

RATIONAL DECISION-MAKING IN BUSINESS ORGANIZATIONS

Nobel Memorial Lecture, 8 December, 1977

by

HERBERT A. SIMON

Carnegie-Mellon University*, Pittsburgh, Pennsylvania, USA

In the opening words of his *Principles*, Alfred Marshall proclaimed economics to be a psychological science.

Political Economy or Economics is a study of mankind in the ordinary business of life; it examines that part of individual and social action which is most closely connected with the attainment and with the use of the material requisites of wellbeing.

Thus it is on the one side a study of wealth; and on the other, and more important side, a part of the study of man. For man's character has been moulded by his every-day work, and the material resources which he thereby procures, more than by any other influence unless it be that of his religious ideals.

In its actual development, however, economic science has focused on just one aspect of Man's character, his reason, and particularly on the application of that reason to problems of allocation in the face of scarcity. Still, modern definitions of the economic sciences, whether phrased in terms of allocating scarce resources or in terms of rational decision making, mark out a vast domain for conquest and settlement. In recent years there has been considerable exploration by economists even of parts of this domain that were thought traditionally to belong to the disciplines of political science, sociology, and psychology.

DECISION THEORY AS ECONOMIC SCIENCE

The density of settlement of economists over the whole empire of economic science is very uneven, with a few areas of modest size holding the bulk of the population. The economic Heartland is the normative study of the international and national economies and their markets, with its triple main concerns of full employment of resources, the efficient allocation of resources, and equity in distribution of the economic product. Instead of the ambiguous and over-general term "economics," I will use "political economy" to designate this Heartland, and "economic sciences" to denote the whole empire, including its most remote colonies. Our principal concern in this paper will be with the important colonial territory known as decision theory. I will have something to say about both its normative and descriptive aspects, and particularly about its applications to the theory of

* I am indebted to Albert Ando, Otto A. Davis, and Benjamin M Friedman for valuable comments on an earlier draft of this Paper.

the firm. It is through the latter topic that the discussion will be linked back to the Heartland of political economy.

Underpinning the corpus of policy-oriented normative economics, there is, of course, an impressive body of descriptive or "positive" theory which rivals in its mathematical beauty and elegance some of the finest theories in the physical sciences. As examples I need only remind you of Walrasian general equilibrium theories and their modern descendants in the works of Henry Schultz, Samuelson, Hicks, and others; or the subtle and impressive body of theory created by Arrow, Hurwicz, Debreu, Malinvaud, and their colleagues showing the equivalence, under certain conditions, of competitive equilibrium with Pareto optimality.

The relevance of some of the more refined parts of this work to the real world can be, and has been, questioned. Perhaps some of these intellectual mountains have been climbed simply because they were there - because of the sheer challenge and joy of scaling them. That is as it should be in any human scientific or artistic effort. But regardless of the motives of the climbers, regardless of real-world veridicality, there is no question but that positive political economy has been strongly shaped by the demands of economic policy for advice on basic public issues.

This too is as it should be. It is a vulgar fallacy to suppose that scientific inquiry cannot be fundamental if it threatens to become useful, or if it arises in response to problems posed by the everyday world. The real world, in fact, is perhaps the most fertile of all sources of good research questions calling for basic scientific inquiry.

Decision Theory in the Service of Political Economy

There is, however, a converse fallacy that deserves equal condemnation: the fallacy of supposing that fundamental inquiry is worth pursuing only if its relevance to questions of policy is immediate and obvious. In the contemporary world, this fallacy is perhaps not widely accepted, at least as far as the natural sciences are concerned. We have now lived through three centuries or more of vigorous and highly successful inquiry into the laws of nature. Much of that inquiry has been driven by the simple urge to understand, to find the beauty of order hidden in complexity. Time and again, we have found the "idle" truths arrived at through the process of inquiry to be of the greatest moment for practical human affairs. I need not take time here to argue the point. Scientists know it, engineers and physicians know it, congressmen and members of parliaments know it, the man on the street knows it.

But I am not sure that this truth is as widely known in economics as it ought to be. I cannot otherwise explain the rather weak and backward development of the descriptive theory of decision making including the theory of the firm, the sparse and scattered settlement of its terrain, and the fact that many if not most of its investigators are drawn from outside economics - from sociology, from psychology, and from political science. Respected and distinguished figures in economics - Edward Mason, Fritz

Machlup, and Milton Friedman, for example - have placed it outside the Pale (more accurately, have placed economics outside *its* Pale), and have offered it full autonomy provided that it did not claim close kinship with genuine economic inquiry.

Thus, Edward Mason, commenting on Papendreu's (1952) survey of research on the behavioral theory of the firm, mused aloud:

...has the contribution of this literature to economic analysis really been a large one? The writer of this critique must confess a lack of confidence in the marked superiority, *for purposes of economic analysis*, of this newer concept of the firm, over the older conception of the entrepreneur. (Mason, 1952, pp. 22 1-2, italics in the original.)

And, in a similar vein, Milton Friedman sums up his celebrated polemic against realism in theory (1953, p. 41, italics supplied):

Complete "realism" is clearly unattainable, and the question whether a theory is realistic "enough" can be settled only by seeing whether it yields predictions that are good enough *for the purpose in hand* or that are better than predictions from alternative theories.

The "purpose in hand" that is implicit in both of these quotations is providing decision-theoretic foundations for positive, and then for normative political economy. In the view of Mason and Friedman, fundamental inquiry into rational human behavior in the context of business organizations is simply not (by definition) economics - that is to say, political economy - unless it contributes in a major way to that purpose. This is sometimes even interpreted to mean that economic theories of decision making are not falsified in any interesting or relevant sense when their empirical predictions of *micropheomena* are found to be grossly incompatible with the observed data. Such theories, we are told, are still realistic "enough" provided that they do not contradict aggregate observations of concern to political economy. Thus economists who are zealous in insisting that economic actors maximize turn around and become satisficers when the evaluation of their own theories is concerned. They believe that businessmen maximize, but they know that economic theorists satisfice.

The application of the principle of satisficing to theories is sometimes defended as an application of Occam's Razor: accept the simplest theory that works.² But Occam's Razor has a double edge. Succinctness of statement is not the only measure of a theory's simplicity. Occam understood his rule as recommending theories that make no more assumptions than

²The phrase "that works" refutes, out of hand, Friedman's celebrated paean of praise for lack of realism in assumptions. Consider his example of the law of falling bodies (Friedman, pp. 16-19). His valid point is that it is advantageous to use the simple law, ignoring air resistance, when it gives a "good enough" approximation. But of course the conditions under which it gives a good approximation are not at all the conditions under which it is unrealistic or a "widely inaccurate descriptive representation of reality." I cannot in this brief space mention, much less discuss, all of the numerous logical fallacies that can be found in Friedman's 40-page essay. For additional criticism, see Simon (1962) and Samuelson (1962).

necessary to account for the phenomena (*Essentia non sunt multiplicanda praeter necessitatem*). A theory of profit or utility maximization can be stated more briefly than a satisfying theory of the sort I shall discuss later. But the former makes much stronger assumptions than the latter about the human cognitive system. Hence in the case before us, the two edges of the razor cut in opposite directions.

In whichever way we interpret Occam's principle, parsimony can be only a secondary consideration in choosing between theories, unless those theories make identical predictions. Hence, we must come back to a consideration of the phenomena that positive decision theory is supposed to handle. These may include both phenomena at the microscopic level of the decision-making agents, or aggregative phenomena of concern to political economy.

Decision Theory Pursued for its Intrinsic Interest

Of course the definition of the word "economics" is not important. Like Humpty Dumpty, we can make words mean anything we want them to mean. But the professional training and range of concern of economists does have importance. Acceptance of the narrow view that economics is concerned only with the aggregative phenomena of political economy defines away a whole rich domain of rational human behavior as inappropriate for economic research.

I do not wish to appear to be admitting that the behavioral theory of the firm *has been* irrelevant to the construction of political economy. I will have more to say about its relevance in a moment. My present argument is counterfactual in form: *even if* there were no present evidence of such relevance, human behavior in business firms constitutes a highly interesting body of empirical phenomena that calls out for explanation as do all bodies of phenomena. And if we may extrapolate from the history of the other sciences, there is every reason to expect that as explanations emerge, relevance for important areas of practical application will not be long delayed.

It has sometimes been implied (Friedman, p. 14) that the correctness of the assumptions of rational behavior underlying the classical theory of the firm is not merely irrelevant, but is not even empirically testable in any direct way, the only valid test being whether these assumptions lead to tolerably correct predictions at the macroscopic level. That would be true, of course, if we had no microscopes, so that the micro-level behavior was not directly observable. But we do have microscopes. There are many techniques for observing decision-making behavior, even at second-by-second intervals if that is wanted. In testing our economic theories, we do not have to depend on the rough aggregate times series that are the main grist for the econometric mill, or even upon company financial statements.

The classical theories of economic decision making and of the business firm make very specific testable predictions about the concrete behavior of decision-making agents. Behavioral theories make quite different predic-

tions. Since these predictions can be tested directly by observation, either theory (or both) may be falsified as readily when such predictions fail as when predictions about aggregate phenomena are in error.

Aggregative Tests of Decision Theory: Marginalism

If some economists have erroneously supposed that microeconomic theory can only be tested by its predictions of aggregate phenomena we should avoid the converse error of supposing that aggregate phenomena are irrelevant to testing decision theory. In particular, are there important, *empirically verified* aggregate predictions that follow from the theory of perfect rationality but that do not follow from behavioral theories of rationality?

The classical theory of omniscient rationality is strikingly simple and beautiful. Moreover, it allows us to predict (correctly or not) human behavior without stirring out of our armchairs to observe what such behavior is like. All the predictive power comes from characterizing the shape of the environment in which the behavior takes place. The environment, combined with the assumptions of perfect rationality, fully determines the behavior. Behavioral theories of rational choice - theories of bounded rationality - do not have this kind of simplicity. But, by way of compensation, their assumptions about human capabilities are far weaker than those of the classical theory. Thus, they make modest and realistic demands on the knowledge and computational abilities of the human agents, but they also fail to predict that those agents will equate costs and returns at the margin.

Have the Marginalist Predictions Been Tested?

A number of empirical phenomena have been cited as providing more or less conclusive support for the classical theory of the firm as against its behavioral competitors (Jorgensen and Siebert, 1968). But there are no direct observations that individuals or firms do actually equate marginal costs and revenues. The empirically verified consequences of the classical theory are always weaker than this. Let us look at four of the most important of them: the fact that demand curves generally have negative slopes, the fact that fitted Cobb-Douglas functions are approximately homogeneous of the first degree, the fact of decreasing returns to scale, and the fact that executive salaries vary with the logarithm of company size. Are these indeed facts? And does the evidence support a maximizing theory against a satisficing theory?

Negatively Sloping Demand Curves. Evidence that consumers actually distribute their purchases in such a way as to maximize their utilities, and hence to equate marginal utilities, is nonexistent. What the empirical data do confirm is that demand curves generally have negative slopes. (Even this "obvious" fact is tricky to verify, as Henry Schultz (1938) showed long years ago.) But negatively sloping demand curves could result from a wide range of behaviors satisfying the assumptions of bounded rationality rath-

er than those of utility maximization. Gary Becker, who can scarcely be regarded as a hostile witness for the classical theory, states the case very well (1962, p. 4):

Economists have long been aware that some changes in the feasible or opportunity sets of households would lead to the same response *regardless of the decision rule used*. For example, a decrease in real income necessarily decreases the amount spent on at least one commodity... It has seldom been realized, however, that the change in opportunities resulting from a change in relative prices also tends to produce a systematic response, regardless of the decision rule. In particular, the fundamental theorem of traditional theory - that demand curves are negatively inclined - largely results from the change in opportunities alone and is largely independent of the decision rule.

Later, Becker is even more explicit, saying (1962, p. 5), "Not only utility maximization but also many other decision rules, incorporating a wide variety of irrational behavior, lead to negatively inclined demand curves because of the effect of a change in prices on opportunities."³

First-Degree Homogeneity of Production Functions. Another example of an observed phenomenon for which the classical assumptions provide sufficient, but not necessary, conditions is the equality between labor's share of product and the exponent of the labor factor in fitted Cobb-Douglas production functions (Simon & Levy, 1963). Fitted Cobb-Douglas functions are homogeneous, generally of degree close to unity and with a labor exponent of about the right magnitude. These findings, however, cannot be taken as strong evidence for the classical theory, for the identical results can readily be produced by mistakenly fitting a Cobb-Douglas function to data that were in fact generated by a linear accounting identity (value of goods equals labor cost plus capital cost) (Phelps-Brown, 1957). The same comment applies to the SMAC production function (Cyert & Simon, 1971). Hence, the empirical findings do not allow us to draw any particular conclusions about the relative plausibility of classical and behavioral theories, both of which are equally compatible with the data.

The Long-Run Cost Curve. Somewhat different is the case of the firm's long-run cost curve, which classical theory requires to be U-shaped if competitive equilibrium is to be stable. Theories of bounded rationality do not predict this - fortunately, for the observed data make it exceedingly doubtful that the cost curves are in fact generally U-shaped. The evidence for many industries shows costs at the high-scale ends of the curves to be essentially constant or even declining (Walters, 1963). This finding is compatible with stochastic models of business firm growth and size (Ijiri & Simon, 1977), but not with the static equilibrium model of classical theory.

Executive Salaries. Average salaries of top corporate executives grow with

³In a footnote Becker indicates that he denotes as irrational "[A]ny deviation from utility maximization." Thus, what I have called "bounded rationality" is "irrationality" in Becker's terminology.

the logarithm of corporate size (Roberts, 1959). This finding has been derived from the assumptions of the classical theory of profit maximization only with the help of very particular *ad hoc* assumptions about the distribution of managerial ability (Lucas, 1978). The observed relation is implied by a simple behavioral theory that assumes only that there is a single, culturally determined, parameter which fixes the average ratio of the salaries of managers to the salaries of their immediate subordinates (Simon, 1957). In the case of the executive salary data, the behavioral model that explains the observations is substantially more parsimonious (in terms of assumptions about exogenous variables) than the classical model that explains the same observations.

Summary: Phenomena That Fail to Discriminate. It would take a much more extensive review than is provided here to establish the point conclusively, but I believe it is the case that specific phenomena requiring a theory of utility or profit maximization for their explanation rather than a theory of bounded rationality simply have not been observed in aggregate data. In fact, as my last two examples indicate, it is the classical, rather than the behavioral form of the theory that faces real difficulties in handling some of the empirical observations.

Failures of Classical Theory. It may well be that classical theory can be patched up sufficiently to handle a wide range of situations where uncertainty and outguessing phenomena do not play a central role - that is, to handle the behavior of economies that are relatively stable and not too distant from a competitive equilibrium. However, a strong positive case for replacing the classical theory by a model of bounded rationality begins to emerge when we examine situations involving decision making under uncertainty and imperfect competition. These situations the classical theory was never designed to handle, and has never handled satisfactorily. Statistical decision theory employing the idea of subjective expected utility, on the one hand, and game theory, on the other, have contributed enormous conceptual clarification to these kinds of situations without providing satisfactory descriptions of actual human behavior, or even, for most cases, normative theories that are actually usable in the face of the limited computational powers of men and computers.

I shall have more to say later about the positive case for a descriptive theory of bounded rationality, but I would like to turn first to another territory within economic science that has gained rapidly in population since World War II, the domain of normative decision theory.

Normative Decision Theory

Decision theory can be pursued not only for the purposes of building foundations for political economy, or of understanding and explaining phenomena that are in themselves intrinsically interesting, but also for the purpose of offering direct advice to business and governmental decision makers. For reasons not clear to me, this territory was very sparsely settled prior to World War II. Such inhabitants as it had were mainly industrial

engineers, students of public administration, and specialists in business functions, none of whom especially identified themselves with the economic sciences. Prominent pioneers included the mathematician, Charles Babbage, inventor of the digital computer, the engineer, Frederick Taylor, and the administrator, Henri Fayol.

During World War II, this territory, almost abandoned, was rediscovered by scientists, mathematicians, and statisticians concerned with military management and logistics, and was renamed "operations research" or "operations analysis." So remote were the operations researchers from the social science community that economists wishing to enter the territory had to establish their own colony, which they called "management science." The two professional organizations thus engendered still retain their separate identities, though they are now amicably federated in a number of common endeavors.

Optimization techniques were transported into management science from economics, and new optimization techniques, notably linear programming, were invented and developed, the names of Dantzig, Kantorovich, and Koopmans being prominent in the early development of that tool.

Now the salient characteristic of the decision tools employed in management science is that they have to be capable of actually making or recommending decisions, taking as their inputs the kinds of empirical data that are available in the real world, and performing only such computations as can reasonably be performed by existing desk calculators or, a little later, electronic computers. For these domains, idealized models of optimizing entrepreneurs, equipped with complete certainty about the world - or, at worst, having full probability distributions for uncertain events - are of little use. Models have to be fashioned with an eye to practical computability, no matter how severe the approximations and simplifications that are thereby imposed on them.

Model construction under these stringent conditions has taken two directions. The first is to retain optimization, but to simplify sufficiently so that the optimum (in the simplified world!) is computable. The second is to construct satisficing models that provide good enough decisions with reasonable costs of computation. By giving up optimization, a richer set of properties of the real world can be retained in the models. Stated otherwise, decision makers can satisfice either by finding optimum solutions for a simplified world, or by finding satisfactory solutions for a more realistic world. Neither approach, in general, dominates the other, and both have continued to co-exist in the world of management science.

Thus, the body of theory that has developed in management science shares with the body of theory in descriptive decision theory a central concern with the ways in which decisions are made, and not just with the decision outcomes. As I have suggested elsewhere (1978b), these are theories of *how* to decide rather than theories of *what* to decide.

Let me cite one example, from work in which I participated, of how

model building in normative economics is shaped by computational considerations (Holt, Modigliani, Muth & Simon, 1960). In the face of uncertain and fluctuating production demands, a company can smooth and stabilize its production and employment levels at the cost of holding buffer inventories. What kind of decision rule will secure a reasonable balance of costs? Formally, we are faced with a dynamic programming problem, which generally pose formidable and often intolerable computational burdens for their solution.

One way out of this difficulty is to seek a special case of the problem that will be computationally tractable. If we assume the cost functions facing the company all to be quadratic in form, the optimal decision rule will then be a linear function of the decision variables, which can readily be computed in terms of the cost parameters. Equally important, under uncertainty about future sales, only the expected values, and not the higher moments, of the probability distributions enter into the decision rule. Hence the assumption of quadratic costs reduces the original problem to one that is readily solved. Of course the solution, though it provides optimal decisions for the simplified world of our assumptions, provides, at best, satisfactory solutions for the real-world decision problem that the quadratic function approximates. In-principle, unattainable optimization is sacrificed for in-practice, attainable satisfaction.

If human decision makers are as rational as their limited computational capabilities and their incomplete information permit them to be, then there will be a close relation between normative and descriptive decision theory. Both areas of inquiry are concerned primarily with procedural rather than substantive rationality (Simon, 1978a). As new mathematical tools for computing optimal and satisfactory decisions are discovered, and as computers become more and more powerful, the recommendations of normative decision theory will change. But as the new recommendations are diffused, the actual, observed, practice of decision making in business firms will change also. And these changes may have macroeconomic consequences. For example, there is some agreement that average inventory holdings of American firms have been reduced significantly by the introduction of formal procedures for calculating reorder points and quantities.

CHARACTERIZING BOUNDED RATIONALITY

The principal forerunner of a behavioral theory of the firm is the tradition usually called institutionalism. It is not clear that all of the writings, European and American, usually lumped under this rubric have much in common, or that their authors would agree with each other's views. At best, they share a conviction that economic theory must be reformulated to take account of the social and legal structures amidst which market transactions are carried out. Today, we even find a vigorous development within

economics that seeks to achieve institutionalist goals within the context of neoclassical price theory. I will have more to say about that a little later.

The name of John R. Commons is prominent - perhaps the most prominent - among American Institutionalists. Commons' difficult writings (e.g., Commons, 1934) borrow their language heavily from the law, and seek to use the *transaction* as their basic unit of behavior. I will not undertake to review Commons' ideas here, but simply remark that they provided me with many insights in my initial studies of organizational decision making (see *Administrative Behavior*, page 136).

Commons also had a substantial influence on the thinking of Chester I. Barnard, an intellectually curious business executive who distilled from his experience as President of the New Jersey Belt Telephone Company, and as executive of other business, governmental, and non-profit organizations, a profound book on decision making titled *The Functions of the Executive* (1938). Barnard proposed original theories, which have stood up well under empirical scrutiny, of the nature of the authority mechanism in organizations, and of the motivational bases for employee acceptance of organizational goals (the so-called "inducements-contributions" theory); and he provided a realistic description of organizational decision making, which he characterized as "opportunistic." The numerous references to Barnard's work in *Administrative Behavior* attest, though inadequately, to the impact he had on my own thinking about organizations.

In Search of a Descriptive Theory

In 1934-5, in the course of a field study of the administration of public recreational facilities in Milwaukee, which were managed jointly by the school board and the city public works department, I encountered a puzzling phenomenon. Although the heads of the two agencies appeared to agree as to the objectives of the recreation program, and did not appear to be competing for empire, yet there was continual disagreement and tension between them with respect to the allocation of funds between physical maintenance, on the one hand, and play supervision on the other. Why did they not, as my economics books suggested, simply balance off the marginal return of the one activity against that of the other?

Further exploration made it apparent that they didn't equate expenditures at the margin because, intellectually, they couldn't. There was no measurable production function from which quantitative inferences about marginal productivities could be drawn; and such qualitative notions of a production function as the two managers possessed were mutually incompatible. To the public works administrator, a playground was a physical facility, serving as a green oasis in the crowded gray city. To the recreation administrator, a playground was a social facility, where children could play together with adult help and guidance.

How can human beings make rational decisions in circumstances like these? How are they to apply the marginal calculus? Or, if it does not apply, what do they substitute for it?

The phenomenon observed in Milwaukee is ubiquitous in human decision making. In organization theory it is usually referred to as *subgoal identification*. When the goals of an organization cannot be connected operationally with actions (when the production function can't be formulated in concrete terms), then decisions will be judged against subordinate goals that can be so connected. There is no unique determination of these subordinate goals. Their formulation will depend on the knowledge, experience, and organizational environment of the decision maker. In the face of this ambiguity, the formulation can also be influenced in subtle, and not so subtle, ways by his self-interest and power drives.

The phenomenon arises as frequently in individual as in social decision making and problem solving. Today, under the rubric of *problem representation*, it is a central research interest of cognitive psychology. Given a particular environment of stimuli, and a particular background of previous knowledge, how will a person organize this complex mass of information into a problem formulation that will facilitate his solution efforts? How did Newton's experience of the apple, if he had one, get represented as an instance of attraction of apple by Earth?

Phenomena like these provided the central theme for *Administrative Behavior*. That study represented "an attempt to construct tools useful in my own research in the field of public administration." The product was actually not so much a theory as a prolegomena to a theory, stemming from the conviction "that decision making is the heart of administration, and that the vocabulary of administrative theory must be derived from the logic and psychology of human choice." It was, if you please, an exercise in problem representation.

On examination, the phenomenon of subgoal identification proved to be the visible tip of a very large iceberg. The shape of the iceberg is best appreciated by contrasting it with classical models of rational choice. The classical model calls for knowledge of all the alternatives that are open to choice. It calls for complete knowledge of, or ability to compute, the consequences that will follow on each of the alternatives. It calls for certainty in the decision-maker's present and future evaluation of these consequences. It calls for the ability to compare consequences, no matter how diverse and heterogeneous, in terms of some consistent measure of utility. The task, then was to replace the classical model with one that would describe how decisions could be (and probably actually were) made when the alternatives of search had to be sought out, the consequences of choosing particular alternatives were only very imperfectly known both because of limited computational power and because of uncertainty in the external world, and the decision maker did not possess a general and consistent utility function for comparing heterogeneous alternatives.

Several procedures of rather general applicability and wide use have been discovered that transform intractable decision problems into tractable ones. One procedure already mentioned is to look for satisfactory choices instead of optimal ones. Another is to replace abstract, global goals

with tangible subgoals, whose achievement can be observed and measured. A third is to divide up the decision-making task among many specialists, coordinating their work by means of a structure of communications and authority relations. All of these, and others, fit the general rubric of "bounded rationality", and it is now clear that the elaborate organizations that human beings have constructed in the modern world to carry out the work of production and government can only be understood as machinery for coping with the limits of man's abilities to comprehend and compute in the face of complexity and uncertainty.

This rather vague and general initial formulation of the idea of bounded rationality called for elaboration in two directions: greater formalization of the theory, and empirical verification of its main claims. During the decade that followed the publication of *Administrative Behavior*, substantial progress was made in both directions, some of it through the efforts of my colleagues and myself, much of it by other research groups that shared the same Zeitgeist.

Empirical Studies

The principal source of empirical data about organizational decision making has been straightforward "anthropological" field study, eliciting descriptions of decision making procedures and observing the course of specific decision-making episodes. Examples are a study, with Guetzkow, Kozmetsky, and Tyndall (1954), of the ways in which accounting data were used in decision-making in large corporations; and a series of studies, with Cyert, March, and others, of specific non-programmed policy decisions in a number of different companies (Cyert, Simon & Trow, 1956). The latter line of work was greatly developed and expanded by Cyert and March and its theoretical implications for economics explored in their important work, *The Behavioral Theory of the Firm* (1963).

At about the same time, the fortuitous availability of some data on businessmen's perceptions of a problem situation described in a business policy casebook enabled Dearborn and me (1958) to demonstrate empirically the cognitive basis for identification with subgoals, the phenomenon that had so impressed me in the Milwaukee recreation study. The businessmen's perceptions of the principal problems facing the company described in the case were mostly determined by their own business experiences - sales and accounting executives identified a sales problem, manufacturing executives, a problem of internal organization.

Of course there is vastly more to be learned and tested about organizational decision-making than can be dealt with in a handful of studies. Although many subsequent studies have been carried out in Europe and the United States, this domain is still grossly undercultivated (For references, see March, 1965, Johnsen, 1968, Eliasson, 1976). Among the reasons for the relative neglect of such studies, as contrasted, say, with laboratory experimentation in social psychology, is that they are extremely costly and time-consuming, with a high grist-to-grain ratio, the methodology for

carrying them out is primitive, and satisfactory access to decision-making behavior is hard to secure. This part of economics has not yet acquired the habits of patience and persistence in the pursuit of facts that is exemplified in other domains by the work, say, of Kuznetz or of the architects of the MIT-SSRC-Pennsylvania econometric models.

Theoretical Inquiries.

On the theoretical side, three questions seemed especially to call for clarification: what are the circumstances under which an employment relation will be preferred to some other form of contract as the arrangement for securing the performance of work; what is the relation between the classical theory of the firm and theories of organizational equilibrium first proposed by C. I. Barnard; and what are the main characteristics of human rational choice in situations where complexity precludes omniscience?

The Employment Relation. A fundamental characteristic of modern industrial society is that most work is performed, not by individuals who produce products for sale, nor by individual contractors, but by persons who have accepted employment in a business firm and the authority relation with the employer that employment entails. Acceptance of authority means willingness to permit one's behavior to be determined by the employer, at least within some zone of indifference or acceptance. What is the advantage of this arrangement over a contract for specified goods or services? Why is so much of the world's work performed in large, hierarchic organizations?

Analysis showed (Simon, 1951) that a combination of two factors could account for preference for the employment contract over other forms of contracts: uncertainty as to which future behaviors would be advantageous to the employer, and a greater indifference of the employee as compared with the employer (within the former's area of acceptance) as to which of these behaviors he carried out. When the secretary is hired, the employer does not know what letters he will want her to type, and the secretary has no great preference for typing one letter rather than another. The employment contract permits the choice to be postponed until the uncertainty is resolved, with little cost to the employee and great advantage to the employer. The explanation is closely analogous to one Marschak had proposed (1949) for liquidity preference. Under conditions of uncertainty it is advantageous to hold resources in liquid, flexible form.

Organizational Equilibrium. Barnard (1938) had described the survival of organizations in terms of the motivations that make their participants (employees, investors, customers, suppliers) willing to remain in the system. In *Administrative Behavior*, I had developed this notion further into a motivational theory of the balance between the inducements that were provided by organizations to their participants, and the contributions those participants made to the organizations' resources.

A formalization of this theory (Simon, 1952-53) showed its close affin-

ity to the classical theory of the firm, but with an important and instructive difference. In comparing the two theories, each inducement-contribution relation became a supply schedule for the firm. The survival conditions became the conditions for positive profit. But while the classical theory of the firm assumes that all profits go to a particular set of participants, the owners, the organization theory treats the surplus more symmetrically, and does not predict how it will be distributed. Hence the latter theory leaves room, under conditions of monopoly and imperfect competition, for bargaining among the participants (e.g., between labor and owners) for the surplus. The survival conditions - positive profits rather than maximum profits - also permits a departure from the assumptions of perfect rationality.

Mechanisms of Bounded Rationality. In *Administrative Behavior*, bounded rationality is largely characterized as a residual category - rationality is bounded when it falls short of omniscience. And the failures of omniscience are largely failures of knowing all the alternatives, uncertainty about relevant exogenous events, and inability to calculate consequences. There was needed a more positive and formal characterization of the mechanisms of choice under conditions of bounded rationality. Two papers (Simon, 1955 and 1956) undertook first steps in that direction.

Two concepts are central to the characterization: *search* and *satisficing*. If the alternatives for choice are not given initially to the decision maker, then he must search for them. Hence, a theory of bounded rationality must incorporate a theory of search. This idea was later developed independently by Stigler (1961) in a very influential paper that took as its example of a decision situation the purchase of a second-hand automobile. Stigler poured the search theory back into the old bottle of classical utility maximization, the cost of search being equated with its marginal return. In my 1956 paper, I had demonstrated the same formal equivalence, using as my example a dynamic programming formulation of the process of selling a house.

But utility maximization, as I showed, was not essential to the search scheme - fortunately, for it would have required the decision maker to be able to estimate the marginal costs and returns of search in a decision situation that was already too complex for the exercise of global rationality. As an alternative, one could postulate that the decision maker had formed some *aspiration* as to how good an alternative he should find. As soon as he discovered an alternative for choice meeting his level of aspiration, he would terminate the search and choose that alternative. I called this mode of selection *satisficing*. It had its roots in the empirically based psychological theories, due to Lewin and others, of aspiration levels. As psychological inquiry had shown, aspiration levels are not static, but tend to rise and fall in consonance with changing experiences. In a benign environment that provides many good alternatives, aspirations rise; in a harsher environment, they fall.

In long-run equilibrium it might even be the case that choice with

dynamically adapting aspiration levels would be equivalent to optimal choice, taking the costs of search into account. But the important thing about the search and satisficing theory is that it showed how choice could actually be made with reasonable amounts of calculation, and using very incomplete information, without the need of performing the impossible - of carrying out this optimizing procedure.

Summary

Thus, by the middle 1950's, a theory of bounded rationality had been proposed as an alternative to classical omniscient rationality, a significant number of empirical studies had been carried out that showed actual business decision making to conform reasonably well with the assumptions of bounded rationality but not with the assumptions of perfect rationality, and key components of the theory - the nature of the authority and employment relations, organizational equilibrium, and the mechanisms of search and satisficing - had been elucidated formally. In the remaining parts of this paper, I should like to trace subsequent developments of decision-making theory, including developments competitive with the theory of bounded rationality, and then to comment on the implications (and potential implications) of the new descriptive theory of decision for political economy.

THE NEOCLASSICAL REVIVAL

Peering forward from the late 1950's, it would not have been unreasonable to predict that theories of bounded rationality would soon find a large place in the mainstream of economic thought. Substantial progress had been made in providing the theories with some formal structure, and an increasing body of empirical evidence showed them to provide a far more veridical picture of decision making in business organizations than did the classical concepts of perfect rationality.

History has not followed any such simple course, even though many aspects of the *Zeitgeist* were favorable to movement in this direction. During and after World War II, a large number of academic economists were exposed directly to business life, and had more or less extensive opportunities to observe how decisions were actually made in business organizations. Moreover, those who became active in the development of the new management science were faced with the necessity of developing decision-making procedures that could actually be applied in practical situations. Surely these trends would be conducive to moving the basic assumptions of economic rationality in the direction of greater realism.

But these were not the only things that were happening in economics in the Postwar period. First, there was a vigorous reaction that sought to defend classical theory from behavioralism on methodological grounds. I have already commented on these methodological arguments in the first

part of my talk. However deeply one may disagree with them, they were stated persuasively and are still influential among academic economists.

Second, the rapid spread of mathematical knowledge and competence in the economics profession permitted the classical theory, especially when combined with statistical decision theory and the theory of games due to von Neuman and Morgenstern, to develop to new heights of sophistication and elegance, and to expand to embrace, albeit in highly stylized form, some of the phenomena of uncertainty and imperfect information. The flowering of mathematical economics and econometrics has provided two generations of economic theorists with a vast garden of formal and technical problems that have absorbed their energies and postponed encounters with the inelegancies of the real world.

If I sound mildly critical of these developments, I should confess that I have also been a part of them, admire them, and would be decidedly unhappy to return to the pre-mathematical world they have replaced. My concern is that the economics profession has exhibited some of the serial one-thing-at-a-time character of human rationality, and has seemed sometimes to be unable to distribute its attention in a balanced fashion among neoclassical theory, macroeconometrics, and descriptive decision theory. As a result, not as much professional effort has been devoted to the latter two, and especially the third, as one might have hoped and expected. The Heartland is more overpopulated than ever, while rich lands in other parts of the empire go untended.

Search and Information Transfer

Let me allude to just three of the ways in which classical theory has sought to cope with some of its traditional limitations, and has even sought to make the development of a behavioral theory, incorporating psychological assumptions, unnecessary. The first was to introduce search and information transfer explicitly as economic activities, with associated costs and outputs, that could be inserted into the classical production function. I have already referred to Stigler's 1961 paper on the economics of information, and my own venture in the same direction in the 1956 essay cited earlier.

In theory of this genre, the decision maker is still an individual. A very important new direction, where decisions are made by groups of individuals, in teams or organizations, is the economic theory of teams developed by Jacob Marschak and Roy Radner (1972). Here we see genuine organizational phenomena - specialization of decision making as a consequence of the costs of transmitting information - emerge from the rational calculus. Because the mathematical difficulties are formidable, the theory remains largely illustrative and limited to very simple situations in miniature organizations. Nevertheless, it has greatly broadened our understanding of the economics of information.

In none of these theories - any more than in statistical decision theory or the theory of games - is the assumption of perfect maximization aban-

done. Limits and costs of information are introduced, not as psychological characteristics of the decision maker, but as part of his technological environment. Hence, the new theories do nothing to alleviate the computational complexities facing the decision maker - do not see him coping with them by heroic approximation, simplifying and satisficing, but simply magnify and multiply them. Now he needs to compute not merely the shapes of his supply and demand curves, but the costs and benefits of computing those shapes to greater accuracy as well. Hence, to some extent, the impression that these new theories deal with the hitherto ignored phenomena of uncertainty and information transmission is illusory. For many economists, however, the illusion has been persuasive.

Rational Expectations Theory

A second development in neoclassical theory on which I wish to comment is the so-called "rational expectations" theory. There is a bit of historical irony surrounding its origins. I have already described the management science inquiry of Holt, Modigliani, Muth and myself that developed a dynamic programming algorithm for the special (and easily computed) case of quadratic cost functions. In this case, the decision rules are linear, and the probability distributions of future events can be replaced by their expected values, which serve as certainty equivalents (Simon, 1956; Theil, 1957).

John F. Muth imaginatively saw in this special case a paradigm for rational behavior under uncertainty. What to some of us in the HMMS research team was an approximating, satisficing simplification, served for him as a major line of defence for perfect rationality. He said in his seminal 1961 *Econometrica* article (p. 316), "It is sometimes argued that the assumption of rationality in economics leads to theories inconsistent with, or inadequate to explain, observed phenomena, especially changes over time . . . Our hypothesis is based on exactly the opposite point of view: that dynamic economic models do not assume enough rationality."

The new increment of rationality that Muth proposed was that (page 316) "expectations, since they are informed predictions of future events, are essentially the same as the predictions of the relevant economic theory." He would cut the Gordian knot. Instead of dealing with uncertainty by elaborating the model of the decision process, he would once and for all - if his hypothesis were correct - make process irrelevant. The subsequent vigorous development of rational expectations theory, in the hands of Sargent, Lucas, Prescott, and others, is well known to most readers (see, for example, Lucas, 1975).

It is too early to render a final verdict on the rational expectations theory. The issue will ultimately be decided, as all scientific debates should be, by a gradual winnowing of the empirical evidence, and that winnowing process has just begun. Meanwhile, certain grave theoretical difficulties have already been noticed. As Muth himself has pointed out, it is only rational (i.e., profit maximizing) to use the "rational expectations" decision

rule if the relevant cost equations are in fact quadratic. I have suggested elsewhere (Simon, 1978a) that it might therefore be less misleading to call the rule a "consistent expectations" rule.

Perhaps even more important, Ando (1979) and Benjamin Friedman (1978, 1979) have shown that the policy implications of the rational expectations rule are quite different under conditions where new information continually becomes available to the system, structural changes occur, and the decision-maker learns, than they are under steady-state conditions. For example, under the more dynamic conditions, monetary neutrality - which generally holds for the static consistent expectations models - is no longer guaranteed for any finite time horizon.

In the recent "revisionist" versions of consistent expectations theory, moreover, where account is taken of a changing environment of information, various behavioral assumptions reappear to explain how expectations are formed - what information decision makers will consider, and what they will ignore. But unless these assumptions are to be made on a wholly *ad hoc* and arbitrary basis, they create again the need for an explicit and valid theory of the decision-making *process* (Simon, 1958a, B. Friedman, 1979).

Statistical Decision Theory and Game Theory

Statistical decision theory and game theory are two other important components of the neoclassical revival. The former addresses itself to the question of incorporating uncertainty (or more properly, risk) into the decision-making models. It requires heroic assumptions about the information the decision maker has concerning the probability distributions of the relevant variables, and simply increases by orders of magnitude the computational problems he faces.

Game theory addresses itself to the "outguessing" problem that arises whenever an economic actor takes into account the possible reactions to his own decisions of the other actors. To my mind, the main product of the very elegant apparatus of game theory has been to demonstrate quite clearly that it is virtually impossible to define an unambiguous criterion of rationality for this class of situations (or, what amounts to the same thing, a definitive definition of the "solution" of a game). Hence, game theory has not brought to the theories of oligopoly and imperfect competition the relief from their contradictions and complexities that was originally hoped for it. Rather, it has shown that these difficulties are ineradicable. We may be able to reach consensus that a certain criterion of rationality is appropriate to a particular game, but if someone challenges the consensus, preferring a different criterion, we will have no logical basis for persuading him that he is wrong.

Conclusion

Perhaps I have said enough about the neoclassical revival to suggest why it has been a highly attractive commodity in competition with the behavioral

theories. To some economists at least, it has held open the possibility and hope that important questions that had been troublesome for classical economics could now be addressed without sacrifice of the central assumption of perfect rationality, and hence also with a maximum of *a priori* inference and a minimum of tiresome grubbing with empirical data. I have perhaps said enough also with respect to the limitations of these new constructs to indicate why I do not believe that they solve the problems that motivated their development.

ADVANCES IN THE BEHAVIORAL THEORY

Although they have played a muted role in the total economic research activity during the past two decades, theories of bounded rationality and the behavioral theory of the business firm have undergone steady development during that period. Since surveying the whole body of work would be a major undertaking, I shall have to be satisfied here with suggesting the flavor of the whole by citing a few samples of different kinds of important research falling in this domain. Where surveys on particular topics have been published, I will limit myself to references to them.

First, there has been work in the psychological laboratory and the field to test whether people in relatively simple choice situations behave as statistical decision theory (maximization of expected utilities) says they do. Second, there has been extensive psychological research, in which Allen Newell and I have been heavily involved, to discover the actual microprocesses of human decision making and problem solving. Third, there have been numerous empirical observations - most of them in the form of "case studies"- of the actual processes of decision making in organizational and business contexts. Fourth, there have been reformulations and extensions of the theory of the firm replacing classical maximization with behavioral decision postulates.

Utility Theory and Human Choice

The axiomatization of utility and probability after World War II and the revival of Bayesian statistics opened the way to testing empirically whether people behaved in choice situations so as to maximize subjective expected utility (SEU). In early studies, using extremely simple choice situations, it appeared that perhaps they did. When even small complications were introduced into the situations, wide departures of behavior from the predictions of SEU theory soon became evident. Some of the most dramatic and convincing empirical refutations of the theory have been reported by Kahneman and Tversky (1972), who showed that under one set of circumstances, decision makers gave far too little weight to prior knowledge and based their choices almost entirely on new evidence, while in other circumstances new evidence had little influence on opinions already formed. Equally large and striking departures from the behavior predicted by the SEU theory were found by Kunreuther and his colleagues (1978)

in their studies of individual decisions to purchase or not to purchase flood insurance. On the basis of these and other pieces of evidence, the conclusion seems unavoidable that the SEU theory does not provide a good prediction - not even a good approximation - of actual behavior.

Notice that the refutation of the theory has to do with the *substance* of the decisions, and not just the process by which they are reached. It is not that people do not go through the calculations that would be required to reach the SEU decision - neoclassical thought has never claimed that they did. What has been shown is that they do not even behave *as if* they had carried out those calculations, and that result is a direct refutation of the neoclassical assumptions.

Psychology of Problem Solving

The evidence on rational decision making is largely negative evidence, evidence of what people do *not* do. In the past twenty years a large body of positive evidence has also accumulated about the processes that people use to make difficult decisions and solve complex problems. The body of theory that has been built up around this evidence is called information processing psychology, and is usually expressed formally in computer programming languages. Allen Newell and I have summed up our own version of this theory in our book, *Human Problem Solving* (1972), which is part of a large and rapidly growing literature that assumes an information processing framework and makes use of computer simulation as a central tool for expressing and testing theories.

Information processing theories envisage problem solving as involving very selective search through problem spaces that are often immense. Selectivity, based on rules of thumb or "heuristics", tends to guide the search into promising regions, so that solutions will generally be found after search of only a tiny part of the total space. Satisficing criteria terminate search when satisfactory problem solutions have been found. Thus, these theories of problem solving clearly fit within the framework of bounded rationality that I have been expounding here.

By now the empirical evidence for this general picture of the problem solving process is extensive. Most of the evidence pertains to relatively simple, puzzle-like situations of the sort that can be brought into the psychological laboratory for controlled study, but a great deal has been learned, also, about professional-level human tasks like making medical diagnoses, investing in portfolios of stocks and bonds, and playing chess. In tasks of these kinds, the general search mechanisms operate in a rich context of information stored in human long-term memory, but the general organization of the process is substantially the same as for the simpler, more specific tasks.

At the present time, research in information-processing psychology is proceeding in several directions. Exploration of professional-level skills continues. A good deal of effort is now being devoted also to determining how initial representations for new problems are acquired. Even in simple

problem domains, the problem solver has much latitude in the way he formulates the problem space in which he will search, a finding that underlines again how far the actual process is from a search for a uniquely determined optimum (Hayes & Simon, 1974).

The main import for economic theory of the research in information processing psychology is to provide rather conclusive empirical evidence that the decision-making process in problem situations conforms closely to the models of bounded rationality described earlier. This finding implies, in turn, that choice is not determined uniquely by the objective characteristics of the problem situation, but depends also on the particular heuristic process that is used to reach the decision. It would appear, therefore, that a model of process is an essential component in any positive theory of decision making that purports to describe the real world, and that the neoclassical ambition of avoiding the necessity for such a model is unrealizable (Simon, 1978a).

Organizational Decision Making

It would be desirable to have, in addition to the evidence from the psychological research just described, empirical studies of the process of decision making in organizational contexts. The studies of individual problem solving and decision making do not touch on the many social-psychological factors that enter into the decision process in organizations. A substantial number of investigations have been carried out in the past twenty years of the decision-making process in organizations, but they are not easily summarized. The difficulty is that most of these investigations have taken the form of case studies of specific decisions or particular classes of decisions in individual organizations. To the best of my knowledge, no good review of this literature has been published, so that it is difficult even to locate and identify the studies that have been carried out.⁴ Nor have any systematic methods been developed and tested for distilling out from these individual case studies their implications for the general theory of the decision-making process.

The case studies of organizational decision making, therefore, represent the natural history stage of scientific inquiry. They provide us with a multitude of facts about the decision-making process - facts that are almost uniformly consistent with the kind of behavioral model that has been proposed here. But we do not yet know how to use these facts to test the model in any formal way. Nor do we quite know what to do with the observation that the specific decision-making procedures used by organizations differ from one organization to another, and within each organization, even from one situation to another. We must not expect from these

⁴For leads into the literature, see March & Simon, 1958; March, 1965; Johnsen, 1968, and Dutton & Starbuck, 1971. However, there are large numbers of specific case studies, some of them carried out as thesis projects, some concerned with particular fields of business application, which have never been recorded in these reference sources (e.g., Eliasson, 1976).

data generalizations as neat and precise as those incorporated in neoclassical theory.

Perhaps the closest approach to a method for extracting theoretically relevant information from case studies is computer simulation. By converting empirical evidence about a decision-making process into a computer program, a path is opened both for testing the adequacy of the program mechanisms for explaining the data, and for discovering the key features of the program that account, qualitatively, for the interesting and important characteristics of its behavior. Examples of the use of this technique are Clarkson's (1963) simulation of the decision making of an investment trust officer, Cyert, Feigenbaum and March's (1959) simulation of the history of a duopoly, and Bonini's (1963) model of the effects of accounting information and supervisory pressures in altering employee motivations in a business firm. The simulation methodology is discussed from a variety of viewpoints in Dutton & Starbuck (1971).⁵

Theories of the Business Firm

The general features of bounded rationality - selective search, satisficing, and so on - have been taken as the starting points for a number of attempts to build theories of the business firm incorporating behavioral assumptions. Examples of such theories would include the theory of Cyert and March (1963), already mentioned; Baumol's (1959) theory of sales maximization subject to minimum profit constraints; Marris's (1964) models of firms whose goals are stated in terms of rates of growth; Leibenstein's (1976) theory of "X-inefficiency" that depresses production below the theoretically attainable; Kornai's (1971) dichotomy between supply-driven and demand-driven management; Williamson's (1975) theory of transactional costs; the evolutionary models of Nelson and Winter (1973); Cyert and DeGroot's (1974) models incorporating adaptive learning; Radner's (1975a, 1975b) explicit satisficing models; and others.

Characterized in this way, there seems to be little commonality among all of these theories and models, except that they depart in one way or another from the classical assumption of perfect rationality in firm decision making. A closer look, however, and a more abstract description of their assumptions, shows that they share several basic characteristics. Most of them depart from the assumption of profit maximization in the short run, and replace it with an assumption of goals defined in terms of targets - that is, they are to greater or lesser degree satisficing theories. If they do retain maximizing assumptions, they contain some kind of mechanism that prevents the maximum from being attained, at least in the short

⁵In addition to simulations of the firm, there are very interesting and potentially important efforts to use simulation to build bridges directly from decision theory to political economy. See Orcutt and Caldwells-Wertheimer (1976) and Eliasson (1978).

run. In the Cyert-March theory, and that of Leibenstein, this mechanism can be viewed as producing "organizational slack," the magnitude of which may itself be a function of motivational and environmental variables.

Finally, a number of these theories assume that organizational learning takes place, so that if the environment were stationary for a sufficient length of time, the system equilibrium would approach closer and closer to the classical profit-maximizing equilibrium. Of course they generally also assume that the environmental disturbances will generally be large enough to prevent the classical solution from being an adequate approximation to the actual behavior.

The presence of something like "organizational slack" in a model of the business firm introduces complexity in the firm's behavior in the short run. Since the firm may operate very far from any optimum, the slack serves as a buffer between the environment and the firm's decisions. Responses to environmental events can no longer be predicted simply by analysing the "requirements of the situation," but depend on the specific decision processes that the firm employs. However well this characteristic of a business firm model corresponds to reality, it reduces the attractiveness of the model for many economists, who are reluctant to give up the process-independent predictions of classical theory, and who do not feel at home with the kind of empirical investigation that is required for disclosing actual real-world decision processes.

But there is another side to the matter. If, in the face of identical environmental conditions, different decision mechanisms can produce different firm behaviors, this sensitivity of outcomes to process can have important consequences for analysis at the level of markets and the economy. Political economy, whether descriptive or normative, cannot remain indifferent to this source of variability in response. At the very least it demands that - before we draw policy conclusions from our theories, and particularly before we act on those policy conclusions - we carry out sensitivity analyses to test how far our conclusions would be changed if we made different assumptions about the decision mechanisms at the micro-level.

If our conclusions are robust - if they are not changed materially by substituting one or another variant of the behavioral model for the classical model - we will gain confidence in our predictions and recommendations; if the conclusions are sensitive to such substitutions, we will use them warily until we can determine which microtheory is the correct one.

As reference to the literature cited earlier in this section will verify, our predictions of the operations of markets and of the economy *are* sensitive to our assumptions about mechanisms at the level of decision processes. Moreover, the assumptions of the behavioral theories are almost certainly closer to reality than those of the classical theory. These two facts, in combination, constitute a direct refutation of the argument that the unreality of the assumptions of the classical theory is harmless. We cannot use the *in vacua* version of the law of falling bodies to predict the sinking of a heavy body in molasses. The predictions of the classical and neoclassical

theories and the policy recommendations derived from them must be treated with the greatest caution.

CONCLUSION

There is a saying in politics that "you can't beat something with nothing." You can't defeat a measure or a candidate simply by pointing to defects and inadequacies. You must offer an alternative.

The same principle applies to scientific theory. Once a theory is well entrenched, it will survive many assaults of empirical evidence that purports to refute it unless an alternative theory, consistent with the evidence, stands ready to replace it. Such conservative protectiveness of established beliefs is, indeed, not unreasonable. In the first place, in empirical science we aspire only to approximate truths; we are under no illusion that we can find a single simple formula, or even a moderately complex one, that captures the whole truth and nothing else. We are committed to a strategy of successive approximations, and when we find discrepancies between theory and data, our first impulse is to patch rather than to rebuild from the foundations.

In the second place, when discrepancies appear, it is seldom immediately obvious where the trouble lies. It may be located in the fundamental assumptions of the theory, but it may as well be merely a defect in the auxiliary hypotheses and measurement postulates we have had to assume in order to connect theory with observations. Revisions in these latter parts of the structure may be sufficient to save the remainder.

What then is the present status of the theory of the firm? There can no longer be any doubt that the microassumptions of the theory - the assumptions of perfect rationality - are contrary to fact. It is not a question of approximation; they do not even remotely describe the processes that human beings use for making decisions in complex situations.

Moreover, there is an alternative. If anything, there is an embarrassing richness of alternatives. Today, we have a large mass of descriptive data, from both laboratory and field, that show how human problem solving and decision making actually take place in a wide variety of situations. A number of theories have been constructed to account for these data, and while these theories certainly do not yet constitute a single coherent whole, there is much in common among them. In one way or another, they incorporate the notions of bounded rationality: the need to search for decision alternatives, the replacement of optimization by targets and satisficing goals, and mechanisms of learning and adaptation. If our interest lies in descriptive decision theory (or even normative decision theory), it is now entirely clear that the classical and neoclassical theories have been replaced by a superior alternative that provides us with a much closer approximation to what is actually going on.

But what if our interest lies primarily in normative political economy rather than in the more remote regions of the economic sciences? Is there

then any reason why we should give up the familiar theories? Have the newer concepts of decision making and the firm shown their superiority "for purposes of economic analysis?"

If the classical and neoclassical theories were, as is sometimes argued, simply powerful tools for deriving aggregative consequences that held alike for both perfect and bounded rationality, we would have every reason to retain them for this purpose. But we have seen, on the contrary, that neoclassical theory does not always lead to the same conclusions at the level of aggregate phenomena and policy as are implied by the postulate of bounded rationality, in any of its variants. Hence, we cannot defend an uncritical use of these contrary-to-fact assumptions by the argument that their veridicality is unimportant. In many cases, in fact, this veridicality may be crucial to reaching correct conclusions about the central questions of political economy. Only a comparison of predictions can tell us whether a case before us is one of these.

The social sciences have been accustomed to look for models in the most spectacular successes of the natural sciences. There is no harm in that, provided that it is not done in a spirit of slavish imitation. In economics, it has been common enough to admire Newtonian mechanics (or, as we have seen, the Law of Falling Bodies), and to search for the economic equivalent of the laws of motion. But this is not the only model for a science, and it seems, indeed, not to be the right one for our purposes.

Human behavior, even rational human behavior, is not to be accounted for by a handful of invariants. It is certainly not to be accounted for by assuming perfect adaptation to the environment. Its basic mechanisms may be relatively simple, and I believe they are, but that simplicity operates in interaction with extremely complex boundary conditions imposed by the environment and by the very facts of human long-term memory and of the capacity of human beings, individually and collectively, to learn.

If we wish to be guided by a natural science metaphor, I suggest one drawn from biology rather than physics (Newell & Simon, 1976). Obvious lessons are to be learned from evolutionary biology, and rather less obvious ones from molecular biology. From molecular biology, in particular, we can glimpse a picture of how a few basic mechanisms - the DNA of the Double Helix, for example, or the energy transfer mechanisms elucidated so elegantly by Professor Mitchell - can account for a wide range of complex phenomena. We can see the role in science of laws of qualitative structure, and the power of qualitative, as well as quantitative explanation.

I am always reluctant to end a talk about the sciences of man in the future tense. It conveys too much the impression that these are potential sciences which may some day be actualized, but that do not really exist at the present time. Of course that is not the case at all. However much our knowledge of human behavior falls short of our need for such knowledge, still it is enormous. Sometimes we tend to discount it, because so many of the phenomena are accessible to us in the very activity of living as human beings among human beings that it seems commonplace to us. Moreover,

it does not always answer the questions for which we need answers. We cannot predict very well the course of the business cycle nor manage the employment rate. (We cannot, it might be added, predict very well the time of the next thunderstorm in Stockholm, or manage the earth's climates.)

With all these qualifications and reservations, we do understand today many of the mechanisms of human rational choice. We do know how the information processing system called Man, faced with complexity beyond his ken, uses his information processing capacities to seek out alternatives, to calculate consequences, to resolve uncertainties, and thereby - sometimes, not always - to find ways of action that are sufficient unto the day, that satisfice.

