analogously to the multinomial model of the preceding section, the optimal incentive scheme will be linear in the end-of-period positions of the components of the process. But there are two significant advantages to the Brownian specification. One is that the optimal sharing rules are especially easy to characterize and compute. The other is that when compensation schemes are restricted to be functions of certain linear aggregates of the enumeration accounts, the optimal schemes are necessarily

linear functions of these aggregates. This is important because such common accounting aggregates as profits, sales, inventories, selling expenses, and cost-of-goods sold are all aggregates of this type.

To motivate our results and illuminate their relationship to the discrete-time theory, we begin with a variant of our earlier discrete-time model.

Assume that the agent acts T times in a period of fixed length such as a year. Assume further that (i) T is large and (ii) that the possible marginal costs of the agent's actions and the possible variations in the profits earned in any single period are quite small as a fraction of risk tolerances of the principal and agent.⁸ Assumption (ii) implies that if the single period profits had mean μ and variance σ^2 , the principal's certain equivalent value for the profit lottery would be approximately $\mu - R\sigma^2/2$ and the agent's would be approximately $\mu - r\sigma^2/2$. We further assume (iii) that P has dimension N and is open relative to the N -simplex of probability distributions, and (iv) that the agent's costs take the form $c = c(p)$, with $c(\cdot)$ differentiable. By Theorem 3, the last two assumptions imply that there is at most a single sharing rule that implements any action $p \in P$ with certain equivalent w .

Consider the agent's single period problem under the assumptions just described. The agent chooses $p \in P$, incurs a cost $c(p)$, and collects a payment s_i if the realized profit level is π_i . (For notational simplicity, we assume here that profit levels are observed.) Suppose that p is implemented by the sharing rule $s = (s_0, \dots, s_N)$. Define $\bar{s}(\hat{p}) = \sum \hat{p}_i s_i$ to be the agent's expected compensation if, faced with sharing rule s , he takes the action \hat{p} . Then the agent's certain equivalent payoff is approximately his mean return minus a risk premium:

$$(17) \quad w(\hat{p}) = -c(\hat{p}) + \bar{s}(\hat{p}) - \frac{1}{2}r \sum_i \hat{p}_i (s_i - \bar{s}(\hat{p}))^2.$$

Let e_i be the $(N+1)$ -vector whose i th component is unity and whose other components are zero. Define

$$c_i(p) = \left. \frac{d}{d\varepsilon} \right|_{\varepsilon=0} c(\varepsilon e_i + (1-\varepsilon)p).$$

By hypothesis, p maximizes the certain equivalent (17) and so satisfies the following first-order necessary conditions:

$$(18) \quad s_j - \bar{s}(p) = c_j(p) + \frac{1}{2}r[(s_j - \bar{s}(p))^2 - \sigma_A^2] \quad \text{for } j=0, \dots, N,$$

where $\sigma_A^2 = \sum p_i (s_i - \bar{s}(p))^2$ is the variance of the agent's compensation. Henceforth, we write \bar{s} for $\bar{s}(p)$.

In (18), the coefficient of r has mean zero (by the definition of σ_A^2). So, it is smaller than $\alpha^2 \equiv \text{Max}_i (s_i - \bar{s})^2$. By assumption (ii), $r\alpha$ is quite small, so the term

⁸ Risk tolerance is the inverse of the coefficient of absolute risk aversion and is measured in units of wealth. Risk neutrality corresponds to a risk tolerance of $+\infty$.

in (18) involving r is much smaller than α . Hence, an approximate solution of (18) is:

$$(19) \quad s_j \cong \bar{s} + c_j(p)$$

with errors much smaller than α , for all j . Computing \bar{s} from (17) and substituting into (19) yields:

$$(20) \quad s_j \cong w + c(p) + r\sigma_A^2/2 + c_j(p) \quad \text{for all } j.$$

Thus, we have determined a simple approximate relation between the action p to be implemented and the unique rule that implements it with certain equivalent w for the case where single period risks are small.

Notice that $\sum p_j c_j(p) = 0$. Consequently, if the sharing rule (20) is used to induce the agent to take action p , the principal's certain equivalent will be approximately:

$$(21) \quad \sum p_i \pi_i - w - c(p) - r\sigma_A^2/2 - R\sigma_P^2/2$$

where $\sigma_P^2 = \sum_i p_i [\pi_i - c_i(p)]^2 - [\sum_i p_i (\pi_i - c_i(p))]^2$ is the variance of the principal's return. The approximate forms (20) and (21) will be exact for the Brownian model (after a suitable change of variables).

To derive a Brownian model that approximates some discrete time model, we change our notation and normalizations as follows: Let $\hat{\pi}_i = (\pi_i - \pi_0)T$ denote the profit in excess of $\pi_0 T$ that would result if outcome θ_i occurred in all T periods. We shall regard $\hat{\pi}_i/T$ as the contribution to profit from a single occurrence of θ_i . Replace the enumeration accounts $\{A'_i; i = 0, \dots, N\}$ by the accounts $Z_i(t/T) = (\hat{\pi}/T)(A'_i - \hat{p}_i t)$ that record the accumulated profits from events of type θ_i , in excess of some "normal" standard, up to time t . Here, \hat{p} is some vector of strictly positive probabilities that defines a suitable standard. The time index for $Z_i(\cdot)$, instead of being an integer designating the period number, now records the fraction of the accounting cycle that has elapsed. Since $Z_0(\cdot) \equiv 0$, we omit it from our list of accounts. Represent the agent's action choice as one of controlling the rate of drift μ of the vector of accounts $Z = (Z_1, \dots, Z_N)$. Thus, $\mu_i = \hat{\pi}_i(p_i - \hat{p}_i)$. Finally, let $c(\mu) = c(p)T$ so that the agent's single period cost of control is $c(\mu)/T$. We assume that μ lies in some open bounded set $M \subset \mathbb{R}^N$.

Suppose the agent chooses a constant action μ . Then, for any interval of time (t, t') , $Z(t') - Z(t) - \mu(t' - t)$ is the sum of approximately $T/(t' - t)$ independent and identically distributed random vectors with finite range and mean zero. Hence, for T large, it has approximately a multivariate normal distribution with a covariance matrix that is proportional to $(t' - t)$. For the Brownian model, we take this approximation to be exact and write:

$$(22) \quad dZ = \mu dt + dB,$$

where B is a driftless N -dimensional vector Brownian motion with covariance matrix Σ . In particular, $B(0) = 0$ and for all $t' > t$, $\text{Var}(B(t') - B(t)) = (t' - t)\Sigma$.

As in Section 3, we assume that the agent chooses $\mu(t)$ at time t in full knowledge of the history $\{Z(\tau) | \tau \leq t\}$. The principal can base the agent's compensation, $s(Z^1)$ on the entire realized path of the Brownian process: $Z^1 \equiv \{Z(\tau) | 0 \leq \tau \leq 1\}$. Given our construction, it is natural to assume that the agent

incurs an instantaneous cost of $c(\mu(t)) dt$. The derivative of c is denoted c' (its values are N -vectors); we assume that c is convex and continuously differentiable and that c' is bounded on M . The Principal's profits over the unit time interval are $\sum Z_i(1)$. The agent's final payoff is $E\{u(s(Z^1) - \int_0^1 c(\mu(t)) dt)\}$ and the principal's is $E\{v(\sum_{i=1}^N Z_i(1) - s(Z^1))\}$, where u and v are exponential utility functions as before and the expectation is taken with respect to the distribution induced by $\{\mu(t)\}$.

Analogously to Theorem 4 and equation (13) we have the following Theorem.

THEOREM 6: *The stochastic process $\{\mu(t); 0 \leq t \leq 1\}$ is implemented with certain equivalent w by sharing rule $s(Z^1)$ only if:*

$$(23) \quad s(Z^1) = w + \int_0^1 c(\mu(t)) dt + \left\{ \int_0^1 c'(\mu(t))^T dZ - \int_0^1 c'(\mu(t))^T \mu(t) dt \right\} \\ + (r/2) \int_0^1 c'(\mu(t))^T \Sigma c'(\mu(t)) dt.$$

The proof is given in the Appendix. The "local" sharing rules which form the integrands in (23) are identical in form to the sharing rules described by (20). The first two terms in (23) provide the agent with the desired certain equivalent plus direct compensation for the costs he would incur by following directions. The $c'(\mu(t))^T dZ$ term in the first line provides the incentive for effort at time t . From that incentive term we subtract its conditional expectation at time t , $c'(\mu(t))^T \mu(t) dt$, so that the component of compensation in the bracketed term in (23) has mean zero and corresponds to the c_j term in (20). Finally, to compensate the agent for the risk he must bear, a risk premium is paid. The agent's incentive income "at time t " has variance $c'(\mu(t))^T \Sigma c'(\mu(t)) dt$ which, multiplied by $r/2$, gives the instantaneous risk premium.

Theorem 6 is stated as an "only if" result. For the special case where μ is constant over time, Theorem 4.1 of Fleming and Rishel (1975) implies that the "if" direction holds as well. We conjecture that Theorem 6 holds as an "if and only if" result for all instructions $\{\mu(t); 0 \leq t \leq 1\}$ such that $\mu(t)$ always lies in M , but we shall not need that result to solve the principal's problem for this model.

The *Principal's Brownian Problem* is defined as follows. The principal chooses instructions $\{\mu(t); 0 \leq t \leq 1\}$ and a sharing rule s to maximize $E\{v(\sum Z_i(1) - s)\}$ subject to the constraint that s implements the instructions with certain equivalent not less than zero. We now show that the solution to the Principal's Brownian Problem corresponds to the solution of a static problem in which the principal is constrained to choose a linear function of the $Z_i(1)$'s and the agent is constrained to choose a constant μ once and for all at time zero. In such a problem, the principal would choose an instruction μ and agent's certain equivalent w to maximize his certain equivalent, given the two constraints that $w \geq 0$ and that the gradient of the sharing rule must be $c'(\mu)$. The principal's certain equivalent for the static problem is then:

$$(24) \quad \sum_{i=1}^N \mu_i - \{w + c + (r/2)c'^T \Sigma c'\} - (R/2)(e - c')^T \Sigma (e - c')$$

where the functions c and c' are to be evaluated at μ and where e is a vector of ones. The certain equivalent (24) has the same general form as (21). It consists of the expected return minus the expected compensation paid to the agent minus a premium for the risk born by the principal.

THEOREM 7: *Suppose that a $\mu^* \in M$ exists that maximizes (24). Then an optimal solution to the Principal's Brownian Problem is to instruct the agent to set $\mu(t) \equiv \mu^*$ and to set:*

$$(25) \quad s(Z^1) = c(\mu^*) + c'(\mu^*)^T (Z(1) - \mu^*) + (r/2) c'(\mu^*)^T \Sigma c'(\mu^*).$$

PROOF SKETCH: Let the horizon of the principal's problem vary, and denote it by τ . Let $v^*(\tau)$ be the principal's optimal value. By Theorem 5, if $\{\mu(t); 0 \leq t \leq 1/n\}$ is an optimal instruction when $\tau = 1/n$, then n repetitions of it form an optimal instruction when $\tau = 1$. Writing $dt = 1/n$ and using Theorem 6 and Ito's lemma:⁹

$$(26) \quad v^*(dt) \approx \max_{\mu, w \geq 0} E_{\mu} \left[-\exp \left\{ -R \left[e^T dZ - \left(w + c dt + c'^T (dZ - \mu dt) + \frac{r}{2} c'^T \Sigma c' dt \right) \right] \right\} \right] \\ \approx -1 + R dt \cdot \max_{\mu} \left\{ e^T \mu - \left(c + \frac{r}{2} c'^T \Sigma c' \right) - \frac{R}{2} (e - c')^T \Sigma (e - c') \right\}$$

where the optimal value of w is zero. It follows that the instructions specified in the Theorem attain the optimum for the $\tau = 1/n$ problem, and hence for the $\tau = 1$ problem. Then, (25) follows from Theorem 6. Q.E.D.

The close parallel between the optimal incentive scheme in the Brownian model and the one in the earlier discrete time model is evident by comparing (25) with (16). In both we have a finite set of time-aggregated accounts ($Z(1)$ in the Brownian case, A^T in the discrete case) and an optimal rule that is a linear function of these accounts. Detailed information about the path of the stochastic process is unnecessary for the construction of optimal incentives. On the other hand, because the coefficients of the individual accounts will normally differ (as a function of the agent's characteristics as well as the technology; cf. footnote 2), further aggregation across accounts is not generally possible. In particular, there is no reason for the agent's optimal incentive scheme in (25) to be a linear function of total end-of-period profits $\sum Z_i(1)$.

If the agent controls a one-dimensional Brownian process, then of course the optimal scheme is linear in profits. But as our parallel development of the discrete and the continuous time models should make clear, the one-dimensional Brownian case only represents fairly a situation where the agent controls a

⁹ This argument proceeds as if the optimal $\{\mu(t)\}$ were known to be continuous, a restrictive assumption that can be dropped for this case by standard arguments in continuous time dynamic programming. See Fleming and Rishel (1975, Chapter 4.1).

frequent Bernoulli process. For the multinomial case, a multi-dimensional Brownian process is necessary in order to keep track of the frequency information of individual profit outcomes, which is crucial for the optimal design in the discrete model.

However, there is a way of expanding the applicability of the one-dimensional Brownian model. Let the agent control a multi-dimensional Brownian process as before, but assume that the principal has less detailed information about the path of the process. In particular, consider the *Principal's Aggregate Problem* in which the principal, rather than observing the vector $Z(t)$ over time, is restricted to observe only the path of some K -dimensional ($K < N$) accounting aggregate $Y(t) = AZ(t)$, where A is a $K \times N$ matrix.

THEOREM 8: *Let*

$$(27) \quad M(A) = \{\mu \in M \mid c(\mu) \leq c(\mu') \text{ for all } \mu' \text{ such that } A(\mu - \mu') = 0\}.$$

An optimal solution to the Principal's Aggregate Problem is to implement $\mu(t) = \mu^$, where*

$$(28) \quad \mu^* \text{ maximizes (24) in } M(A).$$

The sharing rule s is determined from (25).

REMARK: For each μ in $M(A)$, there exists a Lagrange multiplier λ such that $c'(\mu) = \lambda A$, because μ is cost minimizing subject to the constraint that $A\mu$ takes a given value. Hence the linear-in- $Z(1)$ sharing rule with coefficients c' obtained from (25) is also expressible as linear-in- $Y(1)$ sharing rule with coefficient λ .

PROOF SKETCH: Given any instructions $\{\mu(t)\}$, the path of $\{Y(t)\}$ that results from the agent following instructions satisfies $dY = A\mu dt + A dB$. An agent who deviates from the instructions to adopt some $\{\mu^*(t)\}$ which solves (for each t) $\text{Min}_{\mu} c(\mu)$ subject to $A(\mu - \mu(t)) = 0$ generates the same path Y as he would by following instructions, and therefore receives the same compensation $S(Y)$, but at a lower personal cost. Hence, $\{\mu(t)\}$ is not implementable unless it is in $M(A)$.

Maximizing the principal's objective subject to this constraint as in the proof of Theorem 7 and observing (see the preceding Remark) that the resulting sharing rule depends only on Y completes the argument. Q.E.D.

Theorem 8 concludes that an optimal sharing rule for the Principal's Aggregate Problem is a linear function of nonhomogeneous accounting aggregates, rather than just of the pure enumeration aggregates used in Section 3. As a special case, if the principal can only observe the path of total profits (i.e. $A = (1, \dots, 1)$), then the optimal rule is linear in profits even though the agent controls a multi-dimensional Brownian process. A similar conclusion emerges if the agent exercises some discretion over the accounting system. For example, suppose the

agent can choose within broad limits how to record common revenue and expense items. Then, the relevant cost function might be $c(\mu) = c(\mu_1 + \dots + \mu_N)$, so a direct application of (25) establishes that the optimal scheme is linear in profits. On the other hand, if the agent could control the accounting for expenses, but only within the expense categories, and the accounting for revenues, but only within the revenue categories, then the optimal scheme would be a linear function of revenues and expenses separately, but would not necessarily be a linear function of profits. Thus, our model is consistent with the idea that it may be better to use "cost centers" (and "revenue centers") than "profit centers" for evaluating managerial performance.

Theorem 8 has no parallel in other continuous time models. For example, if the agent controls the rates of jump in a vector Poisson process, where different events generate jumps in profits of different sizes, one can recover the number of jumps of each type from the corresponding Poisson profit process. The Brownian model differs from the discrete models and other stationary continuous time control models in that linear aggregation across accounts in the Brownian model using an arbitrary $K \times N$ full rank matrix reduces the dimension of the implementable actions from N to K : In the Brownian model aggregation across accounts destroys information in an intuitive way.

So far, we have been careful to treat our Brownian model as a limiting case of the discrete time multinomial models described in the previous section. Other processes which are "close" to Brownian motion in the traditional sense of weak convergence or probability measures (Billingsley, 1968) may lead to results that are far different from those we have obtained, because weak convergence of distributions does not preserve the subtle informational properties of distributions that are crucial for economic analysis (Milgrom and Weber, 1985). For example, suppose that the agent controls a Brownian motion but can change the drift rate μ only at discrete dates $0 = t_0 < t_1 < \dots < t_n = 1$, and that the principal can observe $Z(t_0), \dots, Z(t_n)$. In that case, Mirrlees' (1974) analysis applies to each time segment of the path, and the first-best solution can be approached arbitrarily closely. Here, the Brownian model fails as an approximation because the principal's information is so very fine compared with the agent's action space that the principal has many ways to induce any desired action. However, a slight change in the timing of the agent's information reverses that conclusion. Suppose that the agent observes $B(t_{k+1}) - B(t_k)$ before choosing $\mu(t_k)$; this enlarges the agent's strategy set. Then, one can show that the unique sharing rule (up to a constant) that induces any particular sequence of pure action choices $\mu(t_k)$ at the times t_k is one that pays $\sum_{k=0}^n c'(\mu(t_k))^T (Z(t_{k+1}) - Z(t_k))$ and that the first-best outcome cannot be approached.

From our analysis and these examples, we find that the Brownian model is properly used to represent those situations in which (i) the principal's information is so meager relative to the agent's opportunities for action that there is at most a single sharing rule to implement any given strategy for the agent, (ii) risk tolerances are large compared to the range of single period profits and to marginal costs of control, and (iii) repetitions are frequent in time.

5. EXAMPLES

Computational ease gives our linearity results substantial methodological value. One can solve the dynamic problem as a static one with the *ad hoc* restriction that the sharing rule is linear. This is particularly convenient since we are pairing normal distributions with exponential utility. Below we offer some examples to illustrate how simple the analysis becomes when the linearity restriction is appropriate. At the same time we wish to caution the reader against indiscriminate application of the *ad hoc* principle. It requires some restrictive assumptions on the information and production technologies; in particular, the agent cannot control privately the covariance of the process.

Sensitivity analysis: We begin by computing a closed form expression for the linear rule when the cost of varying the drift rate μ in the one-dimensional Brownian model is quadratic: $c(\mu) = (k/2)\mu^2$. For simplicity, the principal is assumed risk neutral (i.e., $R = 0$).

As a point of reference, note that the first-best solution in this situation, attainable if the agent's choice of μ were costlessly observable, would entail instructing the agent to choose $\mu = k^{-1}$ and paying the agent a constant wage equal to the cost of his action, $(2k)^{-1}$ (assuming a zero certain equivalent). This yields the profit level: $\pi = (2k)^{-1}$.

When there is moral hazard, a constant payment is no longer optimal, but some other linear rule, $s(z) = \alpha z + \beta$, is. To determine α and β , note that the agent's certain equivalent given $s(z)$ is

$$(29) \quad \alpha\mu + \beta - (k/2)\mu^2 - (r/2)\alpha^2\sigma^2.$$

Maximizing (29) over μ , gives the first-order condition $\alpha = \mu k$. Substituting this into (29) and setting the expression equal to zero determines β . The result is the incentive scheme:

$$s(z) = k\mu(z - \mu) + (r/2)k^2\mu^2\sigma^2 + (k/2)\mu^2.$$

Since $s(\cdot)$ is linear, the principal's expected payoff is $\mu - s(\mu)$ which, maximizing over μ , gives the second best solution

$$(30) \quad \mu^* = (1 + rk\sigma^2)^{-1}k^{-1},$$

implying

$$(31) \quad \pi^* = (1 + rk\sigma^2)^{-1} \cdot (2k)^{-1},$$

$$(32) \quad \alpha^* = (1 + rk\sigma^2)^{-1}.$$

Comparing the first-best values of effort and profit with μ^* and π^* , we see that the moral hazard problem causes a reduction in the action as well as a reduction in the principal's net payoff. (In fact, this holds true for the one-dimensional Brownian model regardless of the cost function; see (24).) Both μ^*/μ and π^*/π are equal to $(1 + rk\sigma^2)^{-1}$. Thus, a small degree of risk aversion or uncertainty, or a flat marginal cost of effort will allow a solution close to the

first-best. All of this accords well with intuition. Nevertheless, the only related result from the general theory is that additional information is valuable for the agency problem (Holmstrom, 1979; Shavell, 1979), which in the case of a normal distribution implies that a reduction in the variance will benefit the principal. This is confirmed by the formulas above. Note though that unlike the general informativeness results, which pertain to costless information acquisition, the closed form formulas can be used to evaluate the net benefits of costly monitoring.

Relative evaluations: Continuing on the previous example, assume that in addition to z , the principal observes another signal y , which could be included in the contract. We may think of y as a market index, the output of some other agents or anything else that has bearing on the agency problem. We assume that (z, y) are jointly normally distributed, that $E[y] = 0$, $\text{Var}(y) = \sigma^2$, and $\text{Cov}(z, y) = \rho\sigma\delta$.

In the underlying dynamic version of the model, (z, y) follows a two-dimensional Brownian process. According to Theorem 7, the optimal sharing rule will therefore be linear in the two accounts: $s(z, y) = \alpha_1 z + \alpha_2 y + \beta$. It is easy to calculate the coefficients along the lines described above (or use (25)) and one finds that:

$$\alpha_1^* = (1 + rk\sigma^2(1 - \rho^2))^{-1},$$

$$\alpha_2^* = -\alpha_1^*(\sigma/\delta)\rho.$$

These induce the following action and profit level:

$$\mu^* = (1 + rk\sigma^2(1 - \rho^2))^{-1}k^{-1}, \quad \text{and}$$

$$\pi^* = (1 + rk\sigma^2(1 - \rho^2))^{-1}(2k)^{-1}.$$

From these formulas we see that observing an additional signal y is equivalent to a reduction in the variance of z from σ^2 to $\sigma^2(1 - \rho^2)$. The reduced variance is the conditional variance of z given y as one might have guessed directly. Notice that if y and z are perfectly correlated ($\rho = +1$ or -1), the first-best can be attained; one can set $\alpha_1 = 1$ and filter out all uncertainty using y .

In general, y is used to reduce risk in accordance with earlier results on relative performance evaluation (see Holmstrom, 1982). In particular, we could form the aggregate performance index $x = z - (\sigma/\delta)\rho y$ and then pay the agent: $s(x) = \alpha_1^* x + \beta$. This index is a sufficient statistic for information about the agent's action.

This example could be extended to incorporate investment decisions. For instance, assume that $z = m + \mu + \theta$, where m is the mean of the project that the agent is controlling. Project choice corresponds to a point (m, σ^2, ρ) . Since the agent was assumed to control the drift, but not the covariance matrix of the Brownian process, project selection must here be assumed undertaken jointly by the principal and the agent; we cannot allow the agent discretion over project choice without altering the linearity results. This is a key restriction of our model. On the other hand, assuming joint selection of projects, the optimal incentive scheme is given as above (only the constant coefficient β is affected by the presence of

m). Maximal profit from a project is then

$$(33) \quad \pi^* = m + (1 + rk\sigma^2(1 - \rho^2))^{-1}(2k)^{-1}.$$

As we mentioned above, $\sigma^2(1 - \rho^2)$ measures the conditional variance of z , that is, the idiosyncratic risk of the project. In contrast to standard portfolio models, we see that idiosyncratic risk enters investment decisions; in fact, with a risk neutral principal, only idiosyncratic risk matters. Equation (33) gives the precise trade-off between mean and idiosyncratic risk.

Allocation of effort: As a final example, consider the case where the agent can allocate his effort between two activities: $z_1 = \mu_1 + \theta_1$ and $z_2 = \mu_2 + \theta_2$. Assume the two activities are stochastically independent and that the cost function is $c(\mu_1, \mu_2) = \mu_1^2/2 + \mu_2^2/2$ (reflecting decreasing returns to each activity). If the principal can only observe $z = z_1 + z_2$, then the solution will obviously be linear in z and the agent will allocate his effort symmetrically, which is an efficient allocation given the informational assumption.

Somewhat surprisingly, the situation is different if z_1 and z_2 can be observed separately. We know that the optimal scheme will be linear, $s(z_1, z_2) = \alpha_1 z_1 + \alpha_2 z_2 + \beta$, but α_1 need not equal α_2 even though the marginal costs and marginal products are identical in the two activities. In fact, the only case for which the two are set equal is if z_1 and z_2 have the same variance. In general, $\alpha_1 > \alpha_2$ if $\sigma_1 < \sigma_2$, which can be seen by noting that the solution (30)–(32) applies to each activity separately.

It follows that there is in general value in observing the agent's activities separately and also that the incentive weights among different projects (or accounts) are not solely a function of their marginal costs and benefits, but also their variances. As an example, if the agent can allocate time between reducing costs or increasing revenues and is equally effective at both, but if the revenues are subject to more exogenous variance than the costs, then the incentive scheme should not be based on profits alone; it should reward cost reductions more highly than revenue increases (implying that more effort will be steered into cost reduction).

6. CONCLUSION

There are two main ideas that motivate the kind of analysis we have pursued. The first is that one need not always use *all* of the information available for an optimal incentive contract. Accounting information which aggregates performance over time is sufficient for optimal compensation schemes in certain classes of environments and it is sometimes possible to aggregate further over the various accounts. The second idea is that optimal rules in a rich environment must work well in a range of circumstances and will therefore not be complicated functions of the outcome; indeed, in our model, linear functions are optimal.

Models that derive optimal rules in which small differences in outcomes lead to large differences in compensation are invariably based on an assumption that

the agent finds it impossible, or very expensive, to cause small changes in individual outcomes. The optimal rule in such cases is usually inordinately sensitive to the distributional assumptions of the model. For example, in the model where the agent makes a one-shot choice that determines the mean of a normal distribution, by changing the distribution of outcomes for each action on a set of arbitrarily small positive probability ε , some of the near-optimal rules derived for that model can be made to perform worse than a flat compensation scheme, which provides no incentives at all for the agent to incur costs to increase production.

Linear rules, in contrast, are strikingly robust. For example, in the Brownian model, the agent's optimal response to a linear rule and the principal's expected payoff do not depend at all on the timing of the agent's information. Nor does this conclusion depend on normality: an (exponential) agent's optimal response to a linear rule where he adds drift to any stochastic process is always the same.

It is probably the great robustness of linear rules based on aggregates that accounts for their popularity. That point is not made as effectively as we would like by our model; we suspect that it cannot be made effectively in any traditional Bayesian model. But issues of robustness lie at the heart of explaining any incentive scheme which is expected to work well in practical environments.

Yale School of Organization and Management, Box 1A, New Haven, CT 06520, U.S.A.

Manuscript received May, 1985; revision received July, 1986.

APPENDIX

PROOF OF THEOREM 6: Suppose the sharing rule s (a random variable measurable with respect to \mathcal{F}_1 , the σ -field generated by the path of $Z(\cdot)$ up to time 1) implements $\{\mu(t)\}$ with certain equivalent w . We refer to the paths of $Z(\cdot)$ as "outcomes." Let E_m , E_μ , etc. refer to expectations using the distribution over outcomes determined by the agent's control strategy $\{m(t)\}$, $\{\mu(t)\}$, etc. Define:

$$(A1) \quad F(\tau; \mu'; m) = E_m \left[u \left(s - \int_0^\tau c(\mu'(t)) dt - \int_\tau^1 c(m(t)) dt \right) \middle| \mathcal{F}_\tau \right] \\ = F(\tau; \mu; m) C(\tau; \mu'),$$

where we define

$$C(\tau; \mu') = \exp \left[r \int_0^\tau [c(\mu'(t)) - c(\mu(t))] dt \right].$$

F represents the agent's conditional expected utility at time τ if he has followed strategy $\{\mu'(t)\}$ to that time and then switches to strategy $\{m(t)\}$. Let $V(\tau; \mu')$ be the maximal value of the agent's dynamic programming problem given the information at time τ when the agent has followed strategy $\{\mu'(t)\}$ up to time τ . Then:

$$(A2) \quad V(\tau; \mu') = \max_m F(\tau; \mu'; m) \\ = \max_m F(\tau; \mu; m) C(\tau; \mu') = V(\tau; \mu) C(\tau; \mu').$$

Since $\{\mu(t)\}$ is an optimal strategy,

$$(A3) \quad V(\tau; \mu) = F(\tau; \mu; \mu) = E_\mu \left[u \left(s - \int_0^1 c(\mu(t)) dt \right) \middle| \mathcal{F}_\tau \right],$$

so that $\{V(\tau; \mu), \mathcal{F}_\tau\}$ is a martingale under E_μ . Moreover, the σ -fields $\{\mathcal{F}_\tau\}$ generated by $\{Z(t)\}$ are identical to the σ -fields generated by the process $Z(t) - \int_0^t \mu(\tau) d\tau$. Under E_μ , this process is a driftless Brownian motion. A theorem of Meyer (see Jacod, 1977) holds that every martingale over the Brownian fields is representable as an Ito stochastic integral with respect to the driftless Brownian motion, as follows:

$$(A4) \quad dV(\tau; \mu) = \gamma(\tau)^T d \left[Z(\tau) - \int_0^\tau \mu(t) dt \right] = \gamma^T dZ - \gamma^T \mu dt,$$

where $\{\gamma(t); t \geq 0\}$ is an N -dimensional stochastic process such that $\gamma(t)$ is \mathcal{F}_t -measurable and $\int \gamma(t)^T \Sigma \gamma(t) dt$ is almost surely finite under the distribution induced by μ . When the agent chooses the control strategy $\{\mu'(t)\}$, we have $dZ = \mu' dt + dB$. So, for any control strategy $\{\mu'(t)\}$ that the agent might choose, we have that $dV(t; \mu) = \gamma^T(\mu' - \mu) dt + \gamma^T dB$. Hence, using (A2):

$$(A5) \quad \begin{aligned} dV(t; \mu') &= d[V(t; \mu)C(\tau; \mu')] \\ &= \{\gamma^T[\mu'(t) - \mu(t)] + r[c(\mu'(t)) - c(\mu(t))]V(t; \mu)\}C(\tau; \mu') dt \\ &\quad + C(\tau; \mu')\gamma^T dB. \end{aligned}$$

By the Principle of Optimality, if $\{\mu'\}$ is an optimal strategy, for all t , $\mu'(t)$ must maximize the drift rate of $V(t; \mu')$ (the coefficient of dt in (A5)). Then since $\{\mu\}$ is by hypothesis an optimal strategy, each $\mu(t)$ must solve: $\max_m \gamma(t)^T m + rc(m)V(t; \mu)$. That leads to the first-order necessary condition:

$$(A6) \quad -rc'(\mu(t)) = \gamma(t)/V(t; \mu).$$

Let $w(t) = u^{-1}(V(t; \mu)) = -\ln(-V(t; \mu))/r$ be the certain equivalent corresponding to the value function $V(t; \mu)$. By Ito's Lemma (suppressing all the arguments), since $Z(\tau) - \int_0^\tau \mu(t) dt$ is a Brownian motion when the agent uses strategy $\{\mu\}$, $-r dw = (\gamma/V) \cdot (dZ - \mu dt) - \frac{1}{2}(\gamma/V)^T \Sigma (\gamma/V) dt$. Using (A6) to eliminate the γ/V terms, one has:

$$(A7) \quad dw(t) = c'(\mu)^T dZ - [c'(\mu)^T \mu - (r/2)c'(\mu)^T \Sigma c'(\mu)] dt.$$

Integrating (A7) and letting $w(0) = w$ leads to:

$$(A8) \quad \begin{aligned} w(1) &= w + \int_0^1 c'(\mu(t))^T dZ - \int_0^1 c'(\mu(t))^T \mu(t) dt \\ &\quad + (r/2) \int_0^1 c'(\mu(t))^T \Sigma c'(\mu(t)) dt. \end{aligned}$$

But $s = w(1) + \int_0^1 c(\mu(t)) dt$, by the construction of $w(1)$. Then, substituting for $w(1)$ from (A8), one obtains the required form (23). Q.E.D.

REFERENCES

- ARROW, KENNETH (1984): "The Economics of Agency," in *Principals and Agents: The Structure of Business*, ed. by John Pratt and Richard Zeckhauser. Boston: Harvard Business School Press, pp. 37-51.
- BILLINGSLEY, PATRICK (1968): *Convergence of Probability Measures*. New York: John Wiley and Sons.
- BULL, CLIVE, ANDREW SCHOTTER, AND KEITH WEIGELT (1985): "Tournaments and Piece Rates: An Experimental Study," Economics Research Report #85-21, New York University.
- FLEMING, WENDELL H., AND RAYMOND W. RISHEL (1985): *Deterministic and Stochastic Optimal Control*. New York: Springer-Verlag.
- FUDENBERG, DREW, BENGT HOLMSTROM, AND PAUL MILGROM (1986): "Repeated Moral Hazard with Borrowing and Saving: The Exponential Case," draft.
- GROSSMAN, SANFORD, AND OLIVER HART (1983): "An Analysis of the Principal-Agent Problem," *Econometrica*, 51, 7-45.
- HOLMSTROM, BENGT (1979): "Moral Hazard and Observability," *Bell Journal of Economics*, 10, 74-91.
- (1982): "Moral Hazard in Teams," *Bell Journal of Economics*, 13, 324-340.
- JACOD, JEAN (1979): "A General Theorem of Representation for Martingales," *Proceedings of Symposia in Pure Mathematics*, 31, 37-53.

- LAFFONT, JEAN-JACQUES, AND JEAN TIROLE (1985): "Auctioning Incentive Contracts," Working Paper, MIT.
- McAFEE, R. PRESTON, AND JOHN McMILLAN (1986): "Competition for Agency Contracts," Working Paper, University of Western Ontario.
- MILGROM, PAUL (1981): "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 12, 380-391.
- MILGROM, PAUL, AND ROBERT WEBER (1985): "Distributional Strategies for Games with Incomplete Information," *Mathematics of Operations Research*.
- MIRRLEES, JAMES (1974): "Notes on Welfare Economics, Information, and Uncertainty," in *Essays on Economic Behavior Under Uncertainty*, ed. by M. Balch, D. McFadden and Shih-Yen Wu. Amsterdam: North-Holland Publishing Co.
- (1971): "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 38, 175-208.
- ROGERSON, WILLIAM (1984): "The First-Order Approach to Principal-Agent Problems," *Econometrica*, 53, 1357-1367.
- SHAVELL, STEVEN (1979): "Risk Sharing and Incentives in the Principal and Agent Relationship," *Bell Journal of Economics*, 10, 55-73.
- TWAIN, MARK (1985): *The Adventures of Huckleberry Finn*. New York: Doubleday and Company.