THE INTERNAL ECONOMY OF LARGE FIRMS*

Roy Radner

On the scale of economic organisation, the internal economy of the firm would seem to be at the opposite end from the economics of international trade, which was Professor Harry Johnson's special interest. In defence of my choice of topic for this occasion, I might point out that the internal economics of multinational firms span the globe and cross many national boundaries. However, I shall not today give any special attention to multinational firms, but just to large ones.

Table 1 gives some idea of the size of the largest U.S. firms, as reported by *Fortune* magazine in its most recent report on 'The Fortune 500.' In terms of employees, the largest is General Motors, with approximately 700,000 employees. This is approximately the same as the number of persons employed in manufacturing in The Netherlands in the same year (1983). Most of the Fortune 500 are, of course, much smaller; in terms of employment, the mean in 28.1 thousand, and the median is 13.1 thousand. Nevertheless, it is clear from the table that the internal economies of these firms are very large indeed.

In this paper I shall review, and try to explain in relatively nontechnical terms, some of the recent developments in economic theory that bear on the internal economics of firms. This review will be partial and selective, reflecting my own research interests. It will focus on two primary themes, *decentralisation* and *incentives*.

Even smaller firms are forced to decentralise information to some extent, given the very limited capacities of decision-makers to observe, process, and

Table 1. *Some Large U.S. Firms, 1983*

Sales Rank	Co.	Sales $ Bn.	Assets $ Bn.	Net Income $ Bn.	Employment Thous.	Employees Rank
1	Exxon	88·6	63·0	4·98	156	(9)
2	G.M.	74·6	45·7	3·73	691	(1)
250	Monfort	1·32	0·199	0·016	3·5	(472)
308	Ex-Cello	0·955	0·690	0·050	13·1	(250)
500	Tandem	0·418	0·415	0·031	4·4	(454)
	Mean	3·37	2·71	0·138	28·1	

Source: 'The 500,' *Fortune*, April 30, 1984, pp. 274–322

[1] The views expressed here are those of the author, and not necessarily those of AT & T Bell Laboratories. This paper was the basis of the Harry G. Johnson Lecture, Annual Meeting of the Royal Economic Society/A.U.T.E., Oxford, March, 1985.

communicate information. This is even true of decision-makers aided by large computer systems. The private information of a decision-maker gives him some measure of private power to pursue his own goals, which may sometimes be in conflict with corporate goals. The private power conferred by informational decentralisation, and the conflicts of interest among the members of the firm, will in general give rise to inefficiencies, whatever the structure of formal organisation, compared to what would be ideally attainable with the given distribution of information among the decision-makers.

After a brief sketch of the history of alternative views of the nature of the firm (Section I), I give an informal description of the features of large firms that distinguish them from a market economy, and also from the simple model of the firm that one finds in undergraduate textbooks on price theory (Section II.).

Large firms are often organised into a collection of profit centres, which in some ways are 'firms within firms.' A first approach to a model of the relations among profit centres is provided by the theory of the core (Section III), which makes a potentially important contribution to the problem of transfer pricing. It can be shown that stable transfer prices exist – i.e. the core is nonempty – for a variety of 'nonneoclassical' technologies (although further exploration of this issue is needed). On the other hand, the core does not seem adequate as a model of truly centralised decision-making. This leads to the consideration of bargaining theory (Section IV) and Groves mechanisms (Section V).

Looking more closely at the authority relationships within a firm, we can distinguish hierarchies, matrix structures, and partnerships; all of these structures are found in some degree in all large firms. I first examine what principal-agent theory has to say about hierarchies, and then examine a model of partnership. (There does not yet appear to be any formal theory of matrix management.)

The following summary, in outline form, lists the main conclusions from the study of these models of decentralised decision-making:

1. Equilibria of decentralised organisations are typically inefficient, relative to what could theoretically be attained with the given distribution of information among the several decision-makers.
2. The inefficiency of equilibria is caused by various incentive problems that arise from conflicts of interest among the decision-makers and/or between the decision-makers and the organiser. These incentive problems can be classified under three headings:
   a. misrepresentation
   b. moral hazard
   c. free riding
3. Misrepresentation can be remedied by suitably designed compensation mechanisms, provided 'budget balancing is not required (Groves – Clarke – Vickrey mechanisms).
4. The inefficiency due to moral hazard alone may be remedied in long-term equilibria to the extent that the players are not too short-sighted, i.e.,

to the extent that they do not discount future utility very much. (Principal-agent games.)

5. Free riding and moral hazard together may be only partly remedied in long-term equilibria with far-sighted players, unless there are individual measures of performance that are publicly observable. (Partnership games.)
6. Efficiency can be improved in long-term equilibria with far-sighted players if they are not too 'greedy.' (Approximate equilibria.)
7. Long-term games typically have infinitely many equilibria.

The theory of the internal economy of the firm is in its infancy, and so this review will turn out to be as much an agenda of problems to be solved as an account of problems already solved. I hope that, at the end, some of you will agree that the agenda is an interesting and challenging one for economists.

II. EARLY VIEWS OF THE NATURE OF THE FIRM

Although an examination of the history of economic thought about the nature of 'the firm' is not one of my goals for this paper, a brief discussion may help put the main body of the paper in some perspective.

In such discussions it is traditional to cite the 1937 article by Coase, 'The Nature of the Firm.'[1] Coase's paper was perhaps more notable for the questions that it posed than for any answers it provided. Essentially, he asked why some economic activities are organised *within* firms rather than *among* them. Without agreeing, he paid great attention to the views of Frank Knight (1921), who stressed the importance of uncertainty and the role of the 'entrepreneur' as a specialist in decision-making. In particular, he quotes passages in which Knight argues that 'the fact of uncertainty brings about.... important characteristics of social organisation.'

> .... When uncertainty is present and the task of deciding what to do and how to do it takes the ascendancy over that of execution the internal organisation of the productive groups is no longer a matter of indifference or a mechanical detail.[2]
>
> The most fundamental change is:
>
> the system under which the confident and venturesome assume the risk or insure the doubtful and timid by guaranteeing to the latter a specified income in return for an assignment of the actual results.... With human nature as we know it it would be impracticable or very unusual for one man to guarantee to another a definite result of the latter's actions without being given power to direct his work. And on the other hand the second party would not place himself under the direction of the first without such a guarantee.... The result of this manifold specialisation of function is the enterprise and wage system of industry. Its existence in the world is the direct result of the fact of uncertainty.[3]

Although uncertainty obviously makes decision-making more complex, the picture of the firm as headed by a rational, profit-maximising entrepreneur persisted – and probably persists – as the mainstream economic model of the firm, although not without dissent. The mainstream model was fortified by the burst of activity in statistical decision theory and operations research that

[1] Full descriptions of all references are gathered at the end of the paper.
[2] Knight (1921), pp. 268–95; quoted in Coase (1937), p. 345, 346.
[3] Knight (1921), pp. 269–70; quoted in Coase (1937), p. 347.

occurred during and after the Second World War. This activity not only showed that very complex decision problems could be solved, but also led to fundamental reconsiderations of the foundations of the theory of probability and uncertainty (Savage, 1954).

One current of dissent was pioneered by Herbert Simon (1957, 1979), and followed by Cyert and March (1963) and others, who stressed the importance of 'bounded rationality' in the decision-making of individuals and groups.

Another current questioned the mainstream picture of the 'unitary' direction of the firm, especially arising out of the separation of management from ownership (Berle and Means, 1932; R. A. Gordon, 1961). This separation, i.e., the fact that the owners and the managers are different persons, gives rise to a loss of control of the managers by the owners. The discretion of the managers to do things other than maximise the profits of the firm is enhanced to the extent that the owners can neither perfectly observe the managers' decisions nor perfectly share the managers' information. It is also enhanced by any monopoly power of the firm, the actual degree of which is better known to the managers than to the owners (Williamson, 1963; Marris, 1964). This partial loss of control extends to the relationship between the management and the workers (or between different levels of management). It is possible for the owners to improve their monitoring of the information and actions of managers and workers, but only with the expenditure of resources, and so the *net* benefit to the owners of such increased monitoring may not be unambiguously positive (Alchian and Demsetz, 1972). Formal models of this loss of control were introduced by Ross (1973) (agency theory) and by Groves (1973) in a more general setting (theory of incentives in teams). The post-Second-World-War development of the theory of games provided the methodology needed for the formal analysis of the conflicts of interest within the firm, and it is the methodology of game theory that underlies the developments that I shall describe in this paper.

Some of you may wonder why I have not mentioned Adam Smith in this brief excursion into the history of economic thought about the nature of the firm. In fact I have saved Adam Smith for the last. After describing the difference between a joint stock company and a 'private copartnery,' he goes on to offer this opinion of joint stock companies:

> The trade of a joint stock company is always managed by a court of directors. This court, indeed, is frequently subject, in many respects, to the control of a general court of proprietors. But the greater part of those proprietors seldom pretend to understand anything of the business of the company, and when the spirit of faction happens not to prevail among them, give themselves no trouble about it, but receive contentedly such half-yearly or yearly dividend as the directors think proper to make to them. ... The directors of such companies, however, being the managers rather of other people's money than of their own, it cannot well be expected that they should watch over it with the same anxious vigilance with which the partners in a private copartnery frequently watch over their own. Like the stewards of a rich man, they are apt to consider attention to small matters as not for their master's honour, and very easily give themselves a dispensation from having it. Negligence and profusion, therefore, must always prevail, more or less, in the management of the affairs of such a company. It is upon this account that joint stock companies for foreign trade have seldom been able to maintain the competition against private adventurers.[1]

[1] A. Smith (1776), vol. 2, pp. 264, 265.

Smith certainly understood the incentive problems inherent in the principal-agent relationship. On the other hand, with hindsight it appears that his judgement of joint-stock companies was too pessimistic. What would he have thought of the Fortune 500?

III. THE ECONOMY OF THE LARGE FIRM

I turn now to a more systematic, but still informal, description of the features that distinguish the internal economy of a large firm from the simple model of the firm that one finds in undergraduate textbooks on price theory, and also distinguish it from the model of a collection of traders in a market economy.

First, as already pointed out in the previous section, there are many 'players' in the firm's internal economy: shareholders, directors, managers, and workers (and sometimes creditors). These different players typically have at last partially divergent interests; hence the difficulty of imputing to the firm a single objective. Among the objectives of the several players in the firm are profits, growth, monetary compensation for managers and workers, quality of work, perquisites, and status. In addition, in the pursuit of these objectives, the different players may have different attitudes towards risk.

Second, the constituent parts of a firm may have technologies that do not conform to the 'neoclassical' hypotheses. The internal economy of the firm will usually be full of instances of increasing returns to scale (up to a point), indivisibilities, and externalities. I am not denying that these technological features may also be important at the level of the entire firm, as well. Nor will I debate whether the neoclassical hypotheses about the technology of the firm are adequate for the analysis of many questions about the interactions of firms in markets. I only emphasise that a theory of the *internal* economy of the firm will not be successful if it holds on to neoclassical assumptions for the firm's constituent parts.

There are many features of the internal technology of the firm that can be nonneoclassical. For example, the provision of overhead services often exhibits increasing returns to scale (that is why such services are centralised!). Some overhead services may also have the nature of a 'local public good.' Although the technology of research and development is not well understood, it is also a source of nonconvexity. Most of the development will probably be done in the form of specific projects, each for a particular line of business. The cost of a development project is usually independent of the use to which it is put; e.g., the cost of designing an integrated circuit will be independent of the number of such circuits ultimately produced. This is a source of increasing returns to scale. Research projects may benefit more than one line of business, and fundamental development may be used in several products. Thus some research and development activities may also produce local public goods.

Third, a study of the internal economy of the firm should pay attention to the structure of information processing within the firm (observation, communication, computation), and also to the structure of authority, which limits the things that the players may and may not do.

Roughly speaking, the subunits of a firm are of two types, *profit centres* and *cost centres*. As the term suggests, a profit centre has imputed to it a profit in each accounting period, and the manager of the profit centre is given incentives to maximise his division's profits, to reach a given target profit, or to reach a given target rate of return. Both the budget and the activities of a cost centre are planned at the beginning of each budget period; the manager of a cost centre is judged by whether he can accomplish his planned activities without exceeding his budget. Typically, the manager of a profit centre has more flexibility in directing the affairs of his unit than does the manager of a cost centre. However, even the profit centre manager is usually constrained by headquarters in matters of investment, location, the hiring of long-term employees, and the selection and development of his product lines.

The rules for the transfers of goods and services from one part of the firm to the other, and for the pricing of such transfers, are not given *a priori*, but can be designed by the firm itself, as can the rules for compensating the managers of different units and at different levels. Indeed, the very division of the activities of the firm into separate units can be determined by the headquarters. Not all inputs and outputs are readily describable in written agreements and contracts; this (among other things) limits the degree to which the firm can be subdivided into profit centres. In particular, if the output of a subunit is difficult to measure quantitatively, it is likely to be treated as a cost centre. These considerations apply, of course, to the internal organisation of a profit centre, as well as to the firm as a whole.

In addition to the classification of subunits into profit centres and cost centres, one can distinguish three main forms of authority relationship within firms: hierarchy, matrix management, and partnership. We are all familiar with hierarchy. Matrix management refers to the superposition of two different hierarchies, so that a middle manager may have two 'bosses'. For example the head of a development project that is designing a new integrated circuit may report to two supervisors, (1) his superior in the research and development unit, and (2) the head of the corresponding department in the integrated circuit line of business. The term 'partnership' refers to a situation in which a group of managers (or workers) jointly produce a particular output, and are held jointly responsible for the activity. Partnerships are usually found at lower levels, e.g., small groups of workers, but some committees and task forces of managers may also serve as partnerships (perhaps temporary).

The two classifications I have just described, (1) cost vs. profit centre, and (2) type of authority relationship, are of course independent dimensions of organisation. A profit centre may make use of one or more authority relationships within it, and the hierarchial organisation of the firm's managers will usually include the managers of both profit and cost centres.

In what follows I shall primarily deal with theories of the relationships among profit centres, and with theories of hierarchies and partnerships. I shall not have the time of discuss the important question of how to divide up the activities of the firm into profit and cost centres, and in any case I am

not aware of much formal theorising about the subject. Nor am I aware of any formal theory of matrix management.

III. TRANSFER PAYMENTS AND THE CORE OF A PROFIT-CENTRE GAME

If a firm is organised as a collection of profit centres, and if some profit centres use goods, services, or resources supplied by other centres, then the issue arises of how the supplier centres will be compensated in accounting dollars. This is sometimes called the problem of 'transfer pricing,' although the term may be misleading since it suggests fixing a price per unit of the commodity that is provided.

If the profit centres – or *divisions*, as I shall call them here – are allowed to bargain freely among themselves regarding the terms of such transfers of commodities (goods, services, resources) and the corresponding transfers of accounting dollars, one might predict the outcome to be in the *core* of a suitably defined game, should the core be nonempty. Roughly speaking, an element of the core is a feasible production plan for the firm, together with an allocation of accounting profits among the divisions such that no subset of divisions obtains a total profit smaller than it could get by operating 'on its own. Thus, an element of the core is in some sense 'stable' against recontracting by groups of divisions.[1] (A formal definition will be given below.)

The definition of the core does not guarantee that it will not be empty. The core has been studied for models of entire economies; in particular, its nonemptiness has been established for models of pure exchange economies, and for economies with production in which every subset of consumers has available to it – in principal – the same set of production possibilities (Debreu and Scarf, 1963; Hildenbrand, 1974). The essence of the typical firm, however, is that different divisions have available *different* production possibilities, each division specialising in a particular set of activities. On the other hand, a firm is simpler than an economy as a whole in that there are no 'consumers' in the usual sense. Finally, as noted above, a realistic model of the internal economics of the firm should allow for so-called 'nonconvexities,' i.e., possibilities of increasing returns to scale, fixed costs, etc. Thus the theories of the core for entire economies, as developed thus far, are not appropriate as they stand for the internal economics of the firm.[2]

Example 1. One Supplier, Three Customers

Here is an example; although not particularly realistic it will illustrate the concept of the core.

[1] The concept of the core was introduced by Edgeworth (1881), in the context of pure exchange.

[2] For completeness, I should mention the following 'neoclassical' theorem: if the optimal production plan for the firm can be supported by shadow prices for all the goods and services that are transferred among the divisions, then the imputation that gives each division its shadow profit is in the core. For a precise statement and proof, see (Radner, 1985*b*). The proposition is implicit in the standard treatments of activity analysis. Of course, this proposition is only of limited interest in our context, since, roughly speaking, the existence of shadow prices (corresponding to an optimal plan) can only be guaranteed under neoclassical hypotheses about technology.

Suppose that there are three customers and one supplier. Each customer can obtain either 1 unit or 0 units from the supplier, and nothing from the outside. The supplier's cost is c_n if it supplies a total of n units ($n = 0, 1, 2,$ or 3). The (incremental) profit to a customer is b if he gets one unit of the input, and zero otherwise. If the supplier provides one unit to each customer, then the profit to the firm as a whole will be $3b - c_3$ ('profit' will here always mean 'incremental profit'). If the supplier provides one unit to each of two customers, and the third customer gets nothing, then the supplier and the two satisfied customers will make a total profit of $2b - c_2$, and the unsatisfied customer will make no profit. Similarly, if only one customer were satisfied, it and the supplier would make a total profit of $b - c_1$, and the unsatisfied customers would make nothing.

Assume that

$$0 = c_0 < c_1 < c_2 < c_3,$$
$$b - c_1 < 0 < 2b - c_2 < 3b - c_3;$$

then the firm's profit is maximised when all three customers are supplied, and furthermore, supplying only one customer would result in a loss. Hence the only coalitions of divisions that can make a positive profit on their own are those that include the supplier and at least 2 customers.

It is clear that all 3 customers should be supplied, but what transfer price (or prices) should they pay the supplier? To put it another way, how should the total maximum profit, $3b - c_3$, be divided among (or 'imputed to') the 4 divisions? Let $x = (x_0, x_1, x_2, x_3)$ denote such an imputation, where x_0 is what the supplier gets, and x_i is what customer i gets ($i = 1, 2, 3$). An imputation $\mathbf{x}$ is in the *core* if it gives each *coalition* (i.e., subgroup) of players at least as much as it could get on its own. Hence, in this example, an imputation $\mathbf{x}$ is in the core if:

$$\begin{aligned} x_0 + x_1 + x_2 + x_3 &= 3b - c_3 \\ x_0 + x_1 + x_2 \quad &\geqslant 2b - c_2 \\ x_0 + x_1 \quad + x_3 &\geqslant 2b - c_2 \\ x_0 \quad + x_2 + x_3 &\geqslant 2b - c_2 \\ x_i &\geqslant 0 \ (i = 0, 1, 2, 3). \end{aligned} \tag{1}$$

Note first that the imputation that gives all the profit ($3b - c_3$) to the supplier is in the core. This is so because (*i*) no coalition containing the supplier can get more than ($3b - c_3$), and (*ii*) no coalition not containing the supplier can get a positive profit. This imputation corresponds to a transfer price equal to b for each customer.

There are other imputations in the core. For example, suppose the transfer price is constrained by company policy to be the same, say t, for all customers; then the profit for the supplier would be $3t - c_3$, and for each customer would be $b - t$. It follows easily from (1) that the transfer price will determine an imputation in the core if and only if it is at least as large as both the incremental cost and the average cost for 3 units, and does not exceed the customers' incremental benefit.

Example 2. Internal Competition

Let us modify the preceding example by supposing that instead of one supplier there are two, with identical cost functions. (Otherwise, we maintain all the assumptions of Example 1). We shall see that in this case (*i*) no core imputation gives either supplier a positive profit, and (*ii*) there are no core imputations at all unless the average cost for 3 units is no greater than the average cost for 2 units.

First, the optimal plan for the firm is to have one of the suppliers produce 3 units, and the other produce nothing. In principle, both suppliers could receive transfer payments even if only one produced. However, any coalition consisting of 3 customers and one supplier can make a profit of $(3b - c_3)$; hence any such coalition (there are two) must receive a total imputed profit of at least $(3b - c_3)$. Suppose that an imputation gave one of the suppliers (say the first) a positive profit; then the second supplier could form a coalition with the three customers and make a greater profit on their own, essentially dividing the first supplier's profit among themselves. Hence a core imputation cannot give either supplier anything but a zero profit.

Suppose now that the average cost for 3 units is greater than that for 2 units. In a core imputation the three customers would have to get a total imputed profit of $(3b - c_3)$, or an average imputed profit of $b - (c_3/3)$. Hence two of them would have to get a total imputed profit no greater than $2b - (2c_3/3)$. However, those two customers and one of the suppliers could, on their own, make a profit of $2b - c_2$, which is larger than $2b - (2c_3/3)$. Hence a core imputation is not possible. On the other hand, it is easy to see that, if the average cost for 2 units is at least that for 3 units, then the following imputation is in the core: each customer gets $b - (c_3/3)$. If the inequality is strict, then there will be other imputations in the core, too.

Example 3. Complementarity

In Example 2 the outputs of the two suppliers were perfect substitutes as inputs for the customers. Suppose instead that their outputs are complementary, so that no customer can obtain any benefit unless he gets one unit from each supplier (in this case, the benefit is b, as before). For simplicity, suppose that the suppliers have identical cost functions, but that each cost function is eactly half of the one in Example 1.

By an argument similar to that used in Example 1 one sees that any imputation that gives each customer zero profit, and the two suppliers a total profit of $3b - c_3$, is in the core. There will typically be other imputations in the core.

The three preceding examples suggest the following general conclusions:

1. The core will be nonempty under a wide variety of 'nonneoclassical' technological conditions, provided there is no competition among internal suppliers who produce close substitutes.
2. The core will be empty if there are alternative suppliers who produce close

substitutes, and their average costs are declining and then increasing in the relevant range of output.

These conclusions can be verified for some models more general than the examples (Radner, 1985*b*), but further research is needed to determine the domain of their validity.

IV. BARGAINING

We must now face up to a fundamental problem with the core as a model of decentralised decision-making. One might think that the division managers would be able to reach a point in the core through a process of bargaining, in the spirit of Edgeworth's process of 'recontracting' (Edgeworth, 1881). However, if the managers had private information about their respective costs and/or benefits, i.e., information that could not be observed directly by the other managers, would they not have an incentive to misrepresent those costs and benefits in order to influence the outcome of the bargaining process? Of course, in order to know how to manipulate the process in his favour, a manager would have to know the rules of the bargaining process, and be able to predict the behaviour of the other managers.

The theory of noncooperative games with incomplete information has been used to analyse this situation. I shall illustrate this approach with a simple example with one supplier and one customer, in the spirit of Example 1.

Example 4

Following the approach pioneered by Harsanyi (1967–68), the customer's benefit and the supplier's cost are assumed to be random variables, say B and C, respectively, with a joint probability distribution that is known to both managers. Before the bargaining begins, the customer observes B but not C, and the supplier observes C but not B.

I shall consider a particular bargaining process studied by Chatterjee and Samuleson (1983), in which the managers submit sealed bids. If the customer's bid is at least as high as the supplier's bid, then one unit is transferred at a price equal to the average of the two bids; otherwise there is no transfer. The customer's *strategy* is a rule that determines his bid as function of his observed benefit; similarly, the supplier's strategy is a function that determines his bid for each observed cost. A pair of strategies is a *Cournot–Nash equilibrium* (or just equilibrium) if neither manager can increase his *ex ante* expected profit by unilaterally changing his strategy.

In the special case in which B and C are independent random variables, and are uniformly distributed between 0 and 1, one can explicitly calculated an equilibrium pair of strategies:

$$\beta(B) = \begin{cases} B & , \quad 0 \leqslant B \leqslant \frac{1}{4}, \\ \left(\frac{2}{3}\right)B + \frac{1}{12} & , \quad \frac{1}{4} \leqslant B \leqslant 1. \end{cases}$$

$$\gamma(C) = \begin{cases} \left(\frac{2}{3}\right)C + \frac{1}{4} & , \quad 0 \leqslant C \leqslant \frac{3}{4}, \\ C & , \quad \frac{3}{4} \leqslant C \leqslant 1. \end{cases}$$

With the equilibrium strategies, the customer systematically offers a price that is below his benefit, and the supplier offers a price that exceeds his cost (in the relevant range of benefit and cost); see Fig. 1. One can show that, for these equilibrium strategies, the customer's bid will exceed the supplier's bid only if *the true benefit exceeds the true cost by one-fourth.* Hence, with positive probability, there will be cases in which there would be positive gains from a transfer but no transfer will take place. Because of this, the expected total profit of the two divisions in equilibrium is 9/64(= 0.140625) whereas the maximum potential expected profit is 1/6. In this sense, the equilibrium of the bargaining process is inefficient.

Example 4 deals with only one particular bargaining mechanism. One can imagine many others. For example, a two-stage mechanism might require that the customer make the first offer, and then that the supplier either accept that offer or make his own offer, which the customer then must finally accept or reject. Myerson and Satterthwaite (1983) have shown how to design, for any given prior probability distribution of B and C, a bargaining mechanism that maximises the total expected profit of the two bargainers. In particular, they show that, for the uniform distribution of Example 4, the Chatterjee–Samuelson mechanism is optimal.

In general, one expects that the optimal bargaining mechanism will depend on the prior distribution of costs and benefits. This is unfortunate, for at least two reasons. First, the same mechanism will be used for many transfers. Second, it is difficult to discover, with any reliability, the prior distributions that express the manager's prior beliefs about benefits and costs. This second point leads to a criticism of the model itself, namely, how can we verify whether or not the different managers have the same prior beliefs?

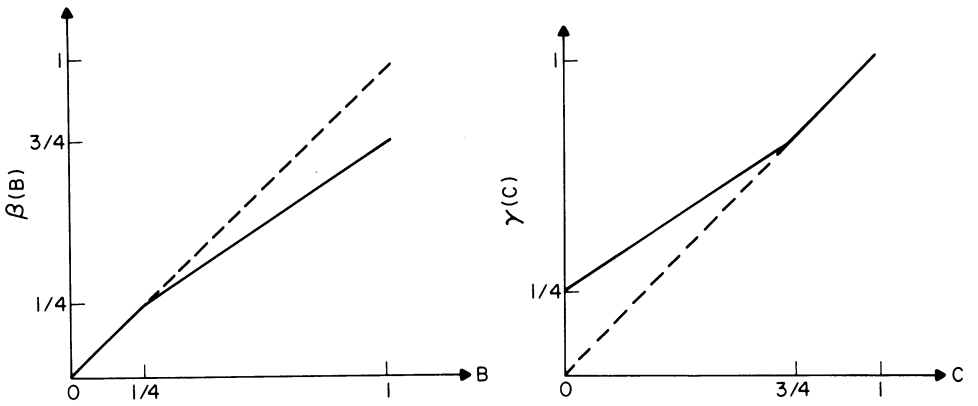

Fig. 1. Equilibrium strategies for C—S Mechanism

Finally, I am not aware of much progress on the theory of bargaining with more than two (but a small number of) participants, which would seem to be needed in order to modify the theory of the core to take account of the decentralisation of information.

V GROVES MECHANISMS PROVIDE ANOTHER METHOD OF TRANSFER PRICING

In the bargaining situation discussed in the preceding section, the transfer payment from the customer to the supplier was a literal 'transfer', in the sense that the amount that the customer paid out was equal to the amount that the supplier received. It was noted (but not proved there) that there exists no fully efficient bargaining mechanism with that property and with the additional property that neither division can be forced to make a loss on transaction. In this section I shall describe a mechanism that is fully efficient, and does not force either the customer or the supplier to make a loss, but does not satisfy the literal transfer requirement. This mechanism is a special case of a class of mechanisms proposed by Groves (1973); similar mechanisms were independently proposed by Vickrey (1961) and Clarke (1971).

To operate what I shall call the *Groves Mechanism* we shall have to bring the headquarters into action. The situation is the same as in the bargaining example except that:

1. the customer sends to headquarters a message, say b, purporting to be his true benefit (which is actually B);
2. the supplier sends to headquarters a message, say c, purporting to be his true cost (which is actually C);
3. if c exceeds b there is no transfer; otherwise the supplier transfers one unit of the commodity to the customer, the headquarters pays the supplier b, and the headquarters charges the customer c.

First note that the customer can guarantee that he will not make a loss, by reporting a benefit that does not exceed his true benefit; if he does this, then in case of a transfer the price he pays will not exceed the benefit he reports, which in turn will not exceed his true benefit. A symmetric argument is applicable to the supplier.

Second, it is optimal for each manager to report truthfully, *no matter now the other manager reports.* To see this, suppose that the supplier reports c. The customer's report will not influence what he pays in case of a transfer, but will only influence whether or not the transfer takes place. The customer wants the transfer to take place only if c does not exceed his true benefit, B. On the other hand, the transfer will take place if c does not exceed the reported benefit, b. The customer can therefore guarantee the decision he wants by reporting his benefit truthfully, i.e., by setting b equal to B. A symmetric argument is applicable to the supplier.

A strategy that is optimal for one player, whatever the strategy of the other

player, is called *dominant*. Thus I have shown that truthful reporting is a dominant strategy for each manager. *A fortiori*, truthtelling is an equilibrium, but in this case a particularly compelling one.

Finally, note that the headquarters makes an accounting loss in any transfer for which the reported benefit strictly exceeds the reported cost, since the headquarters loses the difference between the two numbers. However, in any transfer the total accounting gain of the headquarters and the two divisions is exactly the difference between the true benefit and the true cost. In the dominant strategy equilibrium this difference will be nonnegative if and only if there is a transfer. Hence the decision about the transfer is fully efficient.

The Groves mechanism just described can be generalised to cases in which there are more than two divisions, and to cases in which the decision is more complicated (e.g., about the quantity to be supplied, or about the qualities of the product). (See Groves and Loeb, 1979.) It can also be modified to allocate costs from cost centres to profit centres.

On the other hand, there are problems with the Groves mechanism. First, as we have just seen, in equilibrium each manager is credited with the full profit to the firm that flows from the transfer. This provides the right incentives to the division managers, but in a sense 'overestimates' the contribution that each division is making to the total firm profit. This may be a problem if the division profit is used to compare the manager's performance with some 'standard.'

Second, it can be shown that the Groves mechanism is not immune to collusion between the managers against the headquarters.

Third, it can be shown that the Groves mechanism may fail to lead to truthtelling as a dominant strategy if (1) the managers have only imperfect information about their respective benefits or costs at the time the transfer has to be made, and (2) the customer's benefit and his information about it are suitably correlated with the supplier's cost and his information about it (Williams and Radner, 1985).

Finally, it seems that Groves mechanisms are not widely (or even significantly) used in practice. That may be because they constitute a recent invention, or it may be because of the problems I have described.

VI. HIERARCHY AND PRINCIPAL-AGENT RELATIONSHIPS

The most prevalent form of authority relationship within firms is no doubt hierarchical. I shall not try to give a formal definition of hierarchy here, but we are all familiar with the picture of an inverted tree of relationships between superiors and subordinates. In the firm, this inverted tree starts at the top with the board of directors (representing the owners), and goes down through the chief operating officer, several levels of higher and lower management, foremen, and then workers. (Here 'worker' will include those who work in staff organisations at higher levels as well as workers on the shop floor and salespersons in field.)

I am not aware of a well-developed theory that satisfactorily explains the

prevalence of hierarchy. Note that hierarchy is not the same as division of labour; there are, in principle, alternative ways to organise the division of labour within the firm. Those few economic theorists who study hierarchy tend to take its presence for granted, and try to understand the way it operates. For the most part, this has meant studying models of the relationship between a superior and a subordinate, or between an employer and an employee (Simon (1953) was perhaps the first to do this in a formal way).

Recently, the *principal-agent model* has provided a paradigm for the study of the superior – subordinate relationship. This model focuses on the differences in information available to the principal (superior) and the agent (subordinate), and on the difficulty the principal may have in monitoring the agent's actions and information. This leads to a potential conflict between (1) the need to share risk optimally between the principal and the agent, and (2) the need to provide the agent with effective incentives. An analysis of the model shows how this conflict can lead to a loss of efficiency compared to what is potentially attainable with the given distribution of information. This loss of efficiency is also known as 'loss of control.' The analysis of hierarchies using the principal-agent model was first done by Mirrlees (1976).

I shall here illustrate the principal-agent model in terms of a relationship between an owner and a manager. The owner of an enterprise wants to put it in the hands of a manager. In each of a number of successive periods (months, quarter, year) the profit of the enterprise will depend both on the actions of the manager and on the enviroment in which he is operating. The owner cannot directly monitor the manager's actions, nor can he costlessly observe all of the relevant aspects of the environment. The owner and the manager will have to agree on how the manager is to be compensated, and the owner wants to pick a compensation mechanism that will motivate the manager to provide a good return on his investment, net of the payments to the manager.[1]

Early formal analyses of principal-agent relationship (Ross, 1973; Mirrlees, 1974, 1976) focused on 'short-term' or 'one-period' situations. The simplest of these can be modelled as a two-move game, in which (1) the principal chooses a compensation function (made known to the agent) and (2) the agent chooses the action.[2]

[1] Some other principal-agent relationships in economic life are client – lawyer, customer – supplier, regulator – public utility, and insurer – insured. The insurer – insured relationship is the one that gave rise to the term 'moral hazard', and the first formal economic analysis of moral hazard was probably given by K. J. Arrow (1963, 1965).

[2] More complex models can be formulated to represent situations in which one or more of the following possibilities exist: (1) the principal may obtain some (incomplete or imperfect) information about the environment before choosing the compensation function; (2) the agent may do likewise before choosing his action; and (3) the principal may obtain some information about the agent's action and/or the environment before paying the compensation to the agent. In case (1), the principal's strategy will determine how his choice of compensation function depends on the information he receives. In case (2), the agent's strategy will determine how his action depends on his own information as well as on the announced compensation function. In case (3), the compensation may be made to depend on the additional information available to the principal, as well as on the outcome (i.e., the compensation function may have more arguments). In fact, in many cases a suitable reinterpretation of the constituents of the 'simple' model enables one to accommodate the 'complexities' described above within its framework! For example, if the agent obtains information about the environment before choosing his action, we may reinterpret his 'action' to be a *decision function* that determines his actual action as a function of his information,

Limitations of time and space prevent me from giving a detailed analysis of this game here. (See Grossman and Hart, 1983, for an extended analysis of the model, and references to previous work.) I can, however, give an intuitive explanation of the loss of control alluded to above. Suppose that the principal is neutral towards risk and the agent is averse to risk. Optimal risk-sharing would require that the principal bear all the risk, in which case the agent's compensation would be independent of the outcome of his action. On the other hand, if the principal cannot accurately monitor the agent's information and actions, then in order for the agent to have an incentive to perform well, his compensation must be made to depend on the outcome of his actions, e.g., on the realised profits that the manager produces for the owner. An optimal compromise between these two competing goals satisfies neither goal completely, so that with the best available compensation function the agent must bear some risk, and yet not have incentives that induce him to take actions that are fully optimal for the firm. This argument is most clearly applicable to a *short-term* or one-time-only relationship. The case of a *long-term* relationship can be quite different, as we shall see presently.

An exception to the proposition occurs if the agent is neutral towards risk and is sufficiently wealthy. In this case, an efficient equilibrium is obtained if the principal sells the agent a 'franchise' to the enterprise, i.e., the agent pays the principal a fixed fee, and then keeps the entire outcome.

Are there any remedies for the inefficiency of equilibrium in the principal-agent relationship? One possible remedy is for the principal to expend resources to monitor the agent's action (and, more generally, his information and environment). Whether or not this will improve net efficiency will depend, of course, on the cost of monitoring. The prevalence of *de facto* decentralisation in large organisations suggests that accurate monitoring of agents' actions is too costly to be efficient, or even practicable.

There are several ways in which the two players can exploit a *long-term* principal-agent relationship to escape, at least partially, from the inefficiency of short-term equilibria. A long-term relationship can be modelled as a *supergame* – a many-period game in which the one-period game is repeated over and over. These repetitions give the principal an opportunity to observe the results of the agent's actions over a number of periods, and to use some statistical test to infer whether or not the agent was choosing the appropriate action. The repetitions also provide the principal with opportunities to 'punish' the agent for apparent departures from the appropriate action. The decision rule that the principal uses to adjust his compensation function in any one period in the light of previous observations constitutes his *supergame strategy*. Likewise, the agent will have a supergame strategy for adjusting his actions in the light of the past history of the process.

One can show that there are typically many equilibria of the supergame, even if there is only one equilibium of the one-period game. One can also show that the set of supergame equilibria typically depends on the rates at which the players discount future expected one-period utilities. Finally, if the players' discount rates are small, then there are supergame equilibria that are approximately

efficient, and if the players do not discount the future at all, then there are supergame equilibria that are fully efficient.

In general, the players' strategy spaces are very large, and contain very complex strategies. Little progress has been made in characterising those supergame equilibria that are exactly most efficient in the set of equilibria for any particular pair of players' discount rates. Nevertheless, one can prove the above propositions by analysing the implications of very simple strategies for the principal.

For example, I shall describe such a strategy for the owner–manager situation. First, consider the case in which the owner can *commit himself in advance* to a particular supergame strategy of the following form: the owner pays the manager a fixed compensation w per period until the end of the first period T in which the total of the gross returns in periods 1 to T fall below $T(r+w)$ by an amount at least s(where w, r and s are parameters of the owner's strategy). At the end of such a period T, the manager is replaced and the owner engages another one under the same regime. This can be interpreted as requiring the manager to produce a 'paper' gross return of $(r+w)$ each period (a net return of r), and also allowing him to add any surplus to a (paper) 'cash reserve' and requiring him to subtract any deficit. The manager starts with a positive 'cash reserve' equal to s, and is replaced the first time his cash reserve falls to zero. I call this a 'bankruptcy strategy'.

One can show (Radner, 1986) that, under rather weak conditions, for suitable choice of the parameters w, r, and s, if the manager optimises against the owner's strategy, and if the manager's discount rate is small, then the owner's long-run-average expected return will be close to r, and the expected tenure of the manager will be very long. Using this fact, one can show that there are equilibria of the supergame in which the long-run average utilities of the two players are close to one-period efficiency (Pareto optimality), provided their discount rates are sufficiently small.[1]

If precommitment by the owner to a supergame strategy is not feasible, the players must rely on equilibria that are 'self-enforcing'; one such family of equilibria is examined in (Radner, 1985*c*).

VII. THE PARTNERSHIP MODEL

In complex organisations the outcomes of organisational activity usually depend jointly on the actions of several decision-makers as well as on stochastic aspects of the organisation's environment, and it is often difficult if not impossible to identify separately the individual contributions of the several decision-makers and the environment to the resulting outcome. With the

[1] In the supergame equilibrium, the owner's strategy may be more complicated than the bankruptcy strategy. Nevertheless, the bankruptcy strategy provides a lower bound on what the owner can achieve for himself, while guaranteeing the manager a target level of expected utility. Note that since the 'cash reserve' is only an accounting fiction, the bankruptcy strategy is really only a scoring formula for evaluating the manager's long-term performance, together with a criterion (based on the manager's score) for ending his tenure.

decentralisation of decision-making is associated serious imperfections in the monitoring of individuals' information and actions. The resulting confounding of the effects of the actions of different decision-makers, which we might call the *free-rider* phenomenon, is added to the phenomenon of moral hazard that is aleady present in the principal-agent relationship. This combination of the free-rider and moral-hazard phenomena has been analysed in what I have elsewhere called the 'partnership model.'

The general partnership model is a game in which the players – called 'partners' – share an outcome that depends jointly on the partners' actions and on a stochastic environment. The partners cannot fully monitor each others' information or actions, or the environment. The structure of information is decentralised, and the model explicitly takes account of the partly conflicting goals of the decision-makers. It is assumed that each partner's utility depends directly on his own action and on the size of his share of the (stochastic) outcome.

An *equilibrium* is a combination of partners' actions such that no partner can increase his expected utility by *unilaterally* changing his action. One can show that, under 'realistic' conditions, an equilibrium is not efficient, not even if the partners are neutral towards risk (see Holmstrom, 1982; Radner, 1985*a*). To get a feeling for why equilibria of the partnership game are inefficient, it is convenient to think about the case in which the partners are neutral towards risk. More specifically, suppose that each partner's utility equals his share of the outcome minus the 'disutility' of his effort, and that the stochastic outcome is an increasing function of the partners' efforts. In an equilibrium, each partner will have chosen his level of effort so that, given the equilibrium level of the other partner's efforts, any further small increase in his effort would result in an increase in his expected share of the outcome that would just be balanced by the corresponding increase in his disutility of effort. However, such a small increase in his effort would result in an increase in the other partner's expected share of the outcome, without any concomitant increase in his disutility. Therefore, if *both* partners simultaneously increased their efforts by a small amount, *both* partners would be better off (i.e., have a higher expected utility) than in the equilibrium. Hence an equilibrium cannot be efficient.[1] This argument is easily extended to the case of more than two partners.

In economic jargon, each partner's effort produces a *positive externality* for the other partners, but if he is only interested in his own expected utility he has no incentive to take account of that externality in his own optimisation. To put it more colloquially, each partner tries to be a *free rider* on the efforts of others, with the result that in equilibrium the overall level of effort is too small for efficiency.

One can construct a repeated game corresponding to the one-period partnership game in a way that is analogous to the construction of the repeated principal-agent game. It appears, however, that the incentive to be a free rider in the repeated partnership game may not be so easily reduced as the moral

[1] The mathematically-minded reader will note here that this argument relies on the differentiability of the relevant functions.

hazard in the principal-agent game (given that the partnership game also has moral hazard built into it). I shall try to describe this phenomenon without describing the details of the repeated partnership game.

The first two propositions about the repeated partnership game correspond to analogous results for the repeated principal-agent game:

—If the partners do not discount the future too much, there are equilibria of the repeated game that are more efficient (better for both) than the one-period equilibrium.
—If the partners do not discount the future at all, they can sustain full one period efficiency as a supergame equilibrium.

Recall that in the repeated principal-agent game, the less the players discount the future the closer they can get to efficiency in supergame equilibria, approaching efficiency in the limit as they discount the future less and less. It has been shown (Radner *et al.*, 1983) that this need not be the case in the repeated partnership game:

—If the partners discount the future even slightly, they may not be able to get close to full efficiency in supergame equilibria, no matter how little they discount the future.

I have no intuitive explanation of why there may be no equilibria of the repeated partnership game that are almost efficient when the partners discount the future very little, nor do I know for what general class of games this is so. It is interesting, however, that if one adopts the notion of *approximate* equilibrium, then this difficulty disappears. Roughly speaking a combination of players' strategies is called an approximate equilibrium if no player can improve his own expected utility by more than a small amount by unilaterally changing his own strategy. One can show that, for any given desired degree of approximation, approximate equilibria behave much more nicely than strict equilibria do, as one changes the parameters of the problem. In particular, as the partners' discount rates approach zero, there are corresponding approximate equilibria that approach full efficiency. It seems that, in this situation, 'the best is the enemy of the good'.

VIII. CONCLUDING REMARKS

I began the theoretical part of this paper with a discussion of the core as a model of decision-making and transfer-pricing among profit centres in a firm. Recent research suggests that the core might be nonempty for a wide class of 'nonneoclassical' technologies provided that internal competition is not allowed. On the other hand, the core did not seem adequate as a model of truly decentralised decision-making, which led us to a consideration of bargaining theory and Groves mechanisms.

With regard to authority relationships, we examined the principal-agent and partnership models, in both their short-term and long-term versions. The

preliminary results reported in Section VI and VII suggest that long-term relationships may be less effective in partnerships than in principal-agent relationships, even when the decision-makers are not short-sighted. These results are in accord with my impressions of organisational practice. The folklore of organisational behaviour attaches great importance to measures of individual performance, as contrasted with group performance, even when these measures only imperfectly reflect the individual's true productivity.

However, partnership relationships are also important in organisational practice. Within firms the widespread inability to measure individual performance accurately forces organisations to rely in part on measures of group performance. In addition, many aspects of the *informal* organisation of firms can be interpreted in terms of partnership theory, and in long-lived firms the importance of good long-term cooperative relationships among individuals – both for their own careers and for the effectiveness of the organisation–is well-recognised.

The result mentioned at the end of Section VIII suggests that boundedly rational decision-makers may be able to sustain more efficient long-term equilibria than fully rational ones can. Of course, the boundedness of rationality can take many forms. In this case it takes the form that each player is satisfied with a slightly-less-than-optimal response to the other players' strategies, which leads to the concept of approximate equilibrium discussed in Section VII. Further research is required to explain how the players' notions of what is 'satisfactory' are determined, since these notions are likely to be endogenous in a long-term game. Furthermore, other forms of bounded rationality are likely to affect the equilibria of such games (see Radner, 1975*a*, *b*, 1979).

The fact that long-term games typically have infinitely many equilibria poses a problem for the theory, because it makes the predictions of the theory relatively imprecise. In fact, the situation is even more extreme than is implied by simply saying that the set of equilibria is infinite; equilibria of a long-term game are typically not isolated, so that an equilibrium will have other equilibria arbitrarily close to it. The situation is analogous to that of a ball on a slightly sticky horizontal surface; the ball will stay put wherever you place it, and small shocks to the ball will result in correspondingly small displacements.

This state of affairs is probably not a comfortable one for contemporary economic theory, which tends to try to explain observed economic states as locally stable equilibria in a system with a finite set of equilibria. On the other hand, a theory that leads to a continuum of equilibria is open to completion by adding considerations of history and chance to explain the particular equilibrium that is currently observed. This could lead the economic theory of organisation back towards the older institutional and historical approaches from which it rebelled not so many decades ago.

IX. BIBLIOGRAPHIC NOTE

Section III is based on (Radner, 1985*b*). Sections VI-VIII are based on (Radner,

1985*a*), which contains a fuller account of principal-agent and partnership models on a fairly nontechnical level.

This brief account of some theories of the internal economy of large firms has been deliberately selective. I have omitted important topics of recent research, including the theory of teams, vertical integration, internal labour markets, worker-managed firms, and rank in organisations. I have therefore included in the bibliography a number of items that were not mentioned in the text of the paper, but will direct the reader to material on those omitted topics. I have also included supplementary references on the topics that were discussed. However, I have not tried to be exhaustive, relying instead on key references that will give the interested reader an entry to the literature. (The book by Hess (1983) is particularly valuable in this respect.) This explains why many important papers and books have not been listed!

AT & T Bell Laboratories
Date of receipt of final typescripts: August 1985

References

Alchian, A. A. and Demsetz, H. (1972). 'Production, information costs, and economic organization.' *American Economic Review*, vol. 62, pp. 777–95.

Aoki. M. (1980). 'A model of the firm as a stockholder-employee cooperative game.' *American Economic Review*, vol. 70, pp. 600–10.

Arrow, K. J. (1963). 'Uncertainty and the welfare economics of medical care.' *American Economic Review*, vol. 53, pp. 941–73.

———(1965). *Aspects of the Theory of Risk-Bearing.* Helsinki: Yrjo Johanssonin Saatio, Lecture 3.

———(1972). 'The value of and demand for information.' In *Decision and Organization*, (R. Radner and C. B. McGuire, eds.), Amsterdam: North-Holland, Chapter 6.

———(1974). *The Limits of Organization. New York: Norton.*

Barnard, C.I. (1938). *The Functions of the Executive.* Chambridge: Harvard University Press.

Baumol, W. J. (1959). *Behavior, Value and Growth.* New York: Macmillan.

Beckermann, M. J. (1983). *Tinbergen Lectures on organisation theory.* Heidelberg: Springer-Verlag.

Berle, A. A., Jr. and Means, G. C. (1932). *The Modern Corporation and Private Property.* New York: Macmillan.

Chandler, A .D., Jr. (1962). *Strategy and Structure.* Cambridge: The MIT Press.

Chatterjee, K. and Samuelson, W. (1983). 'Bargaining under incomplete information.' *Operations Research*, vol. 31, pp. 835–51.

Clark, E. (1971). 'Multipart pricing of public goods.' *Public Choice*, vol. 11, pp. 17–33.

Coase, R. H. (1937). 'The nature of the firm.' *Economica*, vol. 4, pp. 386–405, reprinted in *Readings in Price Theory*, (K. E. Boulding and G. J. Stigler, eds.,) Chicago: Irwin, 1952.

Cyert, R. M. and March, J. G. (1963). *Behavioral Theory of the Firm.* Englewood Cliffs: Prentice-Hall.

Debreu, G. and Scarf, H. E. (1963). 'A limit theorem on the core of an economy.' *International Economic Review*, vol. 4, pp. 235–36.

Demski, J. S. and Feldman, G. A. (1978). 'Economic incentives in budgetary control systems.' *The Accounting Review*, vol. 53, pp. 336–59.

Doeringer, P. B. and Piore, M. J.(1971). *Internal Labor Markets and Manpower Analysis.* Lexington: Heath.

Edgeworth, F. Y. (1881). *Mathematical Psychics.* London: Kegan Paul.

Gordon, R. A. (1945). *Business Leadership in the Large Corporation.* Washington, D. C., reprinted by the University of California Press, Berkeley, 1961.

Green, J. R. and Laffont, J.-J. (1979). *Incentives in Public Decision-Making.* Amsterdam: North-Holland.

Grossman, S. J. and Hart, O. D. (1983). 'An analysis of the principal-agent problem.' *Econometrica*, vol. 51, pp. 7–45.

Groves, T. (1973). 'Incentives in teams.' *Econometrica*, vol. 41, pp. 617–31.

———and Loeb, M. (1975). 'Incentives and public inputs.' *Journal of Public Economics*, vol. 4, pp. 211–26.

———and ———(1979). 'Incentives in a divisionalized firm.' *Management Science*, vol. 25, pp. 221–30.

Harsanyi, J. C. (1967–68). 'Games with incomplete information played by Bayesian players.' *Management Science*, vol. 14, pp. 159–82, 321–34, 486–502.

Hess, J. D. (1983). *The Economics of Organization.* Amsterdam: North-Holland.

Hildenbrand, W. (1974). *Core and Equilibria of a Large Economy.* Princeton: Princeton University Press.
Hirshleifer, J. (1956). 'On the economics of transfer pricing.' *Journal of Business*, vol. 29, pp. 172–84.
Holmstrom, B. (1982). 'Moral hazard in teams.' *Bell Journal of Economics*, vol. 13, pp. 324–40.
Ichiishi, T. (1982). 'Management versus ownership, I.' *International Economic Review*, vol. 23, pp. 323–36.
Kleindorfer, P. R. and Sertel, M. R. (1979). 'Profit-maximizing design of enterprises through incentives.' *Journal of Economic Theory*, vol. 20, pp. 318–39.
Knight, F. H. (1921). *Risk, Uncertainty and Profits.* Cambridge: Houghton Miffin Company, reprinted by Kelley and Millman, New York, 1957.
Levinthal, D. (1984). 'A survey of agency models of organizations.' Technical Report No. 443, Institute for Mathematical Studies in the Social Sciences, Stanford University, Stanford. (March).
Linhart, P. B., Radner, R. and Sinden, F. W. (1983). 'A sequential mechanism for direct price regulation.' Bell Laboratories Discussion Paper (unpublished).
Marris, R. (1964). *The Economic Theory of 'Managerial' Capitalism.* New York: Macmillan.
———and Wood, A. (1971). *The Corporate Economy.* Cambridge: Harvard University Press.
Marschak, J. and Radner, R. (1972). *Economic Theory of Teams.* New Haven: Yale University Press.
McGuire, C. B. and Radner, R. (1972). *Decision and Organization.* Amsterdam: North-Holland (to be reprinted by the University of Minnesota Press, Minneapolis, 1985).
Meade, J. (1972). 'The theory of labour-managed firms and of profit sharing.' ECONOMIC JOURNAL, vol. 82, pp. 402–28 reprinted in Vanek (1975).
Myerson, R. B. and Satterthwaite, M. A. (1983). 'Efficient mechanisms for bilateral trading.' *Journal of Economic Theory*, vol. 28, pp. 265–81.
Mirrlees, J. (1974). 'Notes on welfare economics, information, and uncertainty.' in *Essays on Economic Behavior under Uncertainty.* (Balch *et al.*, eds.), Amsterdam: North-Holland Publishing Company.
———(1976). 'The optimal structure of incentives and authority within an organization.' *Bell Journal of Economics*, vol. 7, pp. 105–31.
Miyazaki, H. (1984). 'Internal bargaining, labor contracts, and a Marshallian theory of the firm.' *American Economic Review*, vol. 74, pp. 381–93.
Radner, R. (1975*a*). 'A behavioral model of cost reduction.' *Bell Journal of Economics*, vol. 6, pp. 196–215.
———(1975*b*) 'Satisficing.' *Journal of Mathematical Economics*, vol. 2, pp. 254–62.
———(1979). 'Can Bounded Rationality Resolve the Prisoners' Dilemma,' Bell Laboratories Discussion Paper, Murray Hill, N. J. (to appear in A. Mas-Colell, ed., *Essays in Honor of Gerard Debreu*, Cambridge, Cambridge University Press, 1986).
———(1981*a*). 'Monitoring cooperative agreements in a repeated principal-agent relationship.' *Econometrica*, vol. 49, pp. 1127–48.
———(1981*b*). 'Optimal equilibria in some repeated partnership games with imperfect information' (to appear in the *Review of Economic Studies*).
———(1985a). 'Decentralization and Incentives.' In *Information, Incentives, and Economic Mechanisms.* (T. Groves, R. Radner, and S. Reiter, eds.) Minnesota Press (in press).
———(1985*b*). 'Transfer payments and the core of a profit-center game.' (unpublished).
———(1985*c*). 'Repeated principal-agent games with discounting.' *Econometrica*, vol. 53, pp. 1173–98.
———(1986). 'Repeated moral hazard with low discount rates.' In *Essays in Honor of Kenneth Arrow.* (W. Heller, R. Starr, and D. Starrett, eds.), Cambridge: Cambridge University Press (to appear).
———Myerson, R. and Maskin, E. (1983). 'An example of a repeated partnership game with discounting and with uniformly inefficient equilibria' (to appear in the *Review of Economic Studies*).
Rogerson, W. P. (1982). The structure of wage contracts in repeated agency models.' IMSSS Technical Report No. 388, Stanford, Calif.: Stanford University (unpublished).
———(1983). 'Repeated moral hazard.' Stanford, Calif.: Economics Department, Stanford University (unpublished).
Ross, S. (1973). 'The economic theory of agency: the principal's problem.' *American Economic Review*, vol. 63, pp. 134–9.
Savage, L. J. (1954). *The Foundations of Statistics.* New York: Wiley.
Simon, H. A. (1953). 'A formal theory of the employment relationship.' *Econometrica*, vol. 19, pp. 293–305.
———(1957). *Models of Man.* New York: Wiley.
———(1979). 'Rational decision-making in business organizations.' *American Economic Review*, Vol. 69, pp. 493–513.
Smith, A. (1776). *The Wealth of Nations.* Reprinted by the University of Chicago Press, Chicago (1976) (Cannan Edition).
Solomons, D. (1965). *Divisional Performance: Measurement and Control.* New York: Markus Wiener (reprinted 1983).
Sondermann, D. (1974). 'Economies of scale and equilibria in coalition production economies.' *Journal of Economic Theory*, vol. 8, pp. 259–91.
Spence, A. M. (1975). 'The economics of internal organization: an introduction.' *Bell Journal of Economics*, vol. 6, pp. 163–72.

Stiglitz, J. E. 'Risk, incentives, and the pure theory of moral hazard.' *The Geneva Papers on Risk and Insurance*, vol. 1, pp. 3–33.

Vanek, J. (1970). *The General Theory of Labor-Managed Market Economies.* Ithaca: Cornell University Press.

———(1975). *Self-Management: Economic Liberation of Man.* Harmondsworth: Penguin.

Vickrey, W. (1961). 'Counterspeculation, auctions, and competitive sealed-bid tenders.' *Journal of Finance*, vol. 16, pp. 8–37.

Ward, B. (1958). 'The firm in Illyria: market syndicalism.' *American Economic Review*, vol. 48, pp. 566–89.

Williams, S. R. and Radner, R. (1985). 'Informational externalities and the Groves mechanism.' (unpublished).

Willamson, O. E. (1963). 'Managerial discretion and business behavior.' *American Economic Review*, vol. 53, pp. 1032–57.

———(1971). 'Managerial discretion, organization form, and the multi-division hypothesis.' In Marris and Wood (1971) Chapter 11.

———(1975). *Markets and Hierarchies: Analysis and Antitrust Implications.* New York: Free Press.

Winter, S. G., Jr. (1971). 'Satisficing, selection, and the innovating remnant.' *Quarterly Journal of Economics*, vol. 85, pp. 237–61.

[Footnotes]

[2] **The Nature of the Firm**
R. H. Coase
Economica, New Series, Vol. 4, No. 16. (Nov., 1937), pp. 386-405.
Stable URL:
http://links.jstor.org/sici?sici=0013-0427%28193711%292%3A4%3A16%3C386%3ATNOTF%3E2.0.CO%3B2-B

[3] **The Nature of the Firm**
R. H. Coase
Economica, New Series, Vol. 4, No. 16. (Nov., 1937), pp. 386-405.
Stable URL:
http://links.jstor.org/sici?sici=0013-0427%28193711%292%3A4%3A16%3C386%3ATNOTF%3E2.0.CO%3B2-B

[1] **Uncertainty and the Welfare Economics of Medical Care**
Kenneth J. Arrow
The American Economic Review, Vol. 53, No. 5. (Dec., 1963), pp. 941-973.
Stable URL:
http://links.jstor.org/sici?sici=0002-8282%28196312%2953%3A5%3C941%3AUATWEO%3E2.0.CO%3B2-C

References

LINKED CITATIONS

Production, Information Costs, and Economic Organization
Armen A. Alchian; Harold Demsetz
The American Economic Review, Vol. 62, No. 5. (Dec., 1972), pp. 777-795.
Stable URL:
http://links.jstor.org/sici?sici=0002-8282%28197212%2962%3A5%3C777%3APICAEO%3E2.0.CO%3B2-Q

A Model of the Firm as a Stockholder-Employee Cooperative Game
Masahiko Aoki
The American Economic Review, Vol. 70, No. 4. (Sep., 1980), pp. 600-610.
Stable URL:
http://links.jstor.org/sici?sici=0002-8282%28198009%2970%3A4%3C600%3AAMOTFA%3E2.0.CO%3B2-F

Uncertainty and the Welfare Economics of Medical Care
Kenneth J. Arrow
The American Economic Review, Vol. 53, No. 5. (Dec., 1963), pp. 941-973.
Stable URL:
http://links.jstor.org/sici?sici=0002-8282%28196312%2953%3A5%3C941%3AUATWEO%3E2.0.CO%3B2-C

Bargaining under Incomplete Information
Kalyan Chatterjee; William Samuelson
Operations Research, Vol. 31, No. 5. (Sep. - Oct., 1983), pp. 835-851.
Stable URL:
http://links.jstor.org/sici?sici=0030-364X%28198309%2F10%2931%3A5%3C835%3ABUII%3E2.0.CO%3B2-A

The Nature of the Firm
R. H. Coase
Economica, New Series, Vol. 4, No. 16. (Nov., 1937), pp. 386-405.
Stable URL:
http://links.jstor.org/sici?sici=0013-0427%28193711%292%3A4%3A16%3C386%3ATNOTF%3E2.0.CO%3B2-B

A Limit Theorem on the Core of an Economy
Gerard Debreu; Herbert Scarf
International Economic Review, Vol. 4, No. 3. (Sep., 1963), pp. 235-246.
Stable URL:
http://links.jstor.org/sici?sici=0020-6598%28196309%294%3A3%3C235%3AALTOTC%3E2.0.CO%3B2-0

LINKED CITATIONS

An Analysis of the Principal-Agent Problem
Sanford J. Grossman; Oliver D. Hart
Econometrica, Vol. 51, No. 1. (Jan., 1983), pp. 7-45.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28198301%2951%3A1%3C7%3AAAOTPP%3E2.0.CO%3B2-A

Incentives in Teams
Theodore Groves
Econometrica, Vol. 41, No. 4. (Jul., 1973), pp. 617-631.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28197307%2941%3A4%3C617%3AIIT%3E2.0.CO%3B2-E

Incentives in a Divisionalized Firm
Theodore Groves; Martin Loeb
Management Science, Vol. 25, No. 3. (Mar., 1979), pp. 221-230.
Stable URL:
http://links.jstor.org/sici?sici=0025-1909%28197903%2925%3A3%3C221%3AIIADF%3E2.0.CO%3B2-D

Games with Incomplete Information Played by "Bayesian" Players, I-III. Part I. The Basic Model
John C. Harsanyi
Management Science, Vol. 14, No. 3, Theory Series. (Nov., 1967), pp. 159-182.
Stable URL:
http://links.jstor.org/sici?sici=0025-1909%28196711%2914%3A3%3C159%3AGWIIPB%3E2.0.CO%3B2-P

On the Economics of Transfer Pricing
Jack Hirshleifer
The Journal of Business, Vol. 29, No. 3. (Jul., 1956), pp. 172-184.
Stable URL:
http://links.jstor.org/sici?sici=0021-9398%28195607%2929%3A3%3C172%3AOTEOTP%3E2.0.CO%3B2-P

The Theory of Labour-Managed Firms and of Profit Sharing
J. E. Meade
The Economic Journal, Vol. 82, No. 325, Special Issue: In Honour of E.A.G. Robinson. (Mar., 1972), pp. 402-428.
Stable URL:
http://links.jstor.org/sici?sici=0013-0133%28197203%2982%3A325%3C402%3ATTOLFA%3E2.0.CO%3B2-W

LINKED CITATIONS

Internal Bargaining, Labor Contracts, and a Marshallian Theory of the Firm
Hajime Miyazaki
The American Economic Review, Vol. 74, No. 3. (Jun., 1984), pp. 381-393.
Stable URL:
http://links.jstor.org/sici?sici=0002-8282%28198406%2974%3A3%3C381%3AIBLCAA%3E2.0.CO%3B2-Y

Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship
Roy Radner
Econometrica, Vol. 49, No. 5. (Sep., 1981), pp. 1127-1148.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28198109%2949%3A5%3C1127%3AMCAIAR%3E2.0.CO%3B2-T

Repeated Principal-Agent Games with Discounting
Roy Radner
Econometrica, Vol. 53, No. 5. (Sep., 1985), pp. 1173-1198.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28198509%2953%3A5%3C1173%3ARPGWD%3E2.0.CO%3B2-M

The Economic Theory of Agency: The Principal's Problem
Stephen A. Ross
The American Economic Review, Vol. 63, No. 2, Papers and Proceedings of the Eighty-fifth Annual Meeting of the American Economic Association. (May, 1973), pp. 134-139.
Stable URL:
http://links.jstor.org/sici?sici=0002-8282%28197305%2963%3A2%3C134%3ATETOAT%3E2.0.CO%3B2-D

A Formal Theory of the Employment Relationship
Herbert A. Simon
Econometrica, Vol. 19, No. 3. (Jul., 1951), pp. 293-305.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28195107%2919%3A3%3C293%3AAFTOTE%3E2.0.CO%3B2-2

Rational Decision Making in Business Organizations
Herbert A. Simon
The American Economic Review, Vol. 69, No. 4. (Sep., 1979), pp. 493-513.
Stable URL:
http://links.jstor.org/sici?sici=0002-8282%28197909%2969%3A4%3C493%3ARDMIBO%3E2.0.CO%3B2-%23

LINKED CITATIONS

Counterspeculation, Auctions, and Competitive Sealed Tenders
William Vickrey
The Journal of Finance, Vol. 16, No. 1. (Mar., 1961), pp. 8-37.
Stable URL:
http://links.jstor.org/sici?sici=0022-1082%28196103%2916%3A1%3C8%3ACAACST%3E2.0.CO%3B2-L

The Firm in Illyria: Market Syndicalism
Benjamin Ward
The American Economic Review, Vol. 48, No. 4. (Sep., 1958), pp. 566-589.
Stable URL:
http://links.jstor.org/sici?sici=0002-8282%28195809%2948%3A4%3C566%3ATFIIMS%3E2.0.CO%3B2-Q

Satisficing, Selection, and The Innovating Remnant
Sidney G. Winter
The Quarterly Journal of Economics, Vol. 85, No. 2. (May, 1971), pp. 237-261.
Stable URL:
http://links.jstor.org/sici?sici=0033-5533%28197105%2985%3A2%3C237%3ASSATIR%3E2.0.CO%3B2-J