Incentives, Information, and Organizational Design

Joseph E. Stiglitz*)

Most economic activity — and almost all production — takes place within the context of organizations: Firms, family, and government. In each, there is a large element of collective decision-making. Our mixed attitudes towards this collective decision-making are reflected by a number of popular aphorisms and jokes: While we may claim that “two heads are better than one”, it is also true that “too many cooks spoil the broth”. There is the old joke about an elephant being a horse made by a committee.

The fact that we continue to make so many decisions collectively, that we so frequently assign difficult problems to committees, is suggestive that this form of decision-making must have some virtues, to compensate at least partially for its well recognized vices.

For the past several years, I (together with my co-author, Raaj Sah of Yale University) have been engaged in a research programme attempting to understand better the behaviour of economic organizations(1). We have focused in particular on decision-making. Our concern has not only been with the structure of decision-making within organizations, such as firms, but more broadly, the organization of decision-making in society as a whole.

The contrast between our perspective, and the standard description of “the economic problem” should be clear: We have been concerned not so much with the *rules* for allocating resources (e. g., setting marginal rates of substitution equal to marginal rates of transformation) as with how decisions about resource allocations get made, who makes those decisions, and how those who make those decisions are selected. We are concerned, in a sense, with the *process* by which resources get allocated.

There is a good reason that standard economic theory has paid little attention to these problems. In the standard formulation, there is a fixed set of resources (endowments), and fixed preferences; there is a once-and-for-all problem of allocating resources. Arrow and Debreu had the insight to see that, within that formulation, it made little difference whether there was a static, one-period resource allocation problem or a multi-period resource allocation problem. That result, in itself, should have alerted us that something was seriously flawed about the Walrasian perspective. (This is not to say that the Arrow-Debreu work has not been very useful, but, in retrospect, more in helping us to identify what is wrong with the standard competitive paradigm and in developing the market failures approach to gov-

*) Professor of Economics, Stanford University. NÖG Lecture delivered to the annual meeting of the Austrian Economic Society, Vienna, September 29, 1988. Financial support from the Olin Foundation, the National Science Foundation, and the Hoover Institution is gratefully acknowledged. This paper is partly based on joint work undertaken with Raaj K. Sah, to whom I am most indebted. I have also benefited greatly from conversations with Joseph Farrell. Research assistance of Serge Marquié and Rohit Rahi is gratefully acknowledged.

ernment interventions, than in helping us to understand how resources are actually allocated in modern capitalist economies.)

There is another result that should have had a similar effect: The Lange-Lerner-Taylor theorem, asserting the equivalence of a market socialist economy and a capitalist economy. Given that the theorem was based on an inaccurate description of the market economy and an even worse description of socialist economies, it is not surprising that the result, suggesting an equivalence between the two, should seem so far off the mark. Yet that theorem has not been without influence, as several socialist economies have tried to approach market socialism. The problems encountered, say, by the Chinese even in their limited attempts to move in the direction of market socialism should make us cautious, not only about the policy relevance of the theory, but more fundamentally, about the adequacy of the theory itself(2).

With there being only a once-and-for-all decision concerning the allocation of resources, no wonder that not much attention was paid to the process by which resources get allocated. Indeed, the question of who should be in decision-making roles — a question that occupies much of the time of businessmen, and even academics — is of utterly no importance. The view of economics encapsulated in the Arrow-Debreu framework (and reflected in the contemporaneously written textbooks, such as Samuelson) is what I call “engineering economics” (although in doing so, I intend no slur on the engineering profession). It was, perhaps, most accurately put by Joan Robinson, when she described the job of the manager of a firm as looking up in the book of blueprints the correct page corresponding to current (and possibly future) factor markets.

Even those much heralded results concerning the decentralizability of the economy were not so much results concerning the decentralization of decision-making — for there was really little scope for decision-making within the Arrow-Debreu framework(3) — as they were descriptions of a computational algorithm.

I cannot, within the confines of this lecture, develop fully the alternative approach that we have been exploring. My objective is more limited: First, to clarify some of the important failures of the traditional competitive paradigm; secondly, to describe briefly some of the major ingredients in our alternative approach; and thirdly, to show how our approach yields new insights, both in the new questions which it raises and in the new answers to long standing questions which it provides. We focus our attention in Section 3 in particular on the role of time in resource allocation and on the choice of centralized versus decentralized resource allocation mechanisms(4). Among the questions about which our analysis will have something to say is: Why is it that, in spite of the praise that is traditionally heaped upon the price system — and the abuse to which bureaucratic behaviour is subjected — in times of war almost all countries abandon reliance on the market and resort to some system of direct controls? This seems all the more curious, since it is precisely in times of war that resources need to be husbanded carefully. Moreover, our conventional textbook sto-

Incentives, Information, and Organizational Design

ries stress the role of prices in conveying information(5). When the market at one date looks much the same as on previous dates, there is little “news” to be conveyed. It is in times of change and massive resource reallocations — such as associated with war — when we should have thought that the price system would show its mettle, its distinct advantage (if indeed it has such an advantage) over other ways of conveying information. And yet this is precisely the time when country after country seems to lose its faith in the market.

1. Some failures of the traditional paradigm

There are several obvious — but no less fundamental on that account — criticisms of the traditional theory. I shall touch upon four in this lecture.

1.1 The complexity of the commodity space

The first has to do with the complexity of the commodity space. In the examples we teach our students, we talk about apples, oranges, and wheat. But any farmer can tell you that there is no such thing as a price for an apple. The price depends on the kind of apple, its freshness (and a variety of other quality characteristics), its location, and the time of the year. Industrial commodities are even more difficult, having a myriad of relevant attributes(6).

This has two fundamental implications. First, it makes it virtually impossible for a central planner (a Walrasian auctioneer) to set prices, or to set prices in a way which adequately reflects this diversity of characteristics, which results in products of the right characteristics being produced. Market socialist economies have learned the hard way what happens when the product is incompletely specified. If a price is specified for “nails” short nails made out of any cheap material will be produced. If the length is specified, the producer may still make nails out of a cheap material, which may be excessively brittle. For more complex commodities, almost no matter how many characteristics are specified, there remain scope for discretion, and in particular, cost cutting which adversely affects how well the commodity performs the task for which it is intended. Moreover, it becomes extremely costly to provide complete specifications of very complex commodities (see Note 6). Moreover, if all the inputs (materials, etc.) are fully specified — for instance, the material of which the nail is to be composed — it forecloses opportunities for finding alternative materials which meet the user’s needs as well or better, but which are less expensive. If only the characteristics of the nail were specified (brittleness, hardness, etc.), it often becomes a matter of judgment about whether these standards have been met. And even then, there remain questions of trade-offs: Some material might exceed the original standard in some characteristic, and fall short in another. What price should the producer receive for such a commodity? Market socialism provides no answer — other than requiring the planner to

provide a complete set of prices (an impossible task). The market economy — which is not a pure price system — provides an answer, through the interaction of buyers and sellers in real markets, in which prices are negotiated, not taken as given.

Secondly, it means that markets are frequently — perhaps I should say usually — imperfectly competitive. The products produced by one firm usually differ slightly, in one of the many characteristics, from those produced by others.

The process of production is often more one of "negotiation" than of "price-taking". Firms negotiate delivery times, product characteristics, as well as price. Information (about the needs of the buyers, the technological capabilities of the sellers) is transmitted in the process. Prices play a vital role in this transmission. The qualitative statement, "it would be hard to make a nail which will do what you want it to do", becomes a quantitative statement, "I can do it, but the cost of the nail will be \$ 1.23 per nail". There is competition: The buyer will check with other producers, to see if they can make a better offer. But it is not the kind of competition described by the Arrow-Debreu model.

1.2 Technological change

Nowhere are the failures of the traditional price-taking paradigm more apparent than in the analysis of technological change and of the allocation of time. I shall also have more to say about the allocation of time later in this lecture; here let me make a few comments about technological change.

It is not only that the standard competitive model does not address the problems associated with technological change, and in particular the development of new products; but the framework of that model cannot be extended to incorporate technological change, at least in a meaningful way. How can a firm be a price taker for a commodity, which has yet to be invented? And to obtain a return on its R&D, firms must exercise at least some monopoly power, as Schumpeter long ago emphasized. Moreover, the technical assumptions, such as convexity, underlying standard competitive analysis, are simply inappropriate when there is technological change(7). Many of the problems here are closely related to problems that costly and imperfect information poses for the standard analysis of market economies, to which I shall turn shortly(8).

There are many important links between organizational design, incentives, and technological change — issues which are not even addressed within the standard model. High rates of technological change require greater institutional adaptability, and a standard argument in favour of decentralized systems(9) is their greater adaptability. At the same time, decentralized structures provide both the competition which is a spur to innovation and the close interaction between producers and users which helps direct innovation in an effective manner(10).

Incentives, Information, and Organizational Design

1.3 Information and errors

The third problem with the standard paradigm I have already hinted at: It poses the resource allocation problem as a once-and-for-all decision. In fact, new products, new information, new individuals are constantly entering the market. These events cannot be fully anticipated. As new events occur, new opportunities are opened up and new decisions have to be made. Different economic systems will differ in the speed with which they can respond to the new conditions and opportunities, and in how well they respond, including the errors that they make.

Though it has long been recognized that "to err is human", until recently there have been few attempts to explore systematically the causes of human fallibility and its consequences for the design of organizations. In the Arrow-Debreu framework, individuals may have differing subjective beliefs about the likelihood of different states, but there is no such thing as "right" or "wrong", no attempt to come to terms with the differences in judgments, no process by which views on the likelihood of different states would change.

This is, of course, not the view taken by businessmen: They do not believe that one opinion is necessarily just as good as any other. Firms must decide how much and what information to gather, they must decide on how that information is to be gathered, e. g., to what extent the process of information gathering should be decentralized; and beyond that, they must decide *who* should gather the information and upon whose judgments they should rely.

Inevitably, or at least with a high probability, there will be differences in views: Some will think the project worth undertaking, others that it will not be. Any organization has to have a way of resolving these differences, a way, in other words, of aggregating the beliefs of the different participants.

In designing a structure (or as we called it in our earlier work, the architecture) of decision-making, account must be taken of two aspects of individuals and their relations with each other:

1. Individuals have a finite capacity. (I am particularly aware of this in observing others.) In particular, within any time they can collect only a limited amount of information. Time is of the essence in decision-making, partly because the relevant information is frequently dated: One typically can collect information about current relative prices, but only very limited information about future relative prices; information about relative prices can quickly become dated.

Because of their limited capacities, individuals never have perfect information, and therefore there is always a positive probability of an error of judgment. (Errors arise, of course, not only from the limited collection of information, but also from imperfect processing of the available information, again characteristics which most of us have become

accustomed to noticing in others, if not in ourselves.) The heuristics individuals use in making judgments, in arriving at conclusions from the available information, are particularly error prone when used in contexts which differ from those for which they were originally developed; but judging whether one is in such a situation is itself a question in which opinions may differ, and in which errors frequently arise(11).

2. Communication is limited and imperfect. An individual can seldom communicate all the information he has learned to another individual, and what he communicates is often transmitted with noise; errors enter in the transmission process(12). I cannot give you (to use the term of computer specialists) a "dump" of what is in my brain. I always sense that my understanding of a subject is far better and more complete than what I can communicate to you: You will only be able to know but a fraction of what (I believe) I know.

Concentration of information gathering and decision-making in small groups saves on communication costs and reduces the problems arising from communication errors. But the finiteness of the capacities of individuals means that concentrating decision-making may result in decisions based on more limited information, and hence more subject to errors (on that account).

What is critical is not only the number of individuals involved in the decision-making process, but how they are arranged. In earlier studies, we contrasted the errors arising from hierarchical decision structures (in which, for instance, project approval required approval by all members of a hierarchical chain) with those arising from polyarchical decision structures, in which any unit could approve the project on its own. (Our analysis of polyarchical structures was meant to be suggestive of market economies: In a market economy, any firm can undertake a project, and there are usually several firms in the same industry involved in reviewing similar kinds of projects.) In the simple models we examined, when the standards of acceptance of projects in both hierarchy and polyarchy are the same, so that the likelihood of any single individual (subunit of the organization) recommending that a good project be rejected or a bad project be accepted is the same in both organizational forms, hierarchical decision-making led to more good projects being rejected, but fewer bad projects being accepted. Both of these organizational forms could be viewed as special cases of committee structures, where the size of the committees and the degree of required consensus depends on the costs of various kinds of errors as well as the likelihood that any single decision maker would make a particular error. We showed, for instance, that the more important the decision (that is, a decision where errors were more costly), the larger the size of the committee, explaining the sense of powerlessness often felt by individuals in modern society: As individuals earn the right to participate in more important decisions, their voice becomes only one of an increasingly larger number responsible for making these decisions.

Another critical aspect of the design of decision-making structures is delay. Hierarchical structures, because they require more approval, may be more subject to delay. Moreover,

Incentives, Information, and Organizational Design

the malfunction of any one part of the system may have severe consequences for the overall performance of the organization, and, in particular, lead to long delays.

Delay may be reduced by having a large number of individuals look at a project simultaneously (but independently). But this has a distinct cost disadvantage over the sequential review process typically associated with hierarchical structures(13).

The issues of organizational form and information (error) aggregation arise at all levels in society. We can think of the market economy as one large polyarchival organization, each of the units of which are themselves organizations, many of which are hierarchically organized. Our earlier analyses provided some hints at the circumstances under which one organizational form would be more desirable than another, and an explanation of why we should not be surprised to see the kind of mixture of polyarchy and hierarchy that in fact we observe(14).

1.4 Incentives

Perhaps the most widely recognized failure of the traditional competitive framework is its failure to give sufficient attention to the problem of incentives. Of course, in the obiter dicta surrounding the model, reference is frequently made to the incentives provided by competition. But the incentive problems which the model addresses are not those which, again, most businessmen view themselves as facing. In the traditional model, workers do not have to be motivated: They either do the job which they have contracted to perform, or they do not get paid the contracted amount. The pay that they receive does provide them, to be sure, with an incentive to perform the job, provided it compensates them for the disutility of the task and the foregone leisure.

But the incentive problem which firms face arises from the fact that few individuals are paid on a piece rate basis, on the basis of performance alone. (The reasons for this I discussed extensively in my 1975 paper.)

This is particularly true of individuals involved in decision-making. After all, we have argued that most decision-making is done collectively. It is very difficult to assign responsibility for failures, or successes for that matter. And indeed, many, if not most individuals, go to great trouble to make it difficult to assign blame. They make sure that others have been appropriately consulted, and they are careful about what they put down on paper. Verbal communication is much more subject to reinterpretation — to claiming that, if one's advice turns out to be wrong, the advice was not correctly understood. When individuals put pen to paper, to leave, as the expression goes, a paper trail, they are careful that it reads like a Delphic oracle.

In some sense, the Arrow-Debreu model not only does not address the central incentive issues facing the economy, it gives the wrong impression about the relationship between

competition (at least the peculiar form of competition which goes by the rubric of "perfect competition") and incentives. In the case of technological change, providing incentives requires that the individual be able to appropriate some part of the social returns to the innovation — as the patent system does. But this in turn means that these markets cannot be perfectly competitive. Just as perfectly efficient capital markets would provide no incentive for individuals to acquire information (see *Grossman — Stiglitz*, 1976, 1980A), so perfectly competitive product markets would provide no incentives for innovation.

But while the traditional model of the Arrow-Debreu economy does not pay proper due to the importance of incentives, the Lange-Lerner-Taylor view of market socialism is far worse: Managers are supposed to maximize firm profits, at the established prices, simply because it is their job to do so. Of course, *if* their job was as simplistic as it typically was depicted — looking up the appropriate page in the book of blueprints — incentives might indeed have been a second order issue. But if our view of the importance of decision-making is correct, then incentives are important. And incentives require both property rights and at least the possibility of inequality: For if individuals' standards of living are guaranteed to be essentially the same, regardless of whether one works hard or not, there is little incentive to work hard. Incentive issues are concerned not just with effort, but also with risk-taking: So long as individuals do not bear all the consequences of the risks which their actions can affect, there will be imperfect incentives. These incentive questions are not addressed by the Arrow-Debreu model because, by assumption, individuals do bear all the consequences of their actions.

In recent years, there has been considerable research on the design of good incentive structures(15). These require balancing off, say, the gains from better incentives from relying more extensively on piece rates or more generally performance related compensation schemes, versus the costs in terms of greater risk-bearing by the worker.

One form of incentives which has received some attention is contests — a form of competition much more akin to the kind of competition we see in everyday life than the perfect competition of the standard paradigm. Patent races are contests where the winner takes all, although the loser is often in a good position to compete in the next race. Markets can be viewed as contests, in which the losers — those who can't keep up with the rest — are eliminated(16). The essential nature of these contests is that compensation depends on *relative* performance(17). I now want to say a few words about the virtues of contests, and the concomitant organizational architectures.

2. Contests and polyarchies

In the preceding section, we argued that the view of competition reflected in the Arrow-Debreu model does not provide a good model of modern capitalism; it does not even address some of the central issues facing market economies.

Incentives, Information, and Organizational Design

Modern market economies are characterized by many firms producing many different products. There is competition and there is decentralization of decision-making. A major part of my research programme over the past few years has been concerned with understanding better how competition and decentralized decision-making actually work in modern economies.

In our previous discussions, we have identified three tasks facing any organization: It must select those to be involved in decision-making (the *selection* problem); it must provide incentives; and it must have a decision-making structure, a structure which balances off the various errors, which trades off the costs of errors and of delays with the costs and benefits of obtaining faster and more accurate decisions.

In our work on decision-making structures, we identified some distinct advantages — from the perspective of error aggregation — of polyarchical structures. This was particularly true when the units of the polyarchy adjusted the standards for acceptance of a project to reflect the limited number of reviews to which it was subjected.

Polyarchical structures have two further distinct advantages: They provide a basis of comparison of performance, which can be used for the design of effective incentive structures, and which can be used as the basis of selection.

One of the problems facing standard piece rate systems is the determination of the appropriate piece rate. In contexts where the environment is changing rapidly, where technology is evolving, and where opportunity costs may be varying, the appropriate piece rate must constantly be changing. (Of course, if there were a complete set of prices, the piece rate would simply need to adjust to reflect the price changes. But this brings us back to the first point made earlier — the complexity of the commodity space. There is not a complete set of prices; there are not prices for a car with four thousand seven hundred and eighty two screws tightened and for a car with four thousand, seven hundred and eighty three screws tightened. There is no market price to which we can look for the value added of tightening one screw. And since most production occurs within large firms, most of the value added by individuals — when it can be identified — represents but one piece of the production of a marketed commodity.)

In joint work with Barry Nalebuff (*Nalebuff — Stiglitz, 1983A*), we showed how contests have the advantage of what we called *incentive flexibility*. It was *as if* the piece rate adjusted automatically to reflect the difficulty of the task. Consider a two-person contest. If the task became easier, if one kept his output target fixed, slacking off on effort, it would pay the other to increase his output target, to increase substantially the likelihood he would win. Of course, each participant knows this, and they all adjust their output targets to reflect the increased ease of the task.

Market contests work in a similar way. One of the problems that monopolies face is the difficulty of knowing whether they are producing efficiently. Of course, the standard theory

says that the problem with monopoly is only that it produces too little; whatever it produces it produces efficiently. But this does not seem to accord well with widespread perceptions of managerial slack in monopolies. The problem is that the owners — or the government regulators, in the case of regulated natural monopolies — do not know whether the managers are performing well; and the managers may not know whether their subordinates are performing well. How much should it cost to make a telephone call between New York and Chicago? How expensive should it be to develop a less expensive way of transmitting information? The best, and in some cases, the only meaningful answers can be derived by comparisons. But comparisons require several units undertaking similar tasks.

For some tasks, even a monopoly can set up bases of comparisons: It can organize contests among telephone operators. Polyarchically organized decision-making provides a basis of comparison of the competence of different decision-making units. Letting pay and selection for promotion depend on the comparative evaluations thus both provides incentives and facilitates the selection process.

(Contests have other virtues, and some vices, which we briefly note. The incentives they provide are provided with limited risk-bearing on the part of the participants. Sometimes competition is less effective than we have depicted it — workers band together to limit the degree of competition among themselves, labelling those who work too hard, who play the game as management intended, as rate busters; and sometimes it may even be destructive — one can win not by doing well but by ensuring that one's rival does badly; competition among first-year law students, with pages being ripped out of the relevant law journals to hamper rivals, is clearly not constructive; and Salop (see, e. g., *Salop — Scheffman*, 1983) has shown a variety of circumstances in which, with imperfect competition, efforts are directed at raising rivals' costs rather than lowering one's own costs and he has shown a variety of ways in which firms can and do raise rivals' costs.)

3. Time, incentives, and organizational design

Many, if not most contests that we see in day-to-day life involve time: Both who can do the most (make the most widgets) within a given amount of time, and who can do a particular task (run a quarter mile, swim 100 meters) in the fastest time.

Time is among our scarcest resources — our days on this earth are numbered — and while conventional economics has focused on how markets lead to efficient allocations of individuals' time between leisure and work, they have had little to say about many of the more interesting problems involving time.

In this section, I present four examples in which time is of the essence, in which the standard competitive models provide little if any insight into what is going on, but in which our alternative approach provides several insights. The examples illustrate the ways in which

Incentives, Information, and Organizational Design

markets — and time — address the fundamental problems of incentives, selection, and decision-making.

The first example focuses on incentives, and addresses the question, does the market provide appropriate incentives for R&D? Is the pace of research too fast or too slow?

The second example focuses on selection, and addresses the question, does the market ensure that the least-cost producers actually produce? We construct a simple model showing how time helps solve this problem: In equilibrium lower-cost firms enter markets, on average, at a more rapid rate. But the market's screening is far from perfect: There is some chance that high-cost firms enter even when there are lower-cost firms around. And the market's selection mechanism is far from costless: There is some chance that two firms will enter, even when economic efficiency entails only one firm.

While problems of coordination arise in our second example, issues of coordination and communication are central to our third and fourth examples. In our third example, we show how judgments about the worth of a project made in a decentralized economy are communicated indirectly, and not just by prices, but also by "time": The fact that no one else has undertaken a project conveys important information to any firm or individual contemplating undertaking it. This information affects the decision rules employed. As we shall see, our analysis provides an interesting interpretation of the oft noted observation that firms claim that they require high rates of return in order to undertake a project (after tax returns in the order of 15 percent to 25 percent), far higher than the realized ex-post returns.

The fourth example attempts to provide an answer to one of the questions we posed in the introduction: Why is it that in times of war so many countries abandon the use of the market?

3.1 Patent races: An example of a game of timing

Imagine how an Arrow-Debreu model, with a complete set of markets, might deal with the question, do market economies provide the correct incentives concerning the speed of R&D? Let me simplify the problem for the moment by ignoring risk. First, there would have to exist *prices* for commodities which have not yet been invented, perhaps not yet even conceived. (That, I admit, is a mind-boggling thought.) The commodities would have to be fully described, with a complete list of all their characteristics. (That, again, is a mind-boggling thought.) Next, these prices would have to be defined for each moment of time, that is, there would have to exist prices at all dates, so that the potential producer would know the extra returns to be obtained from inventing the invention a day earlier. Finally, each of these markets would have to be fully competitive: There would have to be many producers standing ready to produce at each date, to produce a commodity which has not yet been invented!

I present this example in some detail, lest there be any among you who have any lingering doubts about the applicability of the Arrow-Debreu paradigm to industrial economies in which technological change is important.

The returns to innovation are, as Schumpeter long ago recognized, the rents associated with monopoly or oligopoly. The question is, does the race to acquire these rents lead to too fast or too slow research? The answer to this question reflects the old joke about economics exams: The questions stay the same, only the answers change. Arrow (1962) suggested that, because in competitive markets, inventors do not capture the consumer surplus associated with the lower prices resulting from innovation, there will be too little innovation; and that with monopoly, this problem is exacerbated, because with monopoly output is lower (than under competition), and therefore the appropriable benefits from cost reductions (the change in costs times the quantity produced) is smaller than under competition. Stiglitz (1971), Barzel (1968), and Dasgupta — Stiglitz (1980A) then showed that there might be *excessive* expenditures on R&D. Their results can be seen as an early version of the by now common argument that in economies with rents, there will be excessive expenditures on rent seeking activity. The return to a patent is a rent. The race to get that rent will dissipate much of the value of the rent.

In the case of a patent race with an incumbent monopolist, Dasgupta and Stiglitz showed that not all rents would be dissipated: The incumbent would spend resources just to the point where it would not pay a rival to enter; the rival's potential profit, entering as a *duopolist*, must be zero, leaving a strictly positive profit for the incumbent monopolist(18).

But upon closer examination, this result in turn had to be qualified: In these earlier studies, decisions about R&D expenditures were once-and-for-all decisions, not made sequentially. Most races — including patent races — are appropriately modelled as entailing sequential decision-making. The participants look at their position, relative to rivals, to decide on how much resources to devote the next period to the race, or whether to drop out. Expenditures on R&D are, for the most part, sunk costs. One of the more important results to emerge from the recent research in industrial structure is the recognition of the importance of sunk costs; with sunk costs potential competition does not suffice, either to ensure economic efficiency or competitive outcomes (zero profits; see Stiglitz, 1988). An incumbent monopolist can position itself so that it can meet challenges from rivals: If it does this, they will have no incentive to enter an R&D contest. And if it is successful, the pace of innovation will — apart from an initial expenditure to put itself in the appropriate position to meet these challenges — be determined essentially as if it were a protected monopoly. Arrow's original conclusion, that markets may provide too little incentives for R&D, is restored, under more general conditions (for, unlike Arrow, we allow the presence of potential competitors), but for quite different reasons(19).

I present this example in some detail both because of its importance for understanding market economies and as an illustration of the general principle, upon which I now wish to expand, that market economies often do not deal well with time.

3.2 Entry races: A second example of a game of timing

In older (pre-Arrow-Debreu) discussions of the virtues and vices of capitalist economies, one criticism of capitalist economies that was often raised related to their inability to coordinate investment, and in particular, entry, decisions. If people believed that there was an increase in the demand for pizza, there might be a rush of entry into pizza parlors, resulting in excess expenditures and social waste. The Arrow-Debreu model made clear the nature of the market failure: There were not equilibrium prices extending even slightly into the future for most commodities, and hence investors had to form *expectations* of future prices. There was no coordinating mechanism to ensure that the appropriate amount of investment would be undertaken. In the early sixties, some economists and the governments that they advised were hopeful that indicative planning would provide that coordinating mechanism, but the failure of the French experiment(20) has, for the most part, dashed those hopes. The lack of attention to this problem in recent years is not because it has disappeared, but arises from the absence of any apparent solution, and the recognition of the marked disadvantages of alternative approaches involving larger governmental roles in coordination.

In this section, I want to argue that decentralized decision-making concerning entry (or to use the term I used early, polyarchy) has both distinct advantages and problems, both of which have received insufficient attention to date.

First, it may partially resolve the selection problem, to which I referred earlier: If there are several potential entrants, which should enter the market? Obviously, if we had a Walrasian auctioneer, the market price would serve to sort among the potential entrants; the firm(s) with the lowest costs would enter. But there is no Walrasian auctioneer.

Similarly, one might imagine a Central Planner, asking each of the potential producers to submit cost estimates, and on the basis of that deciding who should enter. But is there any reason to expect each to submit truthful estimates, or to engage in the requisite research to determine its costs accurately?

The market provides a solution, but not without a cost. Assume that there are fixed, sunk costs of entry S , and that marginal costs are constant, but may either be low, C_L , or high, C_H . Assume there are two potential entrants. Each knows its own costs, but not its rivals'.

With perfect information, the equilibrium would be easy to describe. There are three possible cases — both firms have low costs, both have high costs, or one has low costs, the other high costs. In the last case, the low-cost firm enters, the high-cost firm does not. In the other cases, *one* of the two firms enters (which one makes no difference), the other does not. For both firms to enter entails excessive expenditures on entry costs.

Equilibrium with imperfect information is markedly different. We assume that when both firms enter, prices are determined in a Bertrand game, that is, the price equals C_H if both are high-cost producers, C_L if both are low-cost producers, and $C_H - \varepsilon$ if one is high-cost and the other low-cost. If a firm has a low cost, it enters with probability $\pi_{1,L}$, if it has a high cost, it enters with probability $\pi_{1,H}$. If at date t , it observes no one has entered the previous period, it enters with probability $\pi_{t,i}$. If a low-cost firm enters, it charges price q_L , while if a high-cost firm enters, it charges price q_H . If a low-cost firm observes that his rival has entered and is charging a price q_H , he enters the next period. In this general structure, information is conveyed both by prices and by actions — or more accurately, inaction: The failure of one firm to enter leads the other firm to revise his estimate of whether the firm is a high- or low-cost firm; if the firm fails to enter, it appears more likely that the firm is a high-cost firm. This information obviously affects entry decisions.

We do not propose to provide a general solution to this problem, but rather to illustrate the general principles in a two-period simplified version. For simplicity, we assume the parameters of the problem are such that $\pi_{1,L} = 1$, and the q_i is the monopoly price for firm of type i . Furthermore, we assume that there are only two periods and that the probability of having a high cost is 0.5. Then, the high-cost firm must be indifferent whether to enter the first period or not. There is a fifty-fifty chance his rival will have a high cost, and if he has a high cost, the probability he enters is π_1 (without confusion, we can drop the second subscript). The possible outcomes are summarized in Table 1.

Table 1

High-cost firm enters			
Outcome	Probability	Period 1	Period 2
Rival is low-cost	0.5	0	0
Rival is high-cost and enters	$0.5 \pi_1$	0	0
Rival is high-cost and does not enter	$0.5 (1 - \pi_1)$	R	R

R is the monopoly profits of the high-cost firm. Thus, his expected profits are (ignoring discounting)

$$2 \times 0.5 (1 - \pi_1) R - S$$

and this must equal zero, i. e.,

$$1 - \pi_1 = S/R$$

or

$$\pi_1 = 1 - S/R.$$

When both have high costs, there is a chance $(= (1 - \pi_1)^2)$ that neither will enter in the first period. Then there is again a mixed strategy equilibrium the second period. By the same reasoning, π_2 is such that each firm is indifferent as to entering, i. e., since he only earns a return when the other firm does not enter:

Incentives, Information, and Organizational Design

$$R(1 - \pi_2) = S,$$

or

$$\pi_2 = 1 - S/R.$$

We now see the differences between the decentralized solution and the solution with perfect information:

1. *One firm low-cost, one firm high-cost:*

A fraction of the time, π_1 , the high-cost firm enters; resources equal to S are wasted.

2. *Both firms low-cost:*

Always excessive entry, with resources S wasted.

3. *Both firms high-cost:*

A fraction of the time $(1 - \pi_1)^2 (1 - \pi_2)^2$ no one enters.

A fraction of the time $(1 - \pi_1)^2 (2\pi_2 - \pi_2^2)$ entry is delayed.

A fraction of the time $\pi_1^2 + (1 - \pi_1)^2 \pi_2^2$ both firms enter, and resources S are wasted.

Delay is serving as a partial screening device(21). Entry is not random. But there is a cost to the screening. We could contrast this decentralized screening mechanism with that of a bureaucratic system, which chooses one firm to enter. At the extreme, one could imagine the bureaucracy as simply randomly choosing one firm to enter. Then there would never be delay, there would never be duplication, but a fraction of the time — 0.5 in our example — the high-cost firm would be chosen. Relative to perfect information, both systems make mistakes. Which system is more efficient — that is, which mistakes are more costly — depends on the parameters of the problem. The inefficiency of the bureaucratic system depends on $C_H - C_L$, the difference in costs. The inefficiency of the market system depends on three parameters: 1. the costs of duplication (essentially the magnitude of the sunk costs), 2. the costs of delay (the loss of consumer surplus when no entry occurs the first period, or the first and second period), and 3. the likelihood of duplication and delay, which depends critically on S/R , sunk costs relative to profits. When S/R is high, delay (no entry) is more likely; when S/R is small, duplication is more likely(22).

This result may provide a partial answer to the question of why, in times of war, economies so often resort to bureaucratic selection: Delay may be particularly costly. Moreover, in the *short* run, the selection problem is not as severe as in the long run. If a market economy does a reasonably good job in screening out bad firms (high-cost producers), then as the economy enters an emergency, it begins with a better than random sample of firms from which to choose to make production decisions. Organizations change, and individuals enter and leave, so that the task of screening is a perpetual one. For a while, the economy can

slide along, taking advantage of the screening that occurred prior to the emergency. But over time, the value of this prior screening declines, and accordingly the costs of the failure to screen effectively increase. This is undoubtedly part of the reason why the longer market processes are suspended in an emergency, the more likely it is that the economy will encounter increasing difficulties(23).

3.3 Markets as aggregators of information

Earlier, we identified three problems that all economic organizations must address: *Selection*, *information transmission and aggregation*, and *incentives*. The previous example illustrated how time (delay) could be used as a selection device. My next example illustrates how, in market contexts, time (delay) can be used as a way of transmitting and aggregating information. In earlier work with Sanford Grossman (*Grossman — Stiglitz*, 1976, 1980B) we showed how the price system conveyed and aggregated information. We are concerned here with a context in which the price system, by itself, effectively does neither.

Consider the simplest decision-making problem: An organization must decide whether the expected return to a project is positive or negative. Each individual (subunit) within the organization evaluates the project, but with error. There are three types of projects, very good projects (denoted with a subscript 0), mediocre (denoted with subscript 1), and bad (denoted with subscript 2). Evaluations are far from perfect. The probability that a project of type i will be evaluated as one of type j is denoted by p_{ij} . If different individuals are assigned the task of evaluating the project, there is some probability that their evaluations will differ. The question of information aggregation is, what to do in these circumstances. How are the conflicting reports to be brought together to make an organizational decision?

In a market context, firms face the following quandary. If they think a project is good, it is likely that some other firm will think similarly — if it is a good project. And if others think it is a good idea, then they are likely to enter, driving down the return. If the firm thinks the project is good (because it has erroneously evaluated it) and others think it bad, then it will be the only firm in the market — but because the project is not a good one, again returns are low. This is the same kind of logic that appears in auction markets, under the guise of the “winner’s curse”. In that context, it has been recognized that firms realize that there is information in the others’ bids; that when the individual wins, it means the others have bid lower, and that in turn means that their information about the project is not as positive as the winner’s information. Taking this into account, the winner bids a lower amount.

One can view the decision to enter the market as a “bid”. A firm takes into account the fact that when it is the only firm to enter, it is likely that its information is overly optimistic concerning the project’s returns. This provides an explanation of the commonly observed puzzle that firms *insist* on returns of 15 percent to 25 percent before undertaking projects, yet

Incentives, Information, and Organizational Design

ex-post realized returns average perhaps half that. Firms know that if they are to realize a 12 percent return, their information must indicate a far higher return.

These are examples of how non-price information — other firms' decisions not to enter a market — is taken into account in firm decision-making. The decision of the firm does reflect some of the information gathered by other firms.

I now want to draw attention to what might seem a perverse possibility: That in fact the best projects may not be undertaken as quickly as mediocre projects (and this is *not* because, as we academics sometimes feel, our best ideas, being the most novel, are hardest for others to grasp, while our more mediocre colleagues, with their more simplistic ideas, find it easier to sell their ideas). The reason for this has to do with the absence of a coordination mechanism: We know that if the project is really very good, it is likely that some other firm will think so too, and there is therefore a reasonable chance that both of us will enter, driving down the returns. The fear of excess entry for very good projects leads, on average, to delays in these projects being undertaken.

To see this most dramatically, assume projects, regardless of their quality, have the same fixed, sunk costs, S . To simplify the calculations, we assume that very good projects are either evaluated as very good, or as bad, and mediocre projects are either evaluated as mediocre or as bad. (It is straightforward but tedious to extend the analysis to the more general case.) Let p_0 denote the probability that a very good project will be evaluated as good, and p_1 the probability that a mediocre project will be judged to be mediocre. Firms are, of course, on the look-out not for the mediocre projects, but for the very good ones; the discovery of a mediocre project can be viewed as the unsuccessful outcome of a search for a very good project. Assume that (individuals believe that) more firms will be engaged in evaluating any particular very good project than in evaluating any particular mediocre project. (There are, of course, more mediocre projects around.) For simplicity we ignore here the costs of evaluation.

The return to a very good project, if only one firm undertakes it, is R_0 , and the return to a mediocre project (if only one firm undertakes it) is R_1 . When more than one firm undertakes a project, competition is sufficiently fierce that profits are driven to zero. If R_i is large relative to S , then equilibrium will be characterized by a mixed strategy: Only some of the time when a firm discovers a very good project or a mediocre project will it undertake it. (Firms may, of course, not view themselves as throwing dice to determine their mixed strategy; they may make judgments about the likelihood of other rivals entering, they may look at business forecasts or sunspots, or engage in other forms of "reasoning", the net effect of which is simply to randomize their decisions.)

Mixed strategies mean that a firm that observes a good or mediocre signal must be indifferent to undertaking and not undertaking the project. We focus on symmetric equilibria in which firms act as if they knew the number of other firms that have evaluated a project as

being good (mediocre). The analysis can easily be extended to the case where that number is viewed as a random variable. If π_i is the probability that a firm that believes that his project is of type i undertakes it, and if n_i is the number of firms that have evaluated a particular project of type i as being good, then

$$R_i \times (1 - \pi_i)^{n_i-1} = S.$$

Return if the firm is the only firm to undertake the project	Probability that no other firm enters
--	---

The probability that a project gets undertaken is

$$G = 1 - (1 - \pi_i)^{n_i} = 1 - (S/R_i)^{n_i/(n_i-1)}.$$

To simplify matters, let us assume that n is just a function of R , the expected return to the project. For instance, if there are M firms evaluating the same portfolio of projects, then the expected number of firms giving a favourable review to a project with return R is

$$(1) \quad n(R) = p(R) M,$$

$p(R)$ is what we called in our earlier work the *screening function*; it has the property that $p' > 0$, projects with a higher expected return are more likely to get a favourable approval.

Differentiating G with respect to R , we obtain

$$\frac{dG}{dR} = (1 - G) \left[\left(\frac{1}{n-1} \right)^2 \ln(S/R) \frac{dn}{dR} + \frac{n}{n-1} / R \right] > \text{ or } < 0$$

as

$$\frac{d \ln n}{d \ln R} < \text{ or } > \frac{n-1}{\ln(R/S)},$$

or, if n is simply related to R through the screening function, as described by equation (1)

$$\frac{p' R}{p} < \text{ or } > \frac{n-1}{\ln(R/S)}.$$

Thus, provided that the probability of getting a favourable evaluation goes up fast enough with the quality of the project (its expected return), the probability that a very good project will be undertaken in any period will actually decrease with its quality.

Incentives, Information, and Organizational Design

Of course, even if the probability that a project is undertaken does not decrease, the probability that a project will be undertaken, given that it has had a favourable review, will decrease with the quality of the project under somewhat less stringent conditions; that is

$$\frac{d\pi}{dR} = (1 - \pi) \left[\frac{\ln(S/R)}{(n-1)^2} \frac{dn}{dR} + \frac{1}{n-1} / R \right] < \text{ or } > 0$$

as

$$\frac{d \ln n}{d \ln R} > \text{ or } < \frac{n-1}{n \ln(S/R)},$$

or, if n is simply related to R through the screening function, as

$$\frac{p'R}{p} > \text{ or } < \frac{n-1}{n \ln(S/R)}.$$

3.4 Bureaucratic delay and organizational form

There is a commonly held view that democracies — and polyarchical organizations — are fine if you can afford them; but when things must get done quickly, as in a time of war or in a country engaged in rapid development, more centralized control is called for. It is certainly the case that most countries, in times of war, do resort to more centralization of control, but whether this is because centralization is more efficient in making quick decisions, or because there is a mistaken belief in the efficacy of centralized control(24) — there is perhaps a false sense of security in such times from knowing that someone is in control — has remained a question about which political scientists and historians have commented but which has received scant theoretical attention from economists.

The proclivity to resort to centralized control seems particularly odd, given the widespread perception, at other times, that government bureaucracies are subject to delays and inefficiencies. Why should an organizational form which is so abused in normal times suddenly change its stripes?

I want to argue that in fact there may be some theoretical justification for the standard perceptions. First, however, we need to observe that time, that scarcest of commodities, is frequently not well allocated in organizations.

This should not be surprising: The kinds of prices that are used to allocate scarce resources simply do not exist for the allocation of time within an organization. I once proposed to my department chairman that each chairman of a committee or department within the University be given a time budget, as well as a dollar budget. He could call a meeting,

but he would be charged for the time of the faculty members. I proposed, too, that a clock be placed on the wall, with the units being dollars, rather than minutes, so that each faculty member should weigh whether, for instance, his five minute speech was worth the \$ 1,000 of aggregate faculty time that it used. Needless to say, my suggestions were not taken up.

Hannaway (1989), in her recent study of the behaviour of managers of administrative units of organizations, has emphasized the disparity between organizational objectives and individual objectives, and the difficulty of distinguishing between them. When one manager consults another, it may be to obtain better information or to check on the correctness of his calculations, actions which improve the decision. But it may also be to show the other how well informed he is, or to shift some of the risk — should the decision turn out to be faulty — onto others, actions the return to which may be more private than organizational. Since others' time is an unpriced resource, there is a natural tendency for an excessive amount of consultation. In her studies, the amount of time spent by managers in meetings was enormous — on average almost two-thirds of the time(25).

We argued earlier that some degree of consultation, of collective decision-making, was an inevitable consequence of the limitations on any individual's ability to collect and process information. For most difficult decisions, two heads are better — but costlier — than one. The question is, how are the heads to be arranged?

In our earlier work, focusing on polyarchical versus hierarchical decision-making, we made only limited reference to the issue of the speed of decision-making. Two observations were made: 1. sequential decision-making, while saving on the costs of decision makers, involved delays, and 2. hierarchical decision-making was more vulnerable to delays from "faulty parts", that is, from a weak link in the chain of command; for instance, in a hierarchical structure, if one member of the hierarchy has an in-basket from which papers did not flow out, or came out only with long delays, the organization's performance could be seriously impaired.

Here, I want to emphasize the costs and benefits of coordination in decision-making. We consider a set of projects requiring two pieces of information, which, because of diseconomies of scope, must be collected by two separate individuals (suborganizations). The information may take either one, two, or three periods to collect. At time 0, the individual assigned to collect the information finds out how long it will take to collect his information. The organization cannot tell how long it should have taken to collect the information. Assume that if information is *not* collected immediately prior to being used, it becomes obsolete, and an additional expenditure is required to update it. Accordingly, if each believes that the other will complete the information acquisition in period 3, he will time his completion for period 3(26).

Consider now the gain from having an information coordinator-monitor, who, at time 0, determines how long it should take for the acquisition of each of the two pieces of informa-

Incentives, Information, and Organizational Design

tion. If γ_i is the fraction of the time that it takes i periods to gather the information, there is a probability γ_1^2 that the information will be processed at the end of period 1 and a probability of $(1 - \gamma_3^2 - \gamma_1^2)$ that it will be processed at the end of period 2. Let D_1 represent the value of having a decision one period earlier, D_2 the value of having a decision two periods earlier. Then the total (gross) value of the information coordination is represented by

$$V = D_1 (1 - \gamma_3^2 - \gamma_1^2) + D_2 \gamma_1^2.$$

We will now consider what happens as uncertainty about the time it takes to make a decision becomes greater, i. e., we consider a mean preserving spread in the distributions of γ . Hence, we assume that

$$k = \gamma_1 + 2\gamma_2 + 3\gamma_3 = 2 - \gamma_1 + \gamma_3 = \text{mean time}.$$

Substituting, we obtain

$$V = D_1 [1 - (k - 2)^2 - 2\gamma_1^2 - 2(k - 2)\gamma_1] + D_2 \gamma_1^2.$$

If the distribution is symmetric, $k=2$. It immediately follows that

1. *If the distribution of decision times is not too asymmetric, then a mean preserving increase in the uncertainty associated with information times increases the value of information coordination if $D_2 \geq 2 D_1$; and*
2. *An increase in the value of having information earlier increases the value of information coordination.*

In war time situations, the economy faces a whole range of distinctly new decision problems. The range of times required to make decisions is increased. At the same time, the value of quick decision-making is increased enormously. Our theorem provides a rationale for greater centralized control of decision-making in such situations.

4. Concluding remarks

The Walrasian paradigm of the market economy — and the Arrow-Debreu model to which it has given rise — while it has provided fundamental insights into the role of the price system in coordinating *certain* economic decisions, is incomplete and misleading, both with respect to its view of the nature of competition and decentralization, the efficiency of the economy, and the importance of prices. Much of economic activity takes place within organizations, in which limited use is made of the price system. We have seen how competition — contests — takes on a much richer dimension than reflected by the price-taking

paradigm. Firms compete for new innovations, and they race to enter markets. Individuals compete for positions within organizations. We have also seen how there is much more to decentralization than its role as a computing algorithm, as suggested by the conventional paradigm. Finally, while the traditional model exaggerates the virtues of the market — as *Greenwald — Stiglitz* (1986) have recently shown, whenever markets are incomplete and information is imperfect, market allocations are almost never constrained Pareto efficient — it also understates its virtues: Its ability to solve, if necessarily imperfectly, the problems of selection, incentives, and information gathering and aggregation which are the core problems in organizational design.

Few questions touch us as much in our everyday life as those concerning organizational design, or what we have called, organizational architecture. The sense of control over our own lives, our ability to fulfill ourselves as individuals, our sense of individuality may all depend, in large measure, on the extent of centralization or decentralization in our society. There is undoubtedly a relationship between economic decentralization and political democracy and freedoms. I have not touched upon these issues, not because I believe they are unimportant. But before we can address these broader issues, we must more fully understand the *economic* strengths and weakness of decentralized market economies, looked at not through the rose tinted glasses of the Fundamental Theorems of Welfare Economics, but through the kind of more balanced perspective that I have tried to present to you today(27).

5. References

Arrow, K. J., "Economic Welfare and the Allocation of Resources for Invention", in Nelson, R. (Ed.), *The Rate and Direction of Inventive Activity: Economic and Social Factors*, NBER, Princeton University Press, 1962.

Barzel, Y., "Optimal Timing of Innovations", *Review of Economics and Statistics*, 1968, 35(3), pp. 348-355.

Bolton, P., Farrell, J., "Decentralization, Duplication and Delay", Stanford University, IMSSS Summer Seminar Paper, 1988.

Dasgupta, P., Stiglitz, J. E., "Uncertainty, Market Structure and the Speed of R&D", *Bell Journal of Economics*, 1980A, 11(1), pp. 1-28.

Dasgupta, P., Stiglitz, J. E., "Industrial Structure and the Nature of Innovative Activity", *Economic Journal*, 1980B, 90, pp. 266-293.

Dasgupta, P., Stiglitz, J. E., "Potential Competition, Actual Competition and Economic Welfare", *European Economic Review*, 1988, 32(2/3), pp. 569-577.

Incentives, Information, and Organizational Design

Gilbert, R., Newbery, D., "Pre-emptive Patenting and the Persistence of Monopoly", *American Economic Review*, 1982, 72, pp. 514-526.

Green, J., Stokey, N., "A Comparison of Tournaments and Contests", *Journal of Political Economy*, 1983, 9, pp. 349-365.

Greenwald, B., Stiglitz, J. E., "Information, Finance Constraints and Business Fluctuations", Chung-Hua Institute, Taipei, Proceedings of the Taiwan Conference on Monetary Theory, 1986.

Grossman, S., Stiglitz, J. E., "Information and Competitive Price Systems", *American Economic Review*, 1976, 66(2), pp. 246-253.

Grossman, S., Stiglitz, J. E., "On Value Maximization and Alternative Objectives of the Firm", *Journal of Finance*, 1977, 32(2), pp. 389-402.

Grossman, S., Stiglitz, J. E., "Stockholder Unanimity in the Making of Production and Financial Decisions", *Quarterly Journal of Economics*, 1980A, 94(3), pp. 543-566.

Grossman, S., Stiglitz, J. E., "On the Impossibility of Informationally Efficient Markets", *American Economic Review*, 1980B, 70(3), pp. 393-408.

Hannaway, J., *Managing Managers*, Oxford University Press, 1989.

Holmstrom, B., "Moral Hazard in Teams", *Bell Journal of Economics*, 1982, 13(2), pp. 324-340.

Lazear, E., Rosen, S., "Rank Order Tournaments as Optimum Labor Contracts", *Journal of Political Economy*, 1981, 89, pp. 841-864.

Nalebuff, B., Stiglitz, J. E., "Prizes and Incentives: Towards a General Theory of Compensation and Competition", *The Bell Journal of Economics*, 1983A, 14(1), pp. 21-43.

Nalebuff, B., Stiglitz, J. E., "Information, Competition and Markets", *American Economic Review, Papers and Proceedings*, 1983B, 73, pp. 278-283.

Ross, S. A., "The Economic Theory of Agency: The Principal's Problem", *American Economic Review, Papers and Proceedings*, 1973, 63(2), pp. 134-144.

Sah, R. K., Stiglitz, J. E., "Perpetuation, Self-Reproduction, and Evolution of Organizations: The Selection and Performance of Managers", Paper presented at the 1985 World Congress of the Econometric Society, Cambridge, Mass., 1985A.

Sah, R. K., Stiglitz, J. E., "Human Fallibility and Economic Organization", *American Economic Review*, 1985B, 75, pp. 292-307.

Sah, R. K., Stiglitz, J. E., "The Architecture of Economic Systems: Hierarchies and Polyarchies", *American Economic Review*, 1986, 76, pp. 716-727.

Sah, R. K., Stiglitz, J. E., "Committees, Hierarchies, and Polyarchies", *Economic Journal*, 1988A, 98, pp. 451-470.

Sah, R. K., Stiglitz, J. E., "Qualitative Properties of Profit-Making k-out-of-n Systems Subject to Two Kinds of Failures", *IEEE Transactions on Reliability*, 1988B, 37(5), pp. 515-520.

Sah, R. K., Stiglitz, J. E., "Sources of Technological Divergence between Developed and Less Developed Countries", in Findlay, R., et al. (Eds.), *Debt, Stabilization and Development: Essays in Memory of Carlos Diaz-Alejandro*, Basil Blackwell, Oxford, 1989.

Salop, S. C., Scheffman, D. T., "Raising Rivals' Costs", *American Economic Review, Papers and Proceedings*, 1983, 73(2), pp. 267-271.

Stiglitz, J. E., "On the Optimality of the Stock Market Allocation of Investment", Paper presented at the Far Eastern Meetings of the Econometric Society, Tokyo, 1970.

Stiglitz, J. E., "Perfect and Imperfect Markets", Paper presented at the Winter Meeting of the Econometric Society, Yale University, 1971 (mimeo).

Stiglitz, J. E., "Incentives and Risk-Sharing in Sharecropping", *Review of Economics and Statistics*, 1974, 41, April, pp. 219-255.

Stiglitz, J. E., "Incentives, Risk and Information: Notes Towards a Theory of Hierarchy", *Bell Journal of Economics*, 1975, 6(2), pp. 552-579 (presented at the Berlin Symposium on Planning, August 1973).

Stiglitz, J. E., "Theory of Competition, Incentives and Risk", in Stiglitz, J. E., Mathewson, F. (Eds.), *Developments in the Theory of Market Structure*, Macmillan-M. I. T. Press, 1986, pp. 399-449.

Stiglitz, J. E., "On the Microeconomics of Technical Progress", Paper presented to IDP-Cepal Meetings, Buenos Aires, in Katz, J. M. (Ed.), *Technology Generation in Latin American Manufacturing Industries*, Macmillan, Hampshire, England, 1987A, pp. 56-77.

Stiglitz, J. E., "Human Nature and Economic Organization", Jacob Marschak Lecture, presented to the Far Eastern Meetings of the Econometric Society, Tokyo, 1987B.

Stiglitz, J. E., "Technological Change, Sunk Costs, and Competition", in Baily, M. N., Winston, C. (Eds.), *Special Issue on Macroeconomics*, The Brookings Institution, 1988, pp. 883-947.

6. Notes

(1) Reports on some of our results are presented in *Sah — Stiglitz* (1985A, 1985B, 1986, 1988A, 1988B) and *Stiglitz* (1987B).

(2) In later sections of the paper, I shall touch on some of the problems that have been encountered in implementing market socialism, and I shall be more precise in what ways the Arrow-Debreu model does not provide a good description of how market economies function.

Incentives, Information, and Organizational Design

(3) In the Arrow-Debreu model, firms make *choices* under uncertainty, but those choices are no different from conventional static choices: They simply maximize the value of the firm, using state contingent prices. They do not have to form judgments concerning the likelihood of various contingencies, and they are precluded, by assumption, from gathering information. There is no exploration of new technological possibilities, and no new products, and hence no scope for judgment concerning how the market would react to these new possibilities.

A corollary of these characteristics is that there is always unanimity about what the firm should do. In practice, of course, there are often disagreements among managers about the firm's business strategy (see *Stiglitz*, 1970, and *Grossman — Stiglitz*, 1977, 1980A).

(4) The new analytic results in this paper are contained in this section.

(5) Though not our formal models. It was not until the recent development of information economics that formal models showing how new information affected prices, and how prices could serve to transmit or aggregate information, were developed (see *Grossman — Stiglitz*, 1976, 1980B). These models have quite different welfare properties from the standard Arrow-Debreu model.

(6) An example of the complexity of the product space was recently provided by the U. S. Defense Department, when it put up for bidding a standard, white T-shirt, the kind of commodity that can be purchased in any clothing store for a few dollars. The specifications were 30 small print pages.

(7) For a more extensive discussion of these issues, see *Stiglitz* (1987A).

(8) Indeed, a central part of technological change is the acquisition of information (knowledge), and accordingly, it is not surprising that all the problems that have been encountered in extending the standard competitive paradigm to situations where information is imperfect arise when considering technological change.

(9) Not the stylized decentralization of the Arrow-Debreu model, but the kind of decentralization that actually characterizes market economies; see the preceding section.

(10) These arguments suggest the possibility of multiple equilibria: Decentralized structures create an environment which contributes to their own relative advantage, and conversely with bureaucratic structures. For an elaboration of this idea, see *Sah — Stiglitz* (1989).

(11) The work of Tversky and his co-authors has been particularly insightful in this respect.

(12) The noisiness of communication is often illustrated by the children's game, in which people sit in a circle, and a message is passed quickly around. Usually, the version which returns to the originator of the message is highly garbled.

(13) Thus, the essential difference between a committee of two, in which unanimity is required to proceed with the project, and a hierarchy, is one of delay and wasted effort. The advantage of the hierarchy is that projects that have been rejected by one person are not reviewed by the other — and since both would have to agree to have the project adopted, such a review would have been wasted effort. On the other hand, in a committee structure (as we model it) the two reviews occur at the same time, while in the hierarchy, the second review does not begin until after the first is completed.

(14) For instance, we show that 1. it is always desirable to break up a very long hierarchical chain into two (or more) hierarchies, with the hierarchies being sub-units of a polyarchy; 2. it is always desirable to reorganize a very wide polyarchy (that is, one in which there are many independent decision units), into two sub-units, which bear a hierarchical relationship with one another; and 3. with the appropriate combination of hierarchies and polyarchies, one can attain perfect discrimination between good and bad projects (see *Sah — Stiglitz*, 1986).

(15) See, for instance, *Ross* (1973) or *Stiglitz* (1974).

(16) In on-going work, we have viewed the evolutionary process from this perspective. Two kinds of errors arise: Good organizations get eliminated, and bad organizations survive. We analyse the determinants of the relative frequency and costs of these two kinds of errors (see *Sah — Stiglitz*, 1985A).

(17) Contributors to this recent literature include *Lazear — Rosen* (1981), *Nalebuff — Stiglitz* (1983A, 1983B), *Green — Stokey* (1983), *Holmstrom* (1982), and *Stiglitz* (1986).

(18) See also *Gilbert — Newbery* (1982), who extended the analysis to the case of uncertainty.

Stiglitz (1986) went on to show several other problems with Arrow's original argument: While it is true that if all industries except one are competitive, then production in the one remaining industry will be less than it otherwise would be, it is not necessarily true that if all industries are imperfectly competitive, production in each is less than in the competitive equilibrium. Furthermore, the relevant second-best comparisons must take into account the actual level of production (given that production is restricted, the question is, is R&D too slow in monopolistic economies?) and alternative ways of raising revenues (benefit taxes used to finance R&D would, for instance, have much the same effect on output as monopoly). Finally, as *Dasgupta — Stiglitz* (1980A, 1980B) emphasized, market structure — the extent of competition — should itself be viewed as endogenous, depending, among other things, on the technology of technological change.

(19) There are two other arguments suggesting that R&D may be too slow: Usually there are inappropriate benefits from R&D (the results suggest ideas to others, ideas which have market value but cannot be patented). Indeed, in some cases, imitators have an advantage over innovators. There is a "second mover advantage". Not only does the imitator save on the initial R&D expenditures, and avoids the many unsuccessful projects, but he can also learn from the innovator. He can see, for instance, market reaction to the various attributes of the product introduced by the innovator, and design a better product. Though the innovator could, in principle, redesign his product, it is often more difficult to change a reputation, a market image, than to create a new one.

(20) Partly due to the difficulties of getting the participants to reveal truthfully their investment plans, partly due to the problems we noted earlier arising from the complexity of the commodity space.

(21) *Bolton — Farrell* (1988) have recently formulated a model in which the date of entry — the delay from the first possible time for entry — serves as a perfect screening device; the lowest-cost firm and only the lowest-cost firm enters. The cost of this increased accuracy of screening (and reduced duplication) is increased delay.

(22) The expected loss of the bureaucratic selection mechanism (relative to first-best, perfect information) is approximately $0.25 [(C_H - C_L) Q_H + DWL (C_H - C_L)]$, where Q_H is the output with costs equal to C_H and $DWL (C_H - C_L)$ is the dead-weight loss (the Harberger triangle) arising from the curtailment of production associated with using the high-cost technology.

The expected loss from the market mechanism can be decomposed into three parts. When there is one low-cost firm and one high-cost firm (which occurs with probability 0.5), expected losses are $(1-s) S$, where $s \equiv S/R$. When there are two low-cost firms (which occurs with probability 0.25), the loss is S . And when both firms are high-cost, there are three kinds of losses: With probability s^4 no one enters, and in that case the (present discounted-value) loss is $Z (1 + \delta)$, where Z is the total surplus from producing (when $C = C_H$) the first period, and δZ is the surplus from producing the second; with probability $s^2 (1-s^2)$, entry is delayed, so the loss is just Z ; and with probability $(1-s)^2$ both firms enter the first period, so the loss is S and with probability $s^2 (1-s)^2$ both firms enter the second period, and so the (present discounted value of the) loss is δS .

Incentives, Information, and Organizational Design

(23) This is, of course, only one of several reasons. For instance, the prices and production norms that have been established by market processes prior to the emergency become increasingly out of tune with current situations, and systems of direct control often have difficulty adjusting them.

(24) In this perspective, an explanation for the proclivity for centralized control in times of emergency is more likely to be found in Freudian psychology than in economic analysis.

(25) Her results are consistent with the findings of other scholars.

(26) Presumably, the organization could call for a meeting at time t , and penalize anyone who had not completed his assigned task. But this would impose a risk on the manager — since, if the penalty is large enough to affect his behaviour, it will mean that he will be punished only when it was, in fact, impossible for him to acquire the information.

(27) Certainly, the perspective provided by the Lange-Lerner-Taylor theorem, suggesting an equivalence between socialist and market economies, provides a dramatic example of how traditional approaches can lead one far astray.

Correspondence:

Joseph E. Stiglitz
Department of Economics
Stanford University
Stanford, Cal., 94305-6072
U. S. A.