1984

New Directions in the Economic Theory of Agency

Glenn M. MacDonald

Follow this and additional works at: <https://ir.lib.uwo.ca/economicsresrpt>

Part of the [Economics Commons](#)

Citation of this paper:

MacDonald, Glenn M.. "New Directions in the Economic Theory of Agency." Department of Economics Research Reports, 8409. London, ON: Department of Economics, University of Western Ontario (1984).

10422

ISSN: 0318-725X

ISBN: 0-7714-0536-7

RESEARCH REPORT 8409

NEW DIRECTIONS IN THE ECONOMIC
THEORY OF AGENCY*

by

Glenn M. MacDonald

Department of Economics
University of Western Ontario
London Canada

and

Economics Research Center/NORC
University of Chicago

April, 1984

* I am both pleased and fortunate to be able to acknowledge input from Lorne Carmichael, Jim Heckman, Bengt Holmstrom, Ig Horstmann, Mike Hoy, Harvey King, Peter Kuhn, Ed Lazear, Tracy Lewis, Barbara Mann, Preston McAfee, John McMillan, Michael Parkin, Chris Robinson, Sherwin Rosen, Bob Rosenthal, Glen Stirling, Alan Slivinski, Linda Welling and Richard Zeckhauser.

I. Introduction

This essay is concerned with the economic theory of agency. It seeks to provide an accessible, though not heavily bibliographic, survey of existing results in the form of a rational reconstruction, and argue for pushing the analysis in several directions.

The theory of agency is one part of a broad research program on problems involving asymmetric information. It focuses on utilization of compensation rules with which one player, the principal, seeks to motivate another (or others), the agent, to choose his activities in a way advantageous to the principal. The problem is nontrivial because (i) the agent's activities are known only to himself and there is no immediate incentive for him truthfully to reveal them to the principal; and (ii) the obvious solution involving the assignment of some kind of full residual claimancy to the agent, is suboptimal because the fruit of the agent's efforts is in part stochastically determined, and the agent is risk averse. In this setting there is thus a fundamental tension between risk sharing and considerations of production efficiency.

Why study the economics of agency? First of all, symmetric information is in many cases demonstrably counterfactual. If symmetric information is necessary for the results of the existing models, then they are nonrobust in an important way. This possibility is particularly important for welfare analysis. Second, while symmetric information models are empirically rich, they are either silent on, or rule out, many observed phenomena. Examples are deductibles in insurance, dominance of time rates and piece rates as methods of compensating labor, precommitment to mandatory retirement, academic tenure, executive bonuses, and so on. Put simply, in standard symmetric

information settings with complete markets, there is no reason for one agent to choose a method of compensating another apart from "payment if and only if the goods are received". Yet much of the way in which transactions are structured has, at the very least, the appearance of involving much more, particularly in the sense of seeming to provide motivation for appropriate behavior.

Though the agency framework is quite broad--spanning the continuum from Becker's "rotten kid theorem" to the economics of health--attention will be restricted to agency problems in the context of the firm. That some such restriction is necessary is obvious, and the firm provides the broadest range of applications and probably the most fertile ground for further work.

The paper is comprised of three parts. Section II provides some background on the emergence of agency-type ideas. Section III lays out the agency problem, first focusing on the one period, single agent problem, then proceeding to the multiple agent and/or dynamic frameworks.

Section IV is comprised of suggestions about the directions in which research on this important topic might proceed. For reasons spelled out below, greatest attention is devoted to one period models. The primary orientation is scientific--towards raising the theory's predictive content through suitable enrichment of the basic agency model. Attention is not confined to results which the author can already prove, nor is there any attempt to inundate the reader with every possible extension. Rather, a moderate number of directions are examined, and though there are some safe bets, most ideas are intuitive and conjectural.

The paper closes with a brief conclusion.

Before proceeding, it should be pointed out that in order to provide a common set of assumptions with which to discuss the literature, some compromises are inevitable. In particular, there are several assumptions which are central to the results, and the manner in which they are handled in the various articles is stated, but the role of lesser assumptions is not dealt with in detail, and reference to the original is crucial for an entirely precise representation of the propositions.

II. Background

Compensation rules are not a new topic in economics--Adam Smith was likely not the first economist to argue that sharecropping arrangements would yield insufficient outcomes. Marshall's (1920) discussion of the English Metayer (sharing output) system and Plant's (1932) work on "control" in business are other early examples. This body of work was finished by Johnson (1950) and Cheung (1968), wherein it was shown that in a nonstochastic environment and under free entry, so long as production technologies are independent (i.e. there are no externalities), sharecropping and related compensation rules yield efficient outcomes. Under these assumptions there is no basis upon which to choose between the various compensation rules.

The second-coming for the study of compensation rules had to await the full development of the Arrow-Debreu-McKenzie (ADM) general equilibrium model.

In addition to providing a demonstrably self-consistent framework for the analysis of economic problems under competitive conditions, within the ADM setup the two fundamental theorems of Welfare Economics were established.

And all of this could be easily extended to situations involving uncertainty provided complete markets are available. Even Adam Smith probably did not expect things would turn out so nicely.

From the standpoint of practical significance, that the general competitive equilibrium framework, accompanied by its strong welfare implications, applies under uncertainty is of obvious import. Attention therefore focused on the major prerequisite for the theory to hold under such situations--complete markets. Is it reasonable to assume that if there are gains from trade involving contingencies, such trades will be carried out, or at least that the costs of doing so stand in the same relation to the gains from trade as they are assumed to when contingencies are not involved? One key issue is, of course, verifiability. Needless to say, if it cannot be agreed that some event has or has not occurred, trading on the basis of whether it has cannot occur [Hurwicz (1971)].

So where do the agency ideas fit in? A particular but common situation in which difficulties associated with verifiability are of importance occurs when individuals are endowed with different information and it is not to their individual advantage truthfully to reveal what they know. These asymmetric information problems come in numerous varieties, depending on whether the informational asymmetry involves the constraints individuals face, the actions they take, or the consequences of those actions in conjunction with the constraints. The agency problem can contain elements of all three, but fundamental to it is that some individual actions are not observable by other agents.¹

There are basically three strands of thought (not all equally well worked out) concerning what economic agents might do under such circumstances. One examines activities culminating in direct acquisition of information by uninformed parties--monitoring-type ideas. Another focuses on rearranging the pattern of allocation so that informational asymmetries have a smaller impact--optimal assignment problems are an example. The last analyzes the possibilities for designing methods of exchange based on observables which have (perhaps incidentally) the effect of rendering incentives more compatible, and hence lessening the destructive impact of informational asymmetries. The agency ideas can be viewed as being of this kind. In the agency framework, those "methods of exchange" are the compensation rules referred to above, which describe the distribution of output as a function of the observables.

In brief then, the modern version of the economics of agency arose as part of a general research program on asymmetric information. The focus of attention is the design of compensation rules when individual actions are not observable and basic incentives are in conflict.²

III. Overview of the Theory of Agency

1. The Basic Model

The early papers in the theory of agency [Spence and Zeckhauser (1971), Ross (1973), and Stiglitz (1974, 1975)] and much of the later work focus on a two-player (principal P, and agent A) situation.³ A's action α , together with a random state of nature Ω , where $\Pr(\Omega \leq \omega) \equiv G(\omega)$, are assumed to generate output Q according to

$$Q = f(\alpha, \Omega) \tag{1}$$

where $f_1 > 0$, $f_{12} \neq 0$ and $f_{11} < 0$. Realized output $q = f(\alpha, \omega)$ is to be divided between P and A, with P choosing the method of allocation. α is chosen by A, his choice being made before or after observing $\Omega = \omega$ depending on the model. When α is determined prior to (after) A's learning $\Omega = \omega$, the model will be called type I (II).⁴ In either case, it is too costly for P to observe anything beyond $Q = q$, complete information on which is freely available to all. Consequently, the payment from P to A, denoted $y(\cdot)$, can be a function of output alone-- $y(q)$. In particular, $y(\cdot)$ cannot depend on α or ω .

Both P and A have preferences which may be characterized as von Neuman-Morgenstern utility functions, $U = u(\cdot)$ and $V = v(\cdot)$ respectively. $u(\cdot)$ depends only on income, while $v(\cdot)$ depends on income and α ($v_{\alpha} < 0$); both are weakly concave in income.

The focus of the agency problem is on the compensation rule $y(q)$. There are several ways in which $y(q)$ can be determined. A particularly straightforward one is the partial equilibrium approach of assuming $y(q)$ maximizes P's expected utility given that A must be able to achieve some specified level of expected utility $\bar{V}$ when he chooses α optimally. The $y(q)$ so obtained is thus ex ante Pareto optimal in the class of compensation rules which are functions of q alone; equivalently, $y(q)$ is constrained Pareto optimal given the information structure.

To proceed formally, notice that for any α , Q is a random variable.⁵ Let $\Pr[Q \leq q | \alpha] \equiv H(q; \alpha)$, and $H_q \equiv h(q; \alpha)$ be the associated density; that α is productive implies $H_{\alpha} < 0$. It is assumed that the support of $h(q; \alpha)$, the closure of $\{q | h(q; \alpha) > 0\}$, does not depend on α . If the contrary is assumed, a very strong restriction has been imposed, for then effort can effectively be observed for some output levels. To illustrate, suppose

$\alpha^0 < \alpha^1$, where $H(\bar{q}|\alpha^1) = 0$ for some $\bar{q}$ and $H(\bar{q}|\alpha^0) > 0$. Then, by imposing an appropriate penalty if $q < \bar{q}$, P can ensure $\alpha = \alpha^1$ is chosen. The assumption that the support of $h(q, \alpha)$ depends on α is a very strong one and is instrumental in generating penalty contracts of the type introduced by Mirrlees (1974).⁶ In a sense, such schemes rely on there being no agency problem for some levels of effort.

The formal statement of the agency problem is

$$\max_{y(q)} \int u[q - y(q)] h(q; \hat{\alpha}) dq \quad (2)$$

$$\text{S.T. } \hat{\alpha} = \underset{\alpha}{\operatorname{argmax}} \int v[y(q), \alpha] h(q; \alpha) dq \quad (3)$$

$$\text{and } \int v[y(q), \hat{\alpha}] h(q, \hat{\alpha}) dq \geq \bar{V}. \quad (4)$$

P seeks to maximize the expression (2)--simply the expected utility of income $Q - y(Q)$. In all expected utility calculations, the probability distribution of Q is $h(\cdot)$, which depends on A's choice of α , $\hat{\alpha}$. The manner in which $\hat{\alpha}$ is chosen is one of the constraints faced by P. (3) states that $\hat{\alpha}$ maximizes A's expected utility when he is confronted with the compensation rule $y(q)$; loosely, $\hat{\alpha}$ is a function of $y(q)$.⁷ The final constraint, (4), requires that when A chooses α optimally, his expected utility is at least that available elsewhere, $\bar{V}$.

At present, essentially nothing is known about the form of the maximal $y(q)$. So little in the way of the results is available for this setup for two reasons. First, as written in (4), A's preferences over income lotteries may interact with his choice of effort in an arbitrary and complex fashion. It is not difficult to imagine that showing anything about $y(q)$ may be next to impossible when variations in $y(q)$ have significantly different

effects on A depending on the level of α . Second, the constraint (3) is really a continuum of inequality constraints: (assuming $0 \leq \alpha \leq \bar{\alpha}$ is the feasible set)

$$\forall \alpha \in [0, \bar{\alpha}]: \int v[y(q), \hat{\alpha}] h(q; \hat{\alpha}) dq \geq \int v[y(q), \alpha] h(q; \alpha) dq$$

Loosely, without further restrictions $\hat{\alpha}$ may be very badly behaved as $y(q)$ varies. This possibility presents a technical problem of a fairly unpleasant variety.

Some results do become available if A's preferences are restricted so that his attitudes towards income risk do not depend on α . This separability is accomplished by requiring $v(\cdot)$ to be of the form

$$v(y, \alpha) = -v_0(\alpha) + v_1(\alpha)v_2(y) \quad (5)$$

where $v_1 > 0$ and v_2 is monotone increasing and weakly concave.⁸ When $v_2'' \equiv 0$, in which case A is risk neutral for all income levels, it is readily shown [Harris and Raviv (1979)] that P optimally receives a quantity independent of output, and hence that $y' = 1$. More generally, Grossman and Hart (1983) have demonstrated (for discrete Ω) that while $y' < 0$ or $y' > 1$ can hold for some values of q under an optimal compensation rule, there must be some range of output for which neither holds, in which case $0 < y' < 1$. Put differently, since P and A are sharing q , an increase in q must raise one income; $1 - y' > 0$ or $y' > 0$. The result is that for some q , both incomes rise with q ; $1 - y' > 0$ and $y' > 0$. That $y' < 0$ is a possibility is perhaps surprising. However it is simply a consequence of the fact that $h(\cdot)$ can be such that intermediate levels of output are a strong indication of low

effort; $y' < 0$ for some interval then encourages A to choose greater α in hopes of avoiding realized output falling in the intermediate range.

Further results can be obtained when it is appropriate to characterize the choice of $\hat{\alpha}$ in (3) by a first-order necessary condition for the problem

$$\max_{\alpha} \{-v_0(\alpha) + \int v_2[y(q)]h(q;\alpha)dq\},$$

where (5) has been imposed and $v_1 \equiv \text{unity added}$. The first-order condition so obtained is

$$\int v_2[y(q)]h_{\alpha}[q;\alpha]dq = v'_0(\alpha). \quad (6)$$

Mirrlees (1975), first pointed out that (6) is not in general necessary for (3). Correct treatment of the conditions sufficient for the validity of (6) are supplied by Rogerson (1983).⁹ Specifically, define new units of action $\tilde{\alpha} \equiv v_0(\alpha)$, and let $P[Q \leq q | \tilde{\alpha}] \equiv \tilde{H}(q; \tilde{\alpha})$ with associated density $\tilde{h}(q; \tilde{\alpha})$.

The sufficient conditions are $\frac{\partial}{\partial q} \frac{\tilde{h}_{\tilde{\alpha}}}{\tilde{h}} \geq 0$ (implying $H_{\alpha} \leq 0$) and $H_{\alpha\alpha} \geq 0$.

Under these restrictions, increments to $\tilde{\alpha}$ raise output (stochastically) at a decreasing rate; diminishing marginal product in the appropriate stochastic sense. The analytical consequence is a guarantee that $\hat{\alpha}$ does not vary discontinuously with small perturbations in $y(q)$.¹⁰

Under (6), the determination of $y(q)$ can be treated as an isoperimetric optimal control problem [see Kamien and Schwartz (1981)]. Letting λ and μ be multipliers on (4) and (6) respectively, the Hamiltonian is

$$\mathcal{H} = \int \{u[q - y(q)] + \lambda[-v_0(\alpha) + v_2(y) - \bar{V}] + \mu\{v_2[y(q)]\frac{h_{\alpha}(q,\alpha)}{h(q;\alpha)} - v'_0(\alpha)\}\}h(q;\alpha)dq.$$

Since y' does not appear in $\mathcal{H}$ explicitly, when the optimal $y(q)$ is in the

interior of the set of feasible compensation rules it can be characterized as the solution to the necessary condition $\partial V / \partial y(q) = 0$, which may be rearranged to yield

$$\frac{u' [q-y(q)]}{v_2' [y(q)]} = \lambda + \mu \frac{h_\alpha(q, \hat{\alpha})}{h(q, \hat{\alpha})} . \quad (7)$$

(7), in conjunction with (4) and (6), yields $y(q)$, λ and μ .

Proceeding thus, Shavell (1975) showed that both P and A will invariably bear some of the risk ($y' \neq 0$ or 1) while $0 < y' \leq 1$ for all q was obtained by Harris and Raviv. Further, from P's point of view A always chooses too low a level of α [Holmstrom (1979)].

Two further points are worthy of note. First, an optimal risk-sharing agreement between P and A entails choosing $y(q)$ so that $u' [\cdot] / v_2' [\cdot] = \text{constant}$ [Borch (1962)]. Referring back to (7), the agency framework will not achieve this "first best outcome" unless $\mu = 0$ or $h_\alpha = 0$; that is, unless P and A agree on what α should be chosen (at the optimum choice of $y(q)$ there is no incentive problem) or the density of Q does not depend on α .

Second, this is a good place at which to gain some intuition about the way in which $y(q)$ is manipulated to encourage increased effort, and just how this incentive is related to optimal risk sharing. (7) indicates that $y(q)$ should be higher (lower) than optimal risk sharing would require when $h_\alpha > 0$ (< 0). Let (i) $\bar{y}(q)$ denote the optimal risk sharing compensation rule, (ii) $x(q) \equiv y(q) - \bar{y}(q)$, (iii) $Q^+ = \{q | h_\alpha(q, \hat{\alpha}) > 0\}$, and (iv) $Q^- = \{q | h_\alpha(q, \hat{\alpha}) \leq 0\}$. Then $x(q) > 0$ is equivalent to $q \in Q^+$. (6) can be written

$$\int_{Q^+} v_2 [x(q) + \bar{y}(q)] h_\alpha(q, \hat{\alpha}) dq + \int_{Q^-} v_2 [x(q) + \bar{y}(q)] h_\alpha(q, \hat{\alpha}) dq - v_0' = 0,$$

from which it is immediately clear that the purpose of P's setting $x(q) \neq 0$ is to raise the marginal return to α .

2. Two-Player Extensions of the Basic Model

Given the basic agency model, numerous extensions were forthcoming. It is useful to consider first those extensions that retain the two player, one period setup. The situation with three or more players, and/or more than one period, differs in a more fundamental sense in that there is scope for output comparisons either among distinct agents or for a given agent over time.

As mentioned above, monitoring arrangements are a natural response to informational asymmetries. In the agency context, an important issue involves the conditions under which P's obtaining extra information will lead to a Pareto improvement.¹¹

Harris and Raviv (without (5) and (6)) show that A being risk averse is necessary and sufficient for full observability of α to lead to a Pareto improvement, suggesting potential gains to monitoring. Holmstrom (1979) then considered (using (5) and (6)) a situation wherein a random variable X , freely observable by both parties, is jointly distributed with Q for each α . The compensation rule can then be based on $X=x$ as well as q : $y(q,x)$. Letting $h(q,x;\alpha)$ denote the joint density of X and Q , (7) again emerges with $h_\alpha(q,x;\alpha)/h(q,x;\alpha)$ replacing $h_\alpha(q;\alpha)/h(q;\alpha)$.

It is immediate that a condition necessary and sufficient for observation of $X=x$ to be useless is that h_α/h not depend on $X=x$. The class of functions h for which this occurs is

$$h(q,x;\alpha) = h_1(q;x)h_2(q;\alpha) \quad (8)$$

in which case given q , the density of X does not depend on α . Provided (8) fails, $X=x$ is an "informative monitor" and the compensation rule $y(q,x)$ Pareto-dominates $y(q)$; put differently, there will be positive demand for monitoring arrangements.

This line of argument was examined further (in a much more general informational setting) by Gjesdal (1982). Failure of (8) is closely related to the notion of Blackwell informativeness of $X=x$ for inference about α [Marschak (1971)]. Essentially α is treated as a parameter about which inferences are to be made, at least in the sense of P 's making inferences given q . However, it must be recalled that P 's goal is to maximize his expected utility given some prescribed expected utility for A ; inferences about α are useful but incidental. In fact, it turns out that failure of (8) is no longer necessary for observation of $X=x$ to be useful when (5) does not hold, though it remains sufficient. Spurious randomness may help! The reason is just that when preferences are nonseparable and P would prefer A to raise α , basing y on $X=x$ as well as $Q=q$ raises α if the marginal return to α is convex in x , as it can be under nonseparable $v(\cdot)$. The main point here is that information theorems from decision theory, where agents confront nature, do not carry over straightforwardly to the strategic framework of the agency problem.¹²

Finally, Harris and Raviv demonstrated (using (6)) that when α itself is observable with error, say $X=\alpha+\epsilon$, where ϵ is stochastically independent of Ω , an optimal compensation rule will be of the form

$$y(q,x) = \begin{cases} y_0(q,x) & x > \hat{x} \\ \bar{y} & x \leq \hat{x} \end{cases}$$

where $y_0(\cdot)$ is determined in a fashion similar to the above, and $\bar{y} < 0$ and $\hat{x}$ are constants. An information structure of the type assumed essentially allows direct penalties for low effort. It should be noted however that the specification $X = \alpha + \epsilon$ implies (unless the support of the density of ϵ is unbounded in both directions) the support of the density of X depends on α , in which case, as discussed above, a penalty scheme of some variety will be effective. More general specifications of the stochastic relationship between X , Q and α will typically not have this property.

Another noteworthy extension is provided by Lal and Staelin (1982). They analyze (among other things) the case in which one agent produces two goods. Except under obvious and extremely implausible separabilities, the optimal compensation rule depends non-additively on both outputs. As a practical matter, agents do perform more than one activity. That compensation is expected to depend on both goods in a nontrivial fashion is a coarse refutable hypothesis.

3. Multiple-Agent and Dynamic Extensions of the Basic Model

Multiple-agent and dynamic versions of the agency problem differ fundamentally from the framework studied above. In such settings compensation can be based on entities apart from A 's current output; output comparisons across agents and/or over time are feasible. Moreover, the nature of strategic interaction among agents becomes an important facet of the problem. These additional features render the agency model both more interesting and more difficult. Single-period-multiple-agent and dynamic-single-agent models are currently the subject of considerable research activity. Dynamic-multiple-agent models are only now being developed.

(i) Single-Period-Multiple Agent Models

The single-period-multiple-agent work is the best developed of the three types of extended models, as well as most closely related to the single period analysis. Also, setting out the dynamic models requires significant accumulation of both new concepts and notation. Consequently, attention is primarily confined to the single period multi-agent material, and as for dynamics, only the basic thrust of the models is presented.

In the multiple-agent extensions, P faces n agents $A_i (i=1, \dots, n)$. The models which use this structure can usefully be thought of as falling into three categories. One uses the agency framework to provide insight into why entrepreneurial firms exist; a second recognizes the inherent difficulty of examining fully optimal compensation rules and responds by seeking the optimal rule from a restricted class; and the third approaches the same difficulty by searching for general results through explicitly modeling the problem as a game of incomplete information, perhaps with communication. This section provides an overview of the first two of these approaches. Game theory is taken up in Section IV.

The most important papers concerned with the entrepreneurial firm are by Alchian and Demsetz (1972), which raised some of the basic issues, and Holmstrom (1982), which provided a correct analysis of them.

The central assumption of the entrepreneurial firm problem is that groups of agents form because of production complementarities; a group can produce more than the sum of what each person could produce separately. But in a group, these gains may be difficult to achieve when there is asymmetric information. To see why, assume production is a nonstochastic function of

the efforts (α_i) of the n agents

$$q = f(\alpha_1, \dots, \alpha_n)$$

and that α_i is observable only by agent i , all observing q freely; let $\partial f / \partial \alpha_i \equiv f_i$.¹³ For simplicity, suppose agents' preferences can be written as a special case of (5): $v(y, \alpha) = y - v_{oi}(\alpha_i)$, and that nonparticipation yields $V=0$. Then a Pareto-optimal effort choice for agent i , α_i^* , satisfies $\alpha_i^* \in \operatorname{argmax}_{\alpha_i} q - \sum_i v_{oi}(\alpha_i)$, which (assuming interior α_i^*) implies $f_i = v'_{oi}$, and yields $q^* \equiv f(\alpha_1^*, \dots, \alpha_n^*)$. If there is no principal, and output is shared among the agents according to some differentiable compensation rules $y_i(q)$, where $\sum y_i \equiv q$ (i.e. the shares sum to one), the Nash equilibrium effort levels $\hat{\alpha}_i$ satisfy $y'_i f_i = v'_{oi}$, in which case $\alpha_i^* = \hat{\alpha}_i$ for all i requires $y'_i = 1$ for all i , violating $\sum_i y'_i \equiv 1$ (implied by $\sum_i y_i \equiv q$). If all output is to be distributed among the n agents, none can have the correct marginal incentives.

The difficulty is that $\sum y_i \equiv q$ is an identity. It is easy to see that if $\sum y_i \leq q$ were permitted, Pareto-optimality is easily achievable. Indeed let

$$y_i(q) = \begin{cases} \bar{y}_i & q = f(\alpha_1^*, \dots, \alpha_n^*) \\ 0 & q \neq f(\alpha_1^*, \dots, \alpha_n^*), \end{cases}$$

for $\bar{y}_i > v_{oi}(\alpha_i^*)$, and $\sum \bar{y}_i = q^*$. That is, q is destroyed if it does not take on its Pareto-optimal level. Such penalty contracts (usually called "forcing" in this context) are feasible and generate $\alpha_i = \alpha_i^*$.¹⁴ However, the problem is, of course, that the self-imposed (by the group) threat to destroy the output is

not credible. If $q < q^*$, will the group go ahead and destroy it? That it will not is logically compelling.

P enters the picture at this point. Suppose P chooses the forcing rule having $\bar{y}_i = v_{oi}(\alpha_i^*) + \epsilon$ for some small positive ϵ , and takes on the role of residual claimant.¹⁵ P earns q if $q \neq q^*$ and $q - \sum_i \bar{y}_i$ otherwise. Provided $q - \sum_i \bar{y}_i$ does not fall short of his best alternative, P will be willing to participate and $q = q^*$ follows. His threat to keep the output is entirely credible. Holmstrom's key insight (novel in the agency problem), in contrast to Alchian and Demsetz' emphasis on monitoring, is that P's participation renders $\sum y_i(q) < q$ a credible threat. This gives the entrepreneurial firm a distinct advantage over pure teams, and thus contributes to understanding why it exists. Further, this argument provides an important rationale for studying multi-agent problems--there is an efficiency-based rationale for having a principal in the problem from the outset.

The second body of work on multi-agent settings concentrates on the fundamental difference between single and multiple agent problems--namely that when there is more than one agent, they may be compared to one another. This is critical for two reasons. First, interactions among the agents become important, and cause the problem to change much in the same way that the potential for entry alters monopoly problems. Second, a richer menu of possible compensation rules become available.

As noted above, in an entirely deterministic setting the analysis of optimal compensation rules is straightforward--forcing rules (among other things) are optimal. But when production is subject to stochastic influences, the problem of characterizing optimal compensation rules, nontrivial with just

one agent, becomes horrific.

The "tournaments" literature grew out of responding to the inherent difficulty of characterizing optimal compensation rules by restricting the space of feasible rules. The central works are Lazear and Rosen (1981), Nalebuff and Stiglitz (1983) and Green and Stokey (1983).

The analysis begins by specializing the structure of production to the form

$$Q = \sum_{i=1}^n Q_i,$$

where $Q_i = f^i(\alpha_i, \Omega_i, \bar{\Omega})$, the realization $Q_i = q_i$ is observed by P, and the random vector $(\Omega_1, \dots, \Omega_n, \bar{\Omega})$ has distribution function $G(\omega_1, \dots, \omega_n, \bar{\omega})$.

Ω_i is the A_i -specific or idiosyncratic error, and $\bar{\Omega}$ is a common shock.

Assuming production to be additively separable may seem to dispense with the underlying gains to team production which presumably generate multi-agent firms in the first place, and this perception is correct to some extent.

However, as will become clear, the central requirements for the arguments to follow are that there be some observable entity related to, but not equal to, α_i for all A_i -- Q_i serves this purpose--and that the Q_i be somehow interrelated. That the Ω_i are not assumed stochastically independent provides the required link between the A_i .¹⁶ When the Ω_i are independently distributed, allowing the i^{th} agent's compensation $y_i(\cdot)$ to depend on anything apart from q_i merely adds additional randomness to A_i 's income stream. Under (5), which is assumed, this scheme is unambiguously Pareto-dominated by a purely "individualistic" rule $y_i(q_i)$; Interdependent schemes (i.e., $y_i(q_1, \dots, q_n)$) are ruled out [Holmstrom (1982)].

As indicated above, the central feature of the tournaments literature is that it begins by placing restrictions on the feasible $y_i(\cdot)$. In particular, attention is confined to purely individualistic compensation rules $y_i^I(q_i)$, and "rank order" rules of the form

$$y_i^T(q) = \begin{cases} \bar{y}_1 & \text{if } q_i = r_1 \\ \bar{y}_2 & \text{if } q_i = r_2 \\ . & \\ . & \\ . & \\ \bar{y}_n & \text{if } q_i = r_n, \end{cases}$$

where the elements of $\bar{y} \equiv (\bar{y}_1, \dots, \bar{y}_n)$ are constants satisfying $\bar{y}_1 > \dots > \bar{y}_n$, and r_j is the j^{th} order statistic of the "sample" $q \equiv (q_1, \dots, q_n)$.^{17,18,19}

Given that the compensation rule is restricted to be of the form $y_i^I(q_i)$ (indeed the additional restriction $y_i^I = a + bq_i$ for some constants a and b , a "linear piece rate", is frequently imposed) or $y_i^T(q)$, the issue of what characterizes an optimal compensation rule comes down to which of y_i^I or y_i^T will be chosen under various parameter values.

Under (9), the analysis of $y_i^I(q_i)$ is little different from that presented above for the single agent case since A_i chooses α_i taking α_j ($j \neq i$) as given.

The analysis of y_i^T is not as yet worked out in a general model, and the major papers do not assume a common format. What follows is most closely related to Green and Stokey, primarily because their formulation follows most closely from the models discussed above. Fortunately, the basic results do not appear to be particularly sensitive to which framework is used.

The major results are available for a homogeneous contest: $f^i = f$ for all i , and all A_i are identical. In addition to restricting the A_i 's tastes to be of the form (5), it is assumed that P is risk neutral and that $f(\cdot)$ can be written

$$f(\alpha_i, \Omega_i, \bar{\Omega}) = g(\alpha_i, \Omega_i) + \bar{\Omega};$$

that is, production is subject to an additive common shock $\bar{\Omega}$ and an agent A_i specific idiosyncratic disturbance Ω_i .

For any given vector of effort levels $\alpha \equiv (\alpha_1, \dots, \alpha_n)$, $\Pr(Q_i = r_j | \alpha)$, the probability that Q_i is the j^{th} order statistic, can be computed. Using $\sum_j \Pr(Q_i = r_j | \alpha) = 1$, A_i 's expected utility is ($v_1 \equiv 1$)

$$E(V) = v_2(\bar{y}_n) + \left[\sum_{j=1}^{n-1} [v_2(\bar{y}_j) - v_2(\bar{y}_n)] \Pr(Q_i = r_j | \alpha) \right] - v_0(\alpha_i).$$

A utility level of $v_2(\bar{y}_n) - v_0(\alpha_i)$ is the worst A_i can expect. Increments to this, of size $v_2(\bar{y}_j) - v_2(\bar{y}_n)$ are obtained by producing output greater than that of $n-j$ others.

The A_i are assumed to choose α_i given effort choices made by the other A_i and the compensation rule. When interior, this "best response", $\hat{\alpha}_i$, solves the first-order condition for maximization of $E(V)$ with respect to α_i :

$$\sum_{j=1}^{n-1} [v_2(\bar{y}_j) - v_2(\bar{y}_n)] \frac{\partial \Pr(Q_i = r_j | \alpha)}{\partial \alpha_j} = v'_0(\alpha_i); \quad i=1, \dots, n \quad (10)$$

Though "nonconvexities" are usually of no interest, they become important in the tournament framework in two ways. To illustrate, if Ω_i and $\bar{\Omega}$ have degenerate distributions, given fixed α_j (taking all α_j equal to some $\bar{\alpha}$ for simplicity; $j \neq i$) A_i can guarantee himself $\bar{y}_1$ by choosing $\hat{\alpha}_i$

only slightly in excess of $\bar{\alpha}$; a discrete gain at infinitesimal cost. But given this $\hat{\alpha}_i$, some A_j can achieve the same end by raising $\hat{\alpha}_j$ only slightly above $\hat{\alpha}_i$, and so on in Bertrand-style. The nonexistence of optimal $\hat{\alpha}$ follows. To avoid this problem the random components in the problem must be, loosely speaking, substantial.

The other noteworthy nonconvexity problem arises if $\bar{y}_n$ is sufficiently high that the A_i may wish to "coast". No effort at all guarantees $\bar{y}_n$, and this package may be preferable to trying to win a higher prize.

Provided the nonconvexities are handled properly, the solution of equations (10) yields a unique symmetric Nash equilibrium $\hat{\alpha}_i = \hat{\alpha}_i(\bar{y})$, in which case $\Pr[Q_i = r_j | \hat{\alpha}] = \frac{1}{n}$ for all j . P then chooses $\bar{y}$ to maximize

$$E[Q_i | \alpha_i = \hat{\alpha}(\bar{y})] - \frac{1}{n} \sum_j \bar{y}_j.$$

Given this setup, Green and Stokey provide two major results. The first, anticipated by Lazear and Rosen, is that as the distribution of $\bar{\Omega}$ becomes less concentrated, $y_i^T(\cdot)$ will eventually dominate $y_i^I(\cdot)$. The logic is that increases in the potential variation of $\bar{\Omega}$ generate greater risk for A_i under $y_i^I(\cdot)$, hence reducing its value to the A_i . But since the payment received under y_i^T depends on the difference in outputs (i.e., $q_i \gtrless q_j \Leftrightarrow g(\alpha_i, \omega_i) \gtrless g(\alpha_j, \omega_j)$, which does not depend on $\bar{\omega}$, the value of y_i^T to the A_i does not diminish as $\bar{\Omega}$ becomes more variable. (Recall P is risk neutral.)²⁰

The second result is that tournaments will dominate as the number of agents grows large. Though the formal argument is arduous, an intuitive explanation is as follows. The principal advantage of y_i^T relative to y_i^I is

that y_i^T is flexible in the sense that it works equally well as an incentive mechanism for any realization of $\bar{\Omega}$. Essentially, the compensation rule has a shifting base. But y_i^T 's shortcoming is that it is a blunt tool. Indeed, when $n=2$ there are just two parameters, $\bar{y}_1$ and $\bar{y}_2$, which P can use to manipulate the α_i , whereas under the individualistic scheme the whole function $y_i^I(\cdot)$ can be thought of as continuum many controllable parameters. As n grows, under the tournament compensation rule, A_i (taking the other agents' actions as given) faces a "smoother" compensation rule--the costly bluntness of the compensation rule is alleviated.²¹

Two other interesting results, unique to the homogeneous tournament setting, are shown by Nalebuff and Stiglitz. First, it may be optimal for y_i^T to contain "gaps". That is, under some circumstances P can raise his expected utility without lowering that of any A_i by (focusing on $n=2$ purely for simplicity) choosing a constant $\gamma > 0$ such that

$$y_i^T = \begin{cases} \begin{cases} \bar{y}_1 & \text{if } q_i = r_1 \\ \bar{y}_2 & \text{if } q_i = r_2 \end{cases} & \text{and } r_1 - r_2 \geq \gamma \\ \bar{y} & \text{if } r_1 - r_2 < \gamma \end{cases}$$

The intuition is that for $\gamma > 0$, there are really three prizes, $\bar{y}_1$, $\bar{y}$ and $\bar{y}_2$, providing a smoother yet still flexible compensation rule.

The second result is that the prizes $(\bar{y}_1, \dots, \bar{y}_n)$ need not be distinct. Indeed Nalebuff and Stiglitz provide an example where for any n , $\bar{y}_1 > \bar{y}_2 = \dots = \bar{y}_n$ is optimal. In general this type of result arises when $\partial \Pr(Q_i = r_j | \alpha) / \partial \alpha_i$, evaluated at the common Nash choice $\hat{\alpha}_1 = \hat{\alpha}_2 = \dots = \hat{\alpha}_n$,

does not depend on j for j in a strict subset of $\{1, \dots, n\}$. The additional degree of freedom to specify $\bar{y}_j \neq \bar{y}_{j'}$ ($j \neq j'$) is then not useful.

The homogenous tournament structure is amenable to numerous extensions. [See, for example, O'Keefe, Viscusi and Zeckhauser (1984).] A very promising class of extensions of the basic tournament model allows for heterogeneous A_i . Consider A_i who differ in that $v'_0(\alpha)$ is lower for some than others. Call workers for whom v'_0 is lower "more able". Lazear and Rosen showed that when A_i 's ability is not observed by P , the optimal y_i^T constructed on the assumption that the A_i will agreeably self-select into "leagues" will not in fact succeed in inducing self-selection. That it be individually rational for A_i to self-select places real restrictions on the collection of y_i^T from which P may choose. The spread in high ability leagues must be larger than in an efficient scheme, and high ability individuals work too hard. O'Keefe, Viscusi and Zeckhauser demonstrated that this inefficiency occurred because P 's ability to monitor (treating Ω as what goes unmonitored) was assumed exogenous. If P can choose to observe Ω more or less closely, self-selection can be efficiently induced. The essentials of the argument are straightforward when there are only two ability levels. Suppose P can construct a low ability tournament, having a small range of prizes and high accuracy of monitoring, and a high ability tournament with a wide range of prizes and smaller monitoring accuracy. Then a high ability agent will not seek to "slum" by entering the low ability tournament, since though he is very likely to win there, the prizes are relatively low (close to the average product of a low ability

worker). On the other hand, a low ability worker will not "climb" by entering the high ability tournament because (due to the wide range of prizes in the high ability tournament) he is likely to earn a low prize, far from the average product of a high ability worker. Efficiency is obtained because the freedom to choose monitoring quality as well as y_i^T allows maintenance of the right incentives to work.

A second useful extension is provided by Carmichael (1983). In his analysis risk neutral P and $n=2$ risk averse agents face the technology

$$Q_i = \ell(\alpha_i, \alpha^P) + \Omega_i$$

where α^P is the principal's effort, not observable by the A_i . Note that α^P is a public good. The important addition is that the link between the A_i which is necessary for interdependent compensation rules to provide a gain over purely individualistic rules is provided by a player in the game (P) rather than nature ($\bar{\Omega}$).

The main modification to the earlier results is that some of the incentives to work must be retained by P. The most illuminating conclusion is that as the number of agents rises, the compensation rule is, roughly speaking, full residual claimancy for P, with the A_i participating in a tournament. This result is partially due to the public good aspect of P's efforts (note also that $\bar{\Omega}$ operates in this fashion in the earlier work), which highlights the role of the existence of public goods within the firm in generating the dominance of interdependent compensation schemes as the number of agents grows. The other element operating in favor of interdependent schemes is that, loosely speaking, each A_i competes with the

average A_i . As n grows, the distribution of average output becomes more concentrated, hence lowering the risk faced by A_i and raising the value of an interdependent scheme.

(ii) Dynamic Models

Turning to single-agent dynamic models, the basic issue is whether agency problems wherein the agency relationship lasts over time, or has effects which persist, can generate more efficient outcomes.

Recall that single-period agency problems, like most noncooperative games, rarely generate fully efficient outcomes [see Dubey (1983)]. Put simply, when one agent chooses an action given that chosen by the other player, he does not take into account the implied change in the other player's payoff. Essentially there is an uninternalized externality, naturally leading to inefficiency. Allowing for more than one period allows one player's action to generate a future response by the other player. There is at least the possibility of a more efficient outcome.

There are three basic kinds of approaches to dynamic single-agent models. Radner (1981) [see also, Malueg (1981), Rogerson (1982), Rubenstein and Yaari (1983) and Lambert (1983)] shows that in a T -period repetition of the single period agency problem, a Pareto-optimal pair of actions $\hat{y}(q)$ and $\hat{\alpha}$ can be sustained by the strategy of permanently reverting to the one-period Nash action (y^* or α^*) next period should the other agent not choose the Pareto-efficient action in the present period. Obviously P , not being able to observe the α chosen by A , must have a method for deciding when $\alpha \neq \hat{\alpha}$ has occurred, and Radner shows that a simple decision rule (which looks much like a t -statistic for the hypothesis $\alpha = \hat{\alpha}$) does the trick.²²

Naturally this argument is subject to the standard "end game" difficulties, and this is essentially assumed away through the use of the ϵ -equilibrium notion (though it could have been dealt with in other ways; cf Lazear (1979)).²³ This notwithstanding, the point is that in the presence of agency-type informational asymmetries, rules of inference can be constructed which allow efficient outcomes to be sustained. An important limitation of this line of analysis is its reliance on low rates of discount. Since P punishes A by choosing single-period Nash outcomes in the future in response to current bad behavior, it is crucial that the future not be too heavily discounted by A.

A second dynamic approach is that of Becker and Stigler (1974). Therein (restating their argument in the terms used above) it is shown that A can be induced to behave efficiently by "underpaying" him early in the game and "overpaying" him later. The rules of the game are then that should A behave inefficiently at some point, he loses the future overpayment through being dismissed. It is shown that the payment scheme can be constructed so as to yield efficient behavior.²⁴ The end game problem is then solved by a terminal payment from P to A. A major implication of this work, in addition to the upward sloping (in time) compensation profile, arises from the observation that since A is overpaid as of the end of the game, he would prefer to continue it. It follows that a prespecified maximum duration of employment is necessary--mandatory retirement (Lazear).

Put simply, both the Radner and Becker-Stigler treatments of dynamic agency imply more efficient outcomes than do single-period models because the fact that the players interact more than once generates an incentive

for one player to take account of the effects of his actions on the other's payoffs.

Finally, Fama (1980) argues that competition alone will handle most agency problems if managers are concerned with the capital value of their reputations. Holmstrom (1983) shows that this proposition is correct under particular assumptions, but risk aversion and discounting the future generally render it false.

A very small amount of work on dynamic multiple-agent models is presently available. King (1974) analyzes a two-period tournament model in which winners (losers) of first-period, two-agent tournaments are paired off in second-period tournaments. Central to the model is an explicit treatment of the role of capital markets. Typical of the kind of result available is that when agents can only borrow what they can conceivably pay back, contests are linked over time in a nontrivial way even when there is no capital within the firm; second-period winners' contests differ systematically from losers' contests, as do the effort levels chosen.

IV. New Directions

That there is potential for the theory of agency to contribute to understanding the way firms operate is practically beyond question. Throughout the continuum of intra-firm activities, the possibility for discretionary action under diverse incentives and imperfect observability undoubtedly exists.

Yet the theory of agency surveyed above is not a very compelling one. An ungenerous, but not entirely inaccurate, view is that while the theory

correctly identifies the major issues, it applies well only to a few kinds of salesmen, and then predicts little beyond that their compensation will generally depend positively on their sales and sometimes on the sales of others. Naturally, theory involves simplification, but it would appear that this process has gone much too far in this case. The general direction of the comments to follow is therefore towards enrichment of the theory with a view to raising its predictive content.

It is probably worth mentioning why the theory of agency has so little positive content in its present form. Recall that it grew out of checking the robustness of earlier results against relaxation of the complete markets assumption. To provide a demonstration that a result fails to hold requires an example rather than a general model; hence the somewhat rarefied setting which became the basis for the literature to follow.

The extensions suggested below can be divided into three groups: those related to the influence of external markets on the agency framework; those involving enrichment of the agency setting itself; and those which are primarily modelling issues which could use extension or clarification. The focus is primarily on single-period models. This choice, dictated partly by a desire to prevent the discussion from ranging too widely, is also a result of the notion, fairly strongly suggested by the dynamic literature, that agency issues are most crucial over short periods. Of course, obtaining a clear general statement of this idea is itself an important issue, and work in this direction is underway.

1. External Markets

In the agency framework presented above, the impact of markets external to the firm was entirely summarized by $\bar{V}$.

There are at least two important ways in which such markets might influence the agency setting which are not captured in this specification. The first is through self-selection of the players in the game.²⁵ As modelled, P and A have already decided to participate in some joint venture. Since much of the indeterminacy in the theory arises because few restrictions are placed on $u(\cdot)$ and $v(\cdot)$, imposition of the requirement that the decision to become P or A be rational may be very useful because it can place strong restrictions on the relationship between $u(\cdot)$ and $v(\cdot)$. To illustrate, a trivial two player type I model is one in which (i) all individuals in the economy have preferences of the form (5), where $v_0(\cdot)$ and $v_1(\alpha)$ are identical for all, but $v_2(y)$ varies; (ii) $v_2(y)$ is linear for one half the population and equal to some concave function for everyone else; and (iii) choosing the role of P implies an observable and fixed effort level $\bar{\alpha}$. Then (provided $\bar{\alpha}$ is not so large that autarky dominates production) the risk averse players will all assume the role of P, the rest will become A's, and the compensation rule will be $y(q) = q - \bar{y}$ (i.e., full insurance for P).²⁶

Of course, this very un-Knightian outcome is quite special; however, a less contrived model along the lines of the Khilstrom-Laffont (1979) general equilibrium model of implicit contracting would likely yield a similar result, as well as propositions on the size distribution of firms and entrepreneurial earnings.

The second direction through which external markets can influence the agency relationship is via the actions of "other agents". It is easy to see that the isolation in which P and A interact may be of some importance. Recall that the central problem in the agency framework is that even though A is generally risk averse, he must always face some income risk in order to provide incentives to expend effort. If individuals external to the firm are willing to bear some of that risk for a price, incentives are importantly diminished. The correct way to approach such possibilities requires an equilibrium treatment of them. Indeed allowing for trade with individuals outside the firm in a general equilibrium model of the type described above (wherein choosing the role of external agent is endogenous) will again likely lead to the conclusion that it is the least risk averse individuals who will choose to become agents, for they will be least disposed to utilize trade which removes incentives.

Taking these considerations together, external markets provide one reason why piece rates, risky as they are, are nonetheless the dominant form of output-dependent compensation--self-selection between P and A roles within the firm, plus the reduced propensity to try to eliminate incentives through trade with individuals outside the agency framework, combine to allocate the least risk averse individuals to the role of A in those situations where agency problems are important.

2. Richer Agency Settings

A second reason why the theory of agency has generated few predictions is, as emphasized above, that it contains little on which predictions might be based. The situation described is so sparse that there is little (only $y(q)$) about which to make predictions.

In contrast to the previous subsection, in which the basic setting was left undisturbed, this subsection devotes most attention to several directions in which the agency setting might be pressed to provide greater positive content. The discussion proceeds in order of apparent difficulty.

One point which arises immediately is that the complete "comparative statics" of the agency setting are not completely worked out. A full parameterization of tastes, technology and distribution functions of random effects, and analysis of the effects on $y(q)$ of changes in such parameters is not yet available. Providing one would be a useful if fairly tedious exercise.

At the level of generality of the models presented above, clear results are hard to obtain. It would seem that what is required at this point is a set of special cases which provide insight into the type of restrictions needed to obtain results in the more general settings.

Equally traditional is the point that much could be learned by embedding agency ideas in a more conventional treatment of the firm. To illustrate, suppose aggregate output (Q , from (9), with realization q) of the n agents A_i is an intermediate input used to produce final output (Z) in conjunction with capital (K):

$$Z = \xi(K, Q),$$

where $z = \xi(K, q)$ and $\xi(\cdot)$ has the usual properties. P 's problem then becomes (letting the prices of output and K be unity and R respectively, and assuming identical A_i)

$$\max_{n, K, y_i(q_i)} \int \dots \int u[\xi(K, q) - \sum_{i=1}^n y_i(q_i) - RK] h(q_1, \dots, q_n | \hat{\alpha}) dq_1, \dots, dq_n$$

$$\text{S.T. } \hat{\alpha} = \operatorname{argmax}_{\alpha} \int \dots \int v[y_i(q_i), \alpha_i] h(q_1, \dots, q_n | \tilde{\alpha}_i) dq_1, \dots, dq_n$$

$$\text{and } \int \dots \int v[y_i(q_i), \hat{\alpha}_i] h(q_1, \dots, q_n | \hat{\alpha}_i) dq_1, \dots, dq_n \geq \bar{V},$$

where $h(\cdot | \cdot)$ is the joint density of the q_i given α and $\tilde{\alpha}_i \equiv (\hat{\alpha}_1, \dots, \hat{\alpha}_{i-1}, \alpha, \hat{\alpha}_{i+1}, \dots, \hat{\alpha}_n)$. Note that even if $u' \equiv 1$ this problem does not become equivalent to the standard setting unless $\xi(\cdot)$ is additively separable. Further, written this way the similarity to "random factors of production" model [see Hey (1979), Ch. 21, for example] is striking. Therein the distribution of the intermediate input is exogenous; the agency setting renders this distribution endogenous (since it depends on $y_i(q_i)$) as would pure monitoring or optimal assignment (i.e., other possible responses to informational asymmetry in the firm).

This setup contains many more questions than does the standard agency model. The issues are very standard--what are the effects of changes in the parameters $[R, \bar{V}, \text{parameters of } h(\cdot), u(\cdot), v(\cdot) \text{ and } \xi(\cdot)]$ on the endogenous entities $(n, K, y_i(q_i), \text{ and the distribution of } Z)$? In particular, restricting the choice of $y_i(q_i)$ as per the tournaments literature, is it the case that the substitution of n for K induced by a rise in R (for example) yields a greater reliance on tournaments as predicted when n is exogenous? How does the existence of the agency problem (non-degenerate Ω_i) influence the expected level and dispersion of the firm's output? The list of questions is now quite a long one.

An interesting extension of this firm problem involves restricting the Ω_i to be independent, and allowing the service flow available from K to be random and unknown at the time the α_i are chosen. Since K provides the common

link between the A_i , is it the case that increases in the dispersion of the service flow generates greater reliance on tournament-type compensation rules, as the existing literature would indicate?

The multiple-agent agency setting also provides a convenient framework for the analysis of several issues involving the interaction of agents.

First, consider a simple tournament model with just two A_i . It has been suggested (e.g., Lazear-Rosen) that there is scope for destructive activity. That is, for obvious reasons A_i may wish to destroy part of A_j 's output. This intuition turns out to be correct. To illustrate, suppose each A_i can (in addition to choosing α) surreptitiously destroy k units of the other's output at a dollar cost $C(k)$ ($C' > 0$, $C'' > 0$). It can be shown that even under rational expectations concerning the total output, a symmetric Nash equilibrium will usually involve $k > 0$ for both A_i . This occurs simply because taking the level of the other agent's destructive activities as given, even though output (and hence the total of the prizes) falls with his destructive activity, each A_i can raise the probability that he will win by setting $k > 0$. Destructive activity is another externality making for inefficiency in Nash equilibrium.

An interesting question is just what should P do about this malfeasant behavior? Intuition suggests a compensation rule wherein all A_i receive a low payment if $\sum q_i < \bar{q}$, with $\bar{q}$ chosen such that $\sum q_i < \bar{q}$ is appropriately unlikely when $k_i = 0$ for all i . Note that this is an interdependent compensation rule even though the Ω_i are independent, and (9) holds. It has the form of a penalty-type rule because malfeasant behavior generally alters the support of the density of Q .

The A_i can also interact by colluding, and in a tournament setting incentives to do so appear strong. If it is assumed that the A_i can coordinate effort levels, how do the optimal compensation rules change? In the standard tournament setting, allowing collusion diminishes the gain P obtains by construction of interdependent compensation rules and so generates a greater tendency towards individualistic rules. P may again wish to institute individual penalties based on aggregate output.

A caveat is required concerning the analysis of collusion in the agency framework. The basic agency problem arises because effort is costly to observe. Analysis of collusive arrangements supposes the agents to have an advantage in observing each other and to not be amenable to P 's offers to buy this information. Making such assumptions palatable appears difficult at best.

Another topic to which agency ideas can usefully be extended is the theory of hierarchy. Agency-type notions are prominent in early work on hierarchies (e.g., Tuck (1954), Simon (1965), Williamson (1967)) but are never incorporated in a manner sufficiently precise to generate testable predictions.

The leading example of modern work on hierarchies is Rosen (1982). Therein, pyramidal structures emerge as a result of the public good aspect of management, treated as a factor of production. Management interacts with workers to produce output used as an input in conjunction with higher management time one hierarchial step up, and so on.

What is quite clear from the existing analysis is that the superior-subordinate relationship, which is in large part what is meant by a hierarchy, is not present. As Rosen points out, the operation of the market yields efficient assignment of individuals to spots in a pyramidal pattern of trade, but there is no reason why these individuals need even be in the same firm; the intermediate goods could equally be traded across firms. "Control" is just not an intrinsic part of the setup.

As shown above, utilization of an institutional framework wherein one player determines the compensation rule for the rest (which is just a very "flat" hierarchy) can be quite useful in a setting involving informational asymmetries. The extension of this notion to multiple layers is difficult due to the fact that the basic nature of the problem rules out restrictions such as (9). Nevertheless, the theory of hierarchy and general understanding of the behavior of firms would be vastly enhanced by such extension.

Another extension involves taking account of the observation that the form of the agency problem P faces is not exogenous. Some work has been done on this through analysis of the demand for monitors (e.g. Holmstrom (1982), Gjesdal), but the issue of joint choice of costly monitoring technology and compensation rule has gone unaddressed.

Equally promising is the idea that the agency problem can be altered by varying the activities which individuals are asked to perform. The basic theme is that production requires completion of a collection of tasks, and that the allocation of workers to tasks is not predetermined. P can render the agency problem less severe by altering the assignment of workers to tasks. For example, there may be some tasks for which it is very straightforward to determine A_i 's level of effort if he is the only worker assigned to that task; such assignment could eliminate the agency problem entirely, but at the cost of using assignments which do not exploit workers' comparative advantage well. Analysis of this tradeoff would predict the nature of the situations in which it is necessary to use compensation rules apart from straight time rates. Alternatively, it would be possible to explain why

time rates are as common as they appear to be. Moreover, in such a framework predictions could be made concerning the level of discretion given workers in their job. That is, P may wish to assign A_i a set of tasks and let A_i choose the time allocation across them if A_i has better information (in the model II sense) about the random shocks currently relevant to those tasks. Smith (1937, Book I, Ch. 8) noted that jobs vary greatly in terms of their motivational and discretionary characteristics, and that compensation depended on those aspects of work. But nothing more is presently known about the forces generating this type of heterogeneity.

Another interesting extension is "quasi dynamic".²⁷ Consider a tournament in which the final ranking depends on cumulative output over some prespecified period (T), and the flow rate of output can be adjusted over time (t), $Q_{it} = f(\alpha_{it}, \Omega_{it})$. P observes only $q_i = \sum_t f(\alpha_{it}, \omega_{it})$ and α_{it} is chosen given $q_{i\tau}$ ($\tau=1, \dots, t-1$). If q_{it} can be negative (imagine negative profits for a product line for which A_i is responsible) and $f_{12} > 0$, the A_i who have been lucky up to date t may play a very conservative (reduced α) strategy over $t+1, \dots, T$, simply to avoid the possibility of allowing random influences to lower their standing. This seeming risk aversion on the part of front runners may be the source of the frequently heard claim that North American managers are "afraid to take chances". Similarly, those A_i who have been less fortunate face a vanishing probability of success and so have a tendency to coast. Both phenomena are familiar to sports fans.

Ideas of this type provide a fairly compelling explanation for the popularity of piece rates. Whenever the A_i monitor output and change effort

more frequently than it pays for P to measure output, piece rates, while risky, succeed in retaining incentives for both those in the lead and those left in the dust. Piece rates and periodic lumpy bonuses are a likely outcome.

A related notion involves sequential tournaments. Under declining absolute risk aversion, winners of past tournaments will be less risk averse than losers, and thus will typically face a more risky and incentive oriented compensation rule. Anecdotal evidence on the intra-firm structure of compensation squares with this proposition.

3. Modelling Issues

This subsection deals with future work which is less oriented towards enrichment of the theory and more in the direction of improved modelling methods and clarification of existing theoretical work.

As mentioned above, some work has been done on placing the agency framework in a proper game theoretic framework. In particular, multiple agent problems are posed as games of incomplete information [Harsanyi (1967)]. A brief overview is in order.

Harris and Townsend (1981) is the central paper in the game theoretic approach. An early paper by Groves (1973) has also been influential.

Both address situations vastly more general (with an exception to be noted shortly) than that presented above. For the purposes at hand it is appropriate to specialize them to the agency framework, and to omit detailed development of the game theoretic approach. The omission of detail and stylization of the results is thus more substantial here than elsewhere in

this essay. The interested reader will certainly wish to consult the original works. Further, the Groves analysis focuses on a situation where transferable utility is assumed, which in the present problem has much the same effect as risk neutrality. Attention will be therefore confined to Harris-Townsend.

The analysis begins with an important restriction common to many games of incomplete information--altogether, the players in the game have complete information at the time decisions are taken; there is no "social risk".²⁸ In the agency context, this means the model is of type II; A_i observes Ω_i prior to choosing α_i , and there are no other random disturbances.

The second assumption is that players may communicate with P, and thus there is the possibility that compensation may be based on the information they supply.

Choices (y_i, α_i) made conditional on Ω_i , for the moment ignoring how these might be achieved, are called parameter contingent. The first result is that the only allocations which can be achieved in a game are parameter contingent. The basic idea is that if the choices did not depend on the Ω_i , which after all influence the productive capability of the economy, they could not, in general, be consistent with one another. For example, $y_i = 1$ for all i with P also receiving 1 would exhaust the product for some Ω_i and not others. Thus if an allocation is achievable in a game, it must be parameter contingent.

The second result is that only those parameter contingent allocations for which the A_i rationally reveal Ω_i are achievable in a game. Though nontrivial to demonstrate formally, the basic logic is that if the A_i do not reveal Ω_i , but rather send some other message, some player will not receive what is expected given the messages. Put differently, conditional on the messages all players know what each other will do, what output is, and what they are to receive. The

equilibrium concept is essentially just that players get what they expect (recall that given Ω_i , there is no other uncertainty). If some agent does not truthfully report Ω_i , someone will not receive what he anticipates.

The main import of these results is not in terms of predictions but rather what they imply for modelling strategy. In agency problems of the Model II type, all that need be focused on are situations where the A_i rationally reveal Ω_i , and therein y_i can be conditioned directly on Ω_i . Nothing else is achievable without systematic error on the part of P and A_i .

An important extension of this work is to situations involving social risk. The difficulty which arises is that holiest of grails--finding a good equilibrium concept. When there is social risk, in the sense that at the time y_i and the α_i are chosen Ω_i is not known, a natural candidate for solution concept in a game with communication is the rational expectations equilibrium--when making their decisions, P and the A_i make use of the correct conditional distribution of the unknown entities. While this approach works, the analysis has to be undertaken with extreme caution for the following reason. Rational expectations is an approximation to a dynamic process wherein individuals do not make systematic errors, and is reasonable and appealing in the context of decision theory. Therein the distribution of the unobservables depends only on exogenous entities which individuals learn over time, and use of rational expectations merely collapses this learning. As discussed above, the agency setting is in general poorly approximated by a decision theoretic framework. While it is not inconsistent to assume that players know the correct distribution of outcomes, it is a vastly stronger restriction than in

the decision theoretic setting because the distribution is generated by the strategic choices of other players, which would not remain constant if the players were required to learn it over time.

Another way to see this point is to note that under rational expectations, continual truth telling is likely to emerge as optimal behavior in the presence of social risk, simply because not doing so would generate a distribution distinct from that consistent with the messages sent.²⁹

However, in contrast to the case of no social risk, there is simply no way for the players to check this ex post; one observation does not generally discriminate among distributions.

At present a good solution to this problem is not available. But there is little doubt that the formalism of game theory is required if much further progress is to be made in obtaining very general results in the theory of agency. Thus it is important that the search for a reasonable solution concept in the presence of social risk go ahead.

A second theoretical issue is related to necessary and sufficient conditions for $0 < y' < 1$. At present those conditions are joint restrictions on technology, stochastic structure and tastes $[f(\cdot), G(\cdot) \text{ and } v(\cdot)]$. Since restrictions of this kind are central to achieving useful results in the agency model, determination of independent restrictions on $v(\cdot)$, $f(\cdot)$ and $G(\cdot)$ which yield the required restrictions would be most helpful. At present the convolution of tastes, technology and stochastic structure renders it very difficult to be clear on exactly what is being assumed when the required conditions are imposed. This is hardly a desirable state of affairs given the importance of the restrictions for the results.

Finally, the restriction (5) is a strong one. While what (5) means is clear, it is not at all obvious which results actually require (5). Some effort in this direction would also be useful.

V. Conclusions

This essay set out to provide an accessible overview of the economic theory of agency, and suggest directions in which the theory might be pushed.

Its primary point can be stated as follows. The theory of agency is not well enough developed to be of much scientific interest; there is a framework but little in the way of operationally meaningful hypotheses. Yet the ideas are promising ones, and there appear to be numerous avenues of enrichment which will result in refutable propositions.

Finally, the reader will note that the "New Directions" suggested above do not constitute what might reasonably be called a research agenda. This term refers to an orderly step-by-step progression of results. The theory is not ready for such a restricted plan. Indeed what is required is a multi-pronged attack on a sizable collection of agency issues. Having made some headway in some of these directions, it will be time to argue for an (some?) agenda.

FOOTNOTES

¹It is important to recognize at the outset that the agency problem will differ significantly from a problem in decision theory. Therein, individuals act with incomplete information about the state of nature, rather than about another individual's actions. The basic difference arises because nature will not exploit the informational asymmetry.

²It is worth noting that the efficiency issues, central from the outset, are still not clearly resolved. The reason is simply that in an environment involving informational asymmetries and random occurrences, what efficiency means is far from obvious. See Holmstrom (1977), Wilson (1978), Myerson (1979), and Holmstrom and Myerson (1983).

³Stiglitz (1975) analyzed a model with more than one agent. But since only aggregate output was observable, the problem effectively operated as if there were one agent.

⁴In what follows, attention is directed almost solely to models of the type I variety, for this is the setting in which most results are available.

⁵Focusing on Q rather than α is a useful analytical device introduced by Mirrlees (1975) in a different, but analytically similar, context.

⁶Assuming the support of $h(\cdot)$ independent of α is actually not quite enough to rule out penalty contracts. In fact, except for the case where Ω is a discrete random variable, analyzed by Grossman and Hart (1983), the precise conditions required to eliminate this possibility are not known.

Sufficient conditions for $y(q)$ to be of the penalty form are available.

Holmstrom (1982) gives a very clear account of them. In what follows it will be assumed that $y(q)$ does not involve a penalty.

⁷For simplicity, uniqueness of $\hat{\alpha}$ is assumed.

⁸It is not hard to see why preferences of the form (5) separate attitudes toward risk from α . For given α , the utility of income is of the form $a + bv(y)$. Variations in α involve changing a and b , which is simply an affine transformation of the original preferences. As is well known, affine transformations do not alter attitudes toward risk.

⁹These conditions were first argued to be sufficient by Mirrlees, but the proof was incorrect. Grossman and Hart proved the result assuming P is risk neutral.

¹⁰Grossman and Hart's "spanning condition" implies the above sufficient conditions. Though a stronger restriction, it is quite useful because it allows the construction of special cases to proceed in a straightforward way.

¹¹Since in this formulation P obtains all rents, any Pareto-improving change which P can implement will be undertaken.

¹²This point is sharpened by returning to (7). Notice that h_{α}/h is the derivative of the log likelihood of α when the data are q . Thus it appears as if the fact that $y(q)$ is larger when $h_{\alpha} > 0$ is an attempt by P to induce A to choose an α about which he can more easily make inferences. However, even when P is trying to induce A to work harder but (5) does not hold, it may be that P encourages A to choose α for which h_{α}/h is low as a means of generating spurious randomness. Under (5), decision theory is relevant, but only in a cosmetic sense.

¹³Previously, inclusion of Ω was necessary to prevent P from learning α through observation of q . Here, observation of q does not identify any particular α_i , so Ω can be dropped temporarily. Notice that without Ω , the support (albeit a degenerate one) of the density of Q depends on the α_i . Consequently, a penalty rule is to be expected.

¹⁴Note that this contract is of the "payment if and only if the goods are received" form.

¹⁵This is not the only $y_i(q)$ which will work.

¹⁶Mirrlees (1976) analyzes a situation wherein the Q_i are linked solely by the fact that P must spend time observing the output of each.

¹⁷The j^{th} order statistic of a sample is the j^{th} largest element in the sample.

¹⁸Note that $y_i^T(q_i) > y_i^T(q_i')$ if and only if $q_i > q_i'$ (ignoring ties). It is not known whether this monotonicity is optimal for $n > 2$.

¹⁹Interestingly, the form of y_i^T is closely related to the ad hoc beat-the-average [see Shubik (1983)] behavior once imposed in oligopoly problems. In those models, such schemes are ad hoc because there is no meta-agent to play the motivational role of P.

²⁰This result rests heavily on (i) the additive technology. In Nalebuff-Stiglitz, $Q_i = k(\alpha_i, \bar{\Omega}) + \Omega_i$ and an analogous result does not hold in general; and (ii) the absence of "gaps" (discussed below).

²¹Nalebuff and Stiglitz obtain a very different result when n increases. They restrict y_i^T to the form

$$y_i^T(q_i) = \begin{cases} \bar{y}_1 & \text{if } q_i = r_1 \\ \bar{y}_2 & \text{otherwise} \end{cases}$$

and so do not obtain extra flexibility as n increases.

²²The decision theory flavor is actually a real part of the analysis in Radner's work. P must decide if $\alpha \neq \hat{\alpha}$, where $\hat{\alpha}$ is appropriately thought of as parametric.

²³In Radner's analysis P and A maximize lifetime average (per period) expected utility. In the last period, the best response by A involves choosing $\alpha = \alpha^*$. But choosing $\alpha = \hat{\alpha}$ instead yields lifetime average expected utility only ϵ smaller for a game with many periods. Thus the strategy of playing $\alpha = \hat{\alpha}$ every period is an ϵ -equilibrium.

²⁴This works because all workers are identical and hiring is costless. When this fails, P's threat to fire the worker is not credible; the analysis becomes much more difficult.

²⁵This topic is mentioned in Stiglitz (1975), but is not subjected to analysis. See also Cooper and Hayes (1983).

²⁶Ricketts (1984) provides an analysis of a two player model wherein the choice of which player assumes the role of the principal is examined.

²⁷This idea is due to Bengt Holmstrom.

²⁸Social risk is usually handled in games of incomplete information by adding nature as a player. However, in the present context, where communication is possible, such an addition gives the analysis a distinctly surrealistic flavor which is perhaps best avoided.

²⁹It may be the case that the A_i are understood to send false messages from which the correct data can be inferred. The argument is altered little.

REFERENCES

- Alchian, A. and Demsetz, H., "Production, Information Costs, and Economic Organization," American Economic Review 62 (1972), 777-795.
- Becker, G. S. and Stigler, G. J., "Law Enforcement, Malfeasance, and Compensation of Enforcers," Journal of Legal Studies 3 (1974), 1-18.
- Borch, K., "Equilibrium in a Reinsurance Market," Econometrica 30 (1962), 424-44.
- Carmichael, L. H., "The Agent-Agents Problem: Payment by Relative Output," Journal of Labor Economics 1 (1983), 60-65.
- Cheung, S. N. S., "Private Property Rights and Sharecropping," Journal of Political Economy 76 (1968), 1107-1122.
- Cooper, R. and Hayes, B., "Multiperiod Insurance Contracts," mimeo (1983).
- Dubey, P., "Inefficiency of Nash Equilibria," mimeo (1983).
- Fama, E. F., "Agency Problems and the Theory of the Firm," Journal of Political Economy 8 (1980), 288-307.
- Gjesdal, F., "Information and Incentives: The Agency Information Problem," Review of Economic Studies 64 (1982), 373-390.
- Green, J. R. and Stokey, N. L., "A Comparison of Tournaments and Contracts," Journal of Political Economy 91 (1983), 349-364.
- Grossman, S. J. and Hart, O. D., "An Analysis of the Principal-Agent Problem," Econometrica 51 (1983), 7-45.
- Groves, T. "Incentives in Teams," Econometrica 41 (1973), 617-631.
- Harris, M. and Townsend, R. M., "Resource Allocation Under Asymmetric Information," Econometrica 49 (1981), 33-64.

- Harris, M. and Raviv, A., "Optimal Incentive Contracts with Imperfect Information," Journal of Economic Theory 20 (1979), 231-259.
- Harris, M. and Raviv, A., "Some Results on Incentive Contracts with Applications to Education and Employment, Health Insurance, and Law Enforcement," American Economic Review 68 (1978), 20-30.
- Harsanyi, J. C., "Games of Incomplete Information Played by 'Bayesian' Players," Management Studies 14 (1967), 159-182.
- Hey, J. D., Uncertainty in Microeconomics [New York: New York University Press], 1979.
- Holmstrom, B., "On Incentives and Control in Organizations," Ph.D. Dissertation, Stanford University, 1977.
- Holmstrom, B., "On Incentives and Control in Organizations," Ph.D. Dissertation, 10 (1979), 74-91.
- Holmstrom, B., "Moral Hazard in Teams," Bell Journal of Economics 13 (1981), 324-340.
- Holmstrom, B., "Managerial Incentive Problems--A Dynamic Perspective," mimeo (1983).
- Holmstrom, B. and Myerson, R. B., "Efficient and Durable Decision Rules with Incomplete Information," Econometrica 51 (1983), 1799-1819.
- Hurwicz, L., "On Informationally Decentralized Systems," in Decision and Organization, C. B. McGuire and P. Radner (eds.) [Amsterdam: North-Holland], 1972.
- Johnson, D. G., "Resource Allocation Under Share Contracts," Journal of Political Economy 58 (1950), 111-123.

- Kamien, M. I. and Schwartz, N. L., Dynamic Optimization: The Calculus of Variations and Optimal Control in Economics and Management [New York: North-Holland], 1981.
- Khilstrom, R. E. and Laffont, J. J., "A General Equilibrium Entrepreneurial Theory of Firm Foundation," Journal of Political Economy 87 (1979), 719-747.
- King, H., "Multiperiod Labour Contracts in a Moral Hazard Setting," mimeo (1984).
- Lal, R., Staelin, R., "A Theory of Optimal Salesforce Compensation Plans," mimeo (1981).
- Lambert, R. A., "Long Term Contracts and Moral Hazard," Bell Journal of Economics 14 (1983), 441-452.
- Lazear, E. P., "Why is there Mandatory Retirement?" Journal of Political Economy 87 (1979), 1261-1284.
- Lazear, E. P., and Rosen, S., "Rank-Order Tournaments as Optimum Labor Contracts," Journal of Political Economy 89 (1981), 841-864.
- Malneg, D., "Pareto Optimality of Repeated Insurance Contracts When There are Problems of Moral Hazard and Adverse Selection," mimeo (1981).
- Marshall, A., Principles of Economics [New York: Macmillan], Eighth Edition, 1920.
- Mirrlees, J., "Notes on Welfare Economics, Information and Uncertainty," in Balch et al (eds.), Essays in Economic Behavior Under Uncertainty [Amsterdam: North-Holland], 1974.
- Mirrlees, J., "The Theory of Moral Hazard and Unobservable Behavior - Part I" [Nuffield College, Oxford], 1975.

- Mirrlees, J., "The Optimal Structure of Incentives and Authority Within An Organization," Bell Journal of Economics 7 (1976), 105-131.
- Myerson, R. B., "Incentive Compatibility and the Bargaining Problem," Econometrica 47 (1979), 61-74.
- Nalebuff, B. and Stiglitz, J. E., "Prizes and Incentives: Towards a General Theory of Compensation and Competition," Bell Journal of Economics 14 (1983), 21-43.
- O'Keefe, M., Viscusi, W. K. and Zeckhauser, R. J., "Economic Contests: Comparative Reward Schemes," Journal of Labor Economics 2 (1984), 27-56.
- Plant, A., "Trends in Business Administration," Economica 12 (1932), 45-62.
- Radner, R., "Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship," Econometrica 49 (1981), 1127-1148.
- Ricketts, M., "Principal and Agent, Risk Sharing and Incentives: A Simple Diagrammatic Exposition," mimeo (1984).
- Rogerson, W. D., "The Role of Reputation in a Repeated Agency Problem Involving Information Transmission," Institute for Mathematical Studies in the Social Sciences, Technical Report #377 (1982).
- Rogerson, W. D., "The First-Order Approach to Principal-Agent Problems," mimeo (1983).
- Rosen, S., "Authority, Control and the Distribution of Earnings," Bell Journal of Economics 13 (1982), 311-323.
- Ross, S. A., "The Economic Theory of Agency: The Principal's Problem," American Economic Review 63 (1973), 134-139.

- Rubenstein, A. and Yaari, M. E., "Repeated Insurance Contracts and Moral Hazard," Journal of Economic Theory 30 (1983), 74-97.
- Shavell, S., "Risk Sharing and Incentives in the Principal and Agent Relationship," Bell Journal of Economics 10 (1979), 55-73.
- Shubik, M., Game Theory in the Social Sciences [Cambridge: M.I.T. Press], 1982.
- Simon, H., Administrative Behavior: A Study of Decision Processes in Administrative Organizations, 2nd ed. [New York: Modern Library], 1937.
- Smith, A., The Wealth of Nation [New York: Modern Library], 1937.
- Spence, A. M. and Zeckhauser, R. J., "Insurance, Information and Individual Action," American Economic Review 61 (1971), 380-391.
- Stiglitz, J. E., "Incentives and Risk Sharing in Sharecropping," REStud 41 (1974), 219-255.
- Stiglitz, J. E., "Incentives, Risk, and Information: Notes Towards a Theory of Hierarchy," Bell Journal 6 (1975), 552-579.
- Tuck, R. H., An Essay on the Economic Theory of Rank [Oxford: Basil Blackwell], 1954.
- Williamson, O. E., "Hierarchical Control and Optimum Firm Size," Journal of Political Economy 75 (1967), 123-138.
- Wilson, R., "Information, Efficiency and the Core of an Economy," Econometrica 46 (1978), 807-816.