

From substantive to procedural rationality

Herbert A. Simon

Rational human behavior has been a central object of study in the two distinct disciplines of economics and cognitive psychology. A person unfamiliar with the histories and contemporary research preoccupations of these two disciplines might imagine that there were close relations between them – a constant flow of theoretical concepts and empirical findings from the one to the other and back. In actual fact, communication has been quite infrequent. In the United States, at least, there seem to be no doctoral programs in economics that require their students to master the psychological literature of rational choice, and no psychology programs that insist that their students become acquainted with economic theories of rationality. (I would be gratified to learn that such programs exist, but if they do, they are inconspicuous in the extreme.)

This state of mutual ignorance (perhaps *noblesse oblige* is the right term for it) has a simple explanation. The single term, “rationality,” has had an essentially different meaning in economics from its meaning in cognitive psychology. Traditionally, economists have been interested mostly in what I call “substantive rationality,” while cognitive psychologists have been interested in a quite distinct concept which I shall call “procedural rationality.”

My intent in this paper is, first, to explain the two terms “substantive rationality” and “procedural rationality” – the difference between them, and their relations as well. I shall try to document the fact that during the past 25 years economists have begun to show growing interest in procedural rationality, and to give reasons for believing that procedural rationality will become one of the central concerns of economics over the next 25 years.

Substantive Rationality

Behavior is substantively rational when it is appropriate to the achievement of given goals within the limits imposed by given conditions and constraints.¹ Notice that, by this definition, the rationality of behavior depends upon the actor in only a single respect – his goals. Given these goals, the rational behavior is determined entirely by the characteristics of the environment in which it takes place.

Suppose, for example, that the problem is to minimize the cost of a nutritionally adequate diet, where nutritional adequacy is defined in terms of lower bounds on intakes of certain proteins, vitamins, and minerals, and upper and lower bounds on calories, and where the unit prices and compositions of the obtainable foods are specified. This diet problem can be (and has been) formulated as a straightforward linear-programming problem, and the correct solution found by applying the simplex algorithm or some other computational procedure. Given the goal of minimizing cost and the definition of “nutritionally adequate,” there are no two ways about it – there is only one substantively rational solution.

Classical economic analysis rests on two fundamental assumptions. The first assumption is that the economic actor has a particular goal – *e.g.*, utility maximization or profit maximization. The second assumption is that the economic actor is substantively rational. Given these two assumptions, and given a description of a particular economic environment, economic analysis (descriptive or normative) could usually be carried out using such standard tools as the differential calculus, linear programming, or dynamic programming.

Thus, the assumptions of utility or profit maximization, on the one hand, and the assumption of substantive rationality, on the other, freed economics from any dependence upon psychology. As long as these assumptions went unchallenged, there was no reason why an economist should acquaint himself with the psychological literature on human cognitive processes or human choice. There was absolutely no point at which the findings of psychological research could be injected into the process of economic analysis. The irrelevance of psychology to economics was complete.

1. *Cf.* the entry under “rationality” in J. Gould and W.L. Kolb (9, 1964, pp. 573-574).

Procedural Rationality

Behavior is procedurally rational when it is the outcome of appropriate deliberation. Its procedural rationality depends on the process that generated it. When psychologists use the term "rational," it is usually procedural rationality they have in mind. William James,² for example, uses "rationality" as synonymous with "the peculiar thinking process called reasoning." Conversely, behavior tends to be described as "irrational" in psychology when it represents impulsive response to affective mechanisms without an adequate intervention of thought.

Perhaps because "rationality" resembles "rationalism" too closely, and because psychology's primary concern is with process rather than outcome, psychologists tend to use phrases like "cognitive processes" and "intellective processes" when they write about rationality in behavior. This shift in terminology may have contributed further to the mutual isolation of the concepts of substantive and procedural rationality.

The Study of Cognitive Processes

The process of rational calculation is only interesting when it is non-trivial – that is, when the substantively rational response to a situation is not instantly obvious. If you put a quarter and a dime before a subject and tell him that he may have either one, but not both, it is easy to predict which he will choose, but not easy to learn anything about his cognitive processes. Hence, procedural rationality is usually studied in problem situations – situations in which the subject must gather information of various kinds and process it in different ways in order to arrive at a reasonable course of action, a solution to the problem.

Historically, there have been three main categories of psychological research on cognitive processes: learning, problem solving, and concept attainment. Learning research is concerned with the ways in which information is extracted from one problem situation and stored in such a way as to facilitate the solving of similar problems subsequently. Problem solving research (in this narrower sense) focusses especially upon the complementary roles of trial-and-error procedures and insight

2. W. James (12, 1890, ch. 22).

in reaching problem solutions. Concept attainment research is concerned with the ways in which rules of generalizations are extracted from a sequence of situations and used to predict subsequent situations. Only in recent years, particularly since the Second World War, has there been much unification of theory emerging from these three broad lines of research.

Computational Efficiency

Let us return for a moment to the optimal diet problem which we used to illustrate the concept of substantive rationality. From a procedural standpoint, our interest would lie not in the problem solution – the prescribed diet itself – but in the method used to discover it. At first blush, this appears to be more a problem in the computational mathematics than in psychology. But that appearance is deceptive.

What is the task of computational mathematics? It is to discover the relative efficiencies of different computational processes for solving problems of various kinds. Underlying any question of computational efficiency is a set of assumptions about the capabilities of the computing system. For an omniscient being, there are no questions of computational efficiency, because the consequences of any tautology are known as soon as the premises are stated; and computation is simply the spinning out of such consequences.³

Nowadays, when we are concerned with computational efficiency, we are concerned with the computing time or effort that would be required to solve a problem by a system, basically serial in operation, requiring certain irreducible times to perform an addition, a multiplication, and a few other primitive operations. To compare the simplex method with some other method for solving linear programming problems, we seek to determine how much total computing time each method would need.

The search for computational efficiency is a search for procedural rationality, and computational mathematics is a normative theory of such rationality. In this normative theory, there is no point in prescribing a particular substantively rational solution if there exists no procedure for finding that solution with an acceptable amount of computing effort. So,

3. This statement is a little oversimple in ignoring the distinction between induction and deduction, but greater precision is not needed for our purposes.

for example, although there exist optimal (substantively rational) solutions for combinatorial problems of the travelling-salesman type, and although these solutions can be discovered by a finite enumeration of alternatives, actual computation of the optimum is infeasible for problems of any size and complexity. The combinatorial explosion of such problems simply outraces the capacities of computers, present and prospective.

Hence, a theory of rationality for problems like the travelling-salesman problem is not a theory of best solutions – of substantive rationality – but a theory of efficient computational procedures to find good solutions – a theory of procedural rationality. Notice that this change in viewpoint involves not only a shift from the substantive to the procedural, but a shift also from concern for optimal solutions to a concern for good solutions. I shall discuss this point later.

Computation: Risky Decisions

But now it is time to return to psychology and its concern with computational efficiency. Man, viewed as a thinker, is a system for processing information. What are his procedures for rational choice?

One method of testing a theory of human rational choice is to study choice behavior in relatively simple and well-structured laboratory situations where the theory makes specific predictions about how subjects will behave. This method has been used by a number of investigators – including W. Edwards, G. Pitts, A. Rapaport, and A. Tversky – to test whether human decisions in the face of uncertainty and risk can be explained by the normative concepts of statistical decision theory. This question is particularly interesting because these norms are closely allied, both historically and logically, to the notions of substantive rationality that have prevailed in economics, and make no concessions to computational difficulties – they never choose the computable second-best over the non-computable best.

Time does not permit me to review this extensive literature that this line of inquiry has produced. A recent review by Rapaport⁴ covers experimental tests of SEU (subjective expected utility) maximization, of Bayesian strategies for sequential decisions, and of other models of

4. A. Rapaport and T.S. Wallsten (17, 1972).

rational choice under uncertainty. I think the evidence can be fairly summarized by the statements (a) that it is possible to construct gambles sufficiently simple and transparent that most subjects will respond to them in a manner consistent with SEU theory; but (b) the smallest departures from this simplicity and transparency produce behavior in many or most subjects that *cannot* be explained by SEU or Bayesian models. I will illustrate this statement by just three examples, which I hope are not atypical.

The first is the phenomenon of event matching.⁵ Suppose that you present a subject with a random sequence of *X*'s and *O*'s, of which 70% are *X*'s and 30% *O*'s. You ask the subject to predict the next symbol, rewarding him for the number of correct predictions. "Obviously" the rational behavior is always to predict *X*. This is what subjects almost never do.⁶ Instead, they act as though the sequence were patterned, not random, and guess by trying to extrapolate the pattern. This kind of guessing will lead *X* to be guessed in proportion to the frequency with which it occurs in the sequence. As a result, the sequence of guesses has about the same statistical properties as the original sequence, but the prediction accuracy is lower than if *X* had been predicted each time (58% instead of 70%).

In a recent study by Kahneman & Tversky,⁷ a quite different phenomenon showed up. The rational procedure for combining new information with old is Bayes' Theorem. If a set of probabilities has been assigned to the possible outcomes of an uncertain event, a new evidence is presented, Bayes' Theorem provides an algorithm for revising the prior probabilities to take the new evidence into account. One obvious consequence of Bayes' Theorem is that the more extensive and reliable the new evidence, the greater should be its influence on the new probabilities. Another consequence is that the new probabilities should not depend on the new evidence only, but upon the prior probabilities as well. In the experiments conducted by Kahneman and Tversky, the estimates of subjects were independent of the reliability of the new evidence, and did not appear to be influenced by the prior probabilities at all.

5. J. Feldman (7, 1963).

6. The sole exceptions of which I am aware of are well-known and expert game theorists who served as subjects in this experiment at the Rand Corporation many years ago!

7. D. Kahneman and A. Tversky (14, 1973).

On the other hand, Ward Edwards⁸ has reviewed a large body of experimental evidence describing quite conservative behavior. In these experiments, subjects did not revise prior probability estimates nearly as much as would be called for by Bayes' Theorem. It appears, then that humans can either overrespond to new evidence or ignore it, depending upon the precise experimental circumstances. If these differences in behavior manifest themselves even in laboratory situations so simple that it would be possible for subjects to carry out the actual Bayes calculations, we should be prepared to find variety at least as great when people are required to face the complexities of the real world.

Man's Computational Efficiency

If these laboratory demonstrations of human failure to follow the canons of substantive rationality in choice under uncertainty caused any surprise to economists (and I don't know that they did), they certainly did not to experimental psychologists familiar with human information processing capabilities.

Like a modern digital computer's, Man's equipment for thinking is basically serial in organization. That is to say, one step in thought follows another, and solving a problem requires the execution of a large number of steps in sequence. The speed of his elementary processes, especially arithmetic processes, is much slower, of course, than those of a computer, but there is much reason to think that the basic repertoire of processes in the two systems is quite similar.⁹ Man and computer can both recognize symbols (patterns), store symbols, copy symbols, compare symbols for identity, and output symbols. These processes seem to be the fundamental components of thinking as they are of computation.

For most problems that Man encounters in the real world, no procedure that he can carry out with his information processing equipment will enable him to discover the optimal solution, even when the notion of "optimum" is well defined. There is no logical reason why

8. W. Edwards (6, 1968).

9. In my comparison of computer and Man, I am leaving out of account the greater sophistication of Man's input and output system, and the parallel processing capabilities of his senses and his limbs. I will be primarily concerned here with thinking, secondarily with perceiving, and not at all with sensing or acting.

this needs to be so; it is simply a rather obvious empirical fact about the world we live in – a fact about the relation between the enormous complexity of that world and the modest information-processing capabilities with which Man is endowed. One reason why computers have been so important to Man is that they enlarge a little bit the realm within which his computational powers can match the complexity of the problems. But as the example of the travelling-salesman problem shows, even with the help of the computer, Man soon finds himself outside the area of computable substantive rationality.

The problem space associated with the game of chess is very much smaller than the space associated with the game of life. Yet substantive rationality has so far proved unachievable, both for Man and computer, even in chess. Chess books are full of norms for rational play, but except for catalogues of opening moves, these are procedural rules: how to detect the significant features of a position, what computations to make on these features, how to select plausible moves for dynamic search, and so on.

The psychology of chess playing now has a considerable literature. A pioneer in this research was Professor Adriaan de Groot, of the University of Amsterdam, whose book, *Het Denken van den Schaker*, has stimulated much work on this subject both in Amsterdam, and in our own laboratory at Carnegie-Mellon.¹⁰ These studies have told us a great deal about the thought processes of an expert chess player. First, they have shown how he compensates for his limited computational capacity by searching very selectively through the immense tree of move possibilities, seldom considering as many as 100 branches before making a move. Second, they have shown how he stores in long-term memory a large collection of common patterns of pieces, together with procedures for exploiting the relations that appear in these patterns. The expert chess player's heuristics for selective search and his encyclopedic knowledge of significant patterns are at the core of his procedural rationality in selecting a chess move. Third, the studies have shown how a player forms and modifies his aspirations for a position, so that he can decide when a particular move is "good enough" (satisfices), and can end his search.

Chess is not an isolated example. There is now a large body of data describing human behavior in other problem situations of comparable

10. A. Newell and H.A. Simon (16, 1972); W.G. Chase and H.A. Simon (1, 1973).

complexity. All of the data point in the same direction, and provide essentially the same descriptions of the procedures men use to deal with situations where they are not able to compute an optimum. In all these situations, they use selective heuristics and means-end analysis to explore a small number of promising alternatives. They draw heavily upon past experience to detect the important features of the situation before them, features which are associated in memory with possibly relevant actions. They depend upon aspiration-like mechanisms to terminate search when a satisfactory alternative has been found.

To a moderate extent, this description of choice has been tested outside the laboratory, in even more complex "real-life" situations; and where it has been tested, has held up well. I will only mention as examples Clarkson's wellknown microscopic study of the choices of an investment trust officer,¹¹ and Peer Soelberg's study of the job search and job choice of graduating management students.¹² I cannot supply you with a large number of more recent examples, possibly because they do not exist, or possibly because my own research has taken me away from the area of field studies in recent years. I would be very pleased if some of you could point out to me examples of such work of which I am unaware; because I know that this is a line of inquiry that has been pursued more vigorously during the past decade in Europe, and particularly in Scandinavia, than in the United States.

Contrast this picture of thought processes with the notion of rationality in the classical theory of the firm in its simplest form. The theory assumes that there is given, in addition to the goal of profit maximization, a demand schedule and a cost curve. The theory then consists of a characterization of the substantively rational production decision: for example that the production quantity is set at the level where marginal cost, calculated from that cost curve, equals marginal revenue, calculated from the demand schedule. The question of whether data are obtainable for estimating these quantities or the demand and cost functions on which they are based, is outside the purview of the theory. If the actual demand and cost curves are given, the actual calculation of the optimum is trivial. This portion of economic theory certainly has nothing to do with procedural rationality.

11. G.P.E. Clarkson (3, 1963).

12. P. Soelberg (19, 1967).

Economics' Concern With Procedural Rationality

In my introductory remarks, I said that while economics has traditionally concerned itself with substantive rationality, there has been a noticeable trend, since the Second World War, toward concern also with procedural rationality. This trend has been brought about by a number of more or less independent developments.

The Real World of Business and Public Policy

The first of these developments, which predated the War to some extent, was increasing contact of academic economists with real-world business environments. An early and important product was the 1939 Hall-Hitch paper,¹³ which advanced the heretical proposition that prices are often determined by applying a fixed mark-up to average direct cost rather than by equating them with marginal cost.

I am not concerned here to determine whether Hitch and Hall, or others who have made similar observations, were right or wrong. My point is that first-hand contact with business operations leads to observation of the procedures that are used in reaching decisions, and not simply the final outcomes. Independently of whether the decision processes have any importance for the questions to which classical economics has addressed itself, the phenomena of problem solving and decision making cannot help but excite the interest of anyone with intellectual curiosity who encounters them. They represent a fascinating and important domain of human behavior, which any scientist will wish to describe and explain.

In the United States, in the decade immediately after the Second World War, a number of large corporations invited small groups of academic economists to spend periods of a month or more as "interns" and observers in their corporate offices. Many young economists had their first opportunity, in this way, to try their hands at applying the tools of economic theory to the decisions of a factory department, or a regional sales office.

They found that businessmen did not need to be advised to "set marginal cost equal to marginal revenue." Substantive norms of profit

13. R.L. Hall and C.J. Hitch (10, 1939).

maximization helped real decisions only to the extent that appropriate problem-solving procedures could be devised to implement them. What businessmen needed – from anyone who could supply it – was help in inventing and constructing such procedures, including the means for generating the necessary data. How could the marginal productivity of R & D expenditures be measured? Or of advertising expenditures? And if they could not be, what would be reasonable procedures for fixing these quantities? These – and not abstract questions of profit maximization in a simplified model of the firm – were the questions businessmen wrestled with in their decisions.

Matters were no different with the economists who were increasingly called upon by governments to advise on national fiscal and monetary policy, or on economic development plans. We have the notable example here in the Netherlands of Tinbergen's schemes for target planning¹⁴ – a pioneering example of "satisficing," if I may speak anachronistically. In the face of difficult problems of formulating models, designing appropriate and implementable instruments of measurement, taking account of multidimensional criteria and side conditions, questions of optimization generally faded into the background. The rationality of planning and development models was predominately a procedural rationality.

Operations Research

With the end of the War also, businessmen and government departments began to exhibit an interest in the tools of operations research that had been developed for military application during the War. At the same time, operations analysts began to cast about for peacetime problems to which their skills might be applicable. Since the rapid burgeoning of operations research and management science in industry, and the even more rapid development of powerful analytic tools during the first decade after the War is familiar to all of you, it does not need recounting.

The coincidence of the introduction of the digital computer at the same time undoubtedly accelerated these developments. In fact, it is quite unclear whether operations research would have made any considerable impact on practical affairs if the desk calculator had been its only tool.

14. J. Tinbergen (21, 1952).

Operations research and management science did not alter the economic theory of substantive rationality in any fundamental way. With linear programming and activity analysis it did provide a way of handling the old problems and their solutions without the differential calculus, and the classical theorems of marginalism were soon restated in terms of the new formalism.¹⁵

What was genuinely new for economics in operations research was the concern for procedural rationality – finding efficient procedures for computing actual solutions to concrete decision problems. Let me expand on the specific example with which I am most intimately familiar: decision rules for inventory and work-force smoothing.¹⁶ Here the problem was to devise a decision rule for determining periodically the production level at which a factory should operate. Since the decision for one period was linked to the decisions for the following periods by the inventories carried over, the problem fell in the domain of dynamic programming.

The nub of the problem was to devise a dynamic programming scheme that could actually be carried out using only data that could be obtained in the actual situation. Dynamic programming, in its general formulations, is notoriously extravagant of computational resources. A general algorithm for solving dynamic programming problems would be a non-solution to the real-world decision problem.

The scheme we offered was an algorithm, requiring only a small amount of computing effort, for solving a very special class of dynamic programming problems. The algorithm required the costs to be represented by a quadratic function. This did not mean that we thought real-world cost functions were quadratic; it meant that we thought that many cost functions could be reasonably approximated by a quadratic, and that the deviations from the actual function would not lead to seriously non-optimal decisions. This assumption must, of course, be justified in each individual case, before an application can safely be made. Not only did the quadratic function provide good computational efficiency, but it also greatly reduced the data requirements, because it could be proved that, with this function, only the expected values of predicted variables, and not their higher moments, affected the optimal decision.¹⁷

15. R. Dorfman, P.A. Samuelson and R.M. Solow (5, 1958).

16. C.C. Holt, F. Modigliani, J.F. Muth and H.A. Simon (11, 1960).

17. It is interesting that this same dynamic programming procedure for quadratic cost

This is only part of what was involved in devising a procedurally rational method for making these inventory and production decisions. The problems had also to be solved of translating an aggregate "production level" into specific production schedules for individual products. I will not, however, go into these other aspects of the matter.

Observe of our solution that we constructed a quite classical model for profit maximization, but we did not have the illusion that the model reflected accurately all the details of the real-world situation. All that was expected of the solution was that the *optimal* decision in the world of the model be a *good* decision in the real world. There was no claim that the solution was substantively optimal, but rather that formal optimization in the dynamic programming model was an effective procedural technique for making acceptable decisions (*i.e.*, decisions better than those that would be made without this formal apparatus).

Some operations research methods take the other horn of this dilemma: they retain more of the real-world detail in the model, but then give up, for reasons of computational feasibility, the goal of searching for an optimum, and seek a satisfactory solution instead.¹⁸

Thus, the demands of computability led to two kinds of deviation from classical optimization: simplification of the model to make computation of an "optimum" feasible, or, alternatively, searching for satisfactory, rather than optimal choices. I am inclined to regard both of these solutions as instances of satisficing behavior rather than optimization. To be sure, we can *formally* view these as optimizing procedures by introducing, for example, a cost of computation and a marginal return from computation, and using these quantities to compute the optimal stopping-point for the computation. But the important difference between the new procedures and the classical ones remain. The problem has been shifted from one of characterizing the substantively optimal solution to one of devising practicable computation procedures for making reasonable choices.

functions was invented independently and simultaneously by H. Theil of the Rotterdam School of Economics. See H. Theil (20, 1958). The Rotterdam group was also concerned with concrete applications – in this case to national economic planning in the Netherlands and hence gave a high priority to the demands of procedural rationality in the solutions it developed.

18. I have already mentioned the pioneering work of J. Tinbergen in the Netherlands, who employed national planning models that aimed at target values of key variables instead of an optimum.

Imperfect Competition

More than a century ago, Cournot identified a problem that has become the permanent and ineradicable scandal of economic theory. He observed that where a market is supplied by only a few producers, the notion of profit maximization is ill-defined. The choice that would be substantively rational for each actor depends on the choices made by the other actors; none can choose without making assumptions about how others will choose.

Cournot proposed a particular solution for the problem, which amounted to an assumption about the *procedure* each actor would follow: each would observe the quantities being produced by his competitors, and would assume these quantities to be fixed in his own calculations. The Cournot solution has often been challenged, and many alternative solutions have been proposed – conjectural variations, the kinky demand curve, market leadership, and others. All of them rest on postulates about the decision process, in particular, about the information each decision maker will take into account, and the assumptions he will make about the reactions of the others to his behavior.

I have referred to the theory of imperfect competition as a “scandal” because it has been treated as such in economics, and because it is generally conceded that no defensible formulation of the theory stays within the framework of profit maximization and substantive rationality. Game theory, initially hailed as a possible way out, provided only a rigorous demonstration of how fundamental the difficulties really are.

If perfect competition were the rule in the markets of our modern economy, and imperfect competition and oligopoly rare exceptions, the scandal might be ignored. Every family, after all, has some distant relative it would prefer to forget. But imperfect competition is not a “distant relative,” it is the characteristic form of market structure in a large part of the industries in our economy.

In the literature on oligopoly and imperfect competition one can trace a gradual movement toward more and more explicit concern with the processes used to reach decisions, even to the point – unusual in most other areas of economics – of trying to obtain empirical data about these processes. There remains, however, a lingering reluctance to acknowledge the impossibility of discovering at last “The Rule” of substantively rational behavior for the oligopolist. Only when the hope of that

discovery has been finally extinguished will it be admitted that understanding imperfect competition means understanding procedural rationality.

This change in viewpoint will have large effects on many areas of economic research. There has been a great burgeoning, for example, of "neo-classical" theories of investment – theories that undertake to deduce the rates of investment of business firms from the assumptions of profit maximization and substantive rationality. Central to such theories is the concept of "desired capital" – that is, the volume of capital that would maximize profits. Jorgenson, for example, typically derives "desired capital" by an argument that assumes a fixed price for the firm's products and a production function of the Cobb-Douglas type, all in the absence of uncertainty.¹⁹ Under these assumptions, he shows that the optimal level of capital is proportional to output.

Since the data which Jorgenson and others use to test these theories of investment derive mostly from oligopolistic industries, their definitions of rationality are infected with precisely the difficulties we have been discussing. Can we speak of the capital desired by General Motors or the American Can Company without considering their expectations for size and share of market or the interactions of these expectations with price policies and with the responses of competitors?²⁰ Under conditions of imperfect competition, one can perhaps speak of the procedural rationality of an investment strategy, but surely not of its substantive rationality. At most, the statistical studies of investment behavior show that some business firms relate their investments to output; they do not show that such behavior is predictable from an objective theory of profit maximization. (And if that is what is being demonstrated, what is the advantage of doing it by means of elaborate statistical studies of public data, rather than by making inquiries or observations of the actual decision processes in the firms themselves?)

Expectations and Uncertainty

Making guesses about the behavior of a competitor in an oligopolistic industry is simply a special case of forming expectations in order to make

19. D.W. Jorgenson (13, 1963).

20. R.M. Cyert, E.A. Feigenbaum and J.G. March (4, 1959).

decisions under uncertainty. As economics has moved from statics to dynamics – to business cycle theory, growth theory, dynamic investment theory, theory of innovation and technological change – it has become more and more explicit in its treatment of uncertainty.

Uncertainty, however, exists not in the outside world, but in the eye and mind of the beholder. We need not enter into philosophical arguments as to whether quantum-mechanical uncertainty lies at the very core of nature, for we are not concerned with events at the level of the atom. We are concerned with how men behave rationally in a world where they are often unable to predict the relevant future with accuracy. In such a world, their ignorance of the future prevents them from behaving in a substantively rational manner; they can only adopt a rational choice procedure, including a rational procedure for forecasting or otherwise adapting to the future.

In a well-known paper, my former colleague, John F. Muth,²¹ proposed to objectify the treatment of uncertainty in economics by removing it from the decision maker to nature. His hypothesis is “that expectations of firms (or, more generally, the subjective probability distribution of outcomes) tend to be distributed, for the same information set, about the prediction of the theory (or the ‘objective’ probability distributions of outcomes).” In application this hypothesis involves setting the expected value (in the statistical sense) of a future economic variable equal to its predicted value.

Muth’s proposal is ingenious and important. Let us see exactly what it means. Suppose that a producer has an accurate knowledge of the consumer demand function and the aggregate supply function of producers in his industry. Then he can estimate the equilibrium price – the price at which the quantities that producers will be induced to offer will just balance demand. Muth proposes essentially that each producer takes this equilibrium price as his price forecast. If random shocks with zero expected value are now introduced into the supply equation, and if producers continue to act on price forecasts made in the manner just described, then the forecast price will equal the expected value of the actual price.

Notice that the substantively rational behavior for the producer would be to produce the quantity that would be optimal for the price that is *actually* realized. The assumption of Muth’s model that the random

21. J.F. Muth (15, 1961).

shocks are completely unpredictable makes this impossible. The producer then settles for a procedure that under the assumptions of the model will give him an unbiased prediction of the price. Nor, as Muth himself notes, will this procedure be optimal, even under uncertainty, unless the loss function is quadratic.

Uncertainty plays the same innocuous role in the optimal linear production smoothing rule I described earlier,²² which is closely related to Muth's analysis. Here the explicit assumption of a quadratic cost function makes it possible to prove that only the expected values and not the higher moments of predicted variables are relevant to decision. This does not mean that action based on unbiased estimates is substantively rational, independently of the variances of those estimates. On the contrary, performance can always be improved if estimation errors can be reduced.

Even if it turns out to be empirically true that the forecasts of business firms and other economic actors are unbiased forecasts of future events, this finding will have modest implications for the nature of human rationality. Unbiased estimation can be a component of all sorts of rational and irrational behavior rules.

In an earlier section I commented on the psychological evidence as to human choice in the face of uncertainty. Only in the very simplest situations does behavior conform reasonably closely to the predictions of classical models of rationality. But even this evidence exaggerates the significance of those classical models for human affairs; for all of the experiments are limited to situations where the alternatives of choice are fixed in advance, and where information is available only from precisely specified sources.

Once we become interested in the procedures – the rational processes – that economic actors use to cope with uncertainty, we must broaden our horizons further. Uncertainty not only calls forth forecasting procedures; it also calls forth a whole range of actions to reduce uncertainty, or at least to make outcomes less dependent upon it. These actions are of at least four kinds:

1. Intelligence actions to improve the data on which forecasts are based, to obtain new data, and to improve the forecasting models;
2. Actions to buffer the effects of forecast errors: holding inventories, insuring, and hedging, for example;

22. C.C. Holt, F. Modigliani, J.F. Muth and H.A. Simon (11, 1960).

3. Actions to reduce the sensitivity of outcomes to the behavior of competitors: steps to increase product and market differentiation, for example;
4. Actions to enlarge the range of alternatives whenever the perceived alternatives involve high risk.

A theory of rational choice in the face of uncertainty will have to encompass not only the topic of forecasting, but these other topics as well. Moreover, it will have to say something about the circumstances under which people will (or should) pursue one or the other of these lines of action.

Confronting a list of contingencies of this sort fills many economists with malaise. How can an unique answer be found to the problem of choice if all of these considerations enter it? How much more attractive is classical economics, in allowing strong conclusions to be drawn from a few *a priori* assumptions, with little need for empirical observation!

Alas, we must take the world as it is. As economics becomes more concerned with procedural rationality, it will necessarily have to borrow from psychology or build for itself a far more complete theory of human cognitive processes than it has had in the past. Even if our interest lies in normative rather than descriptive economics, we will need such a theory. There are still many areas of decision – particularly those that are ill-structured – where human cognitive processes are more effective than the best available optimization techniques or artificial intelligence methods. Every Class A chessplayer plays a far better game than any existing chess-playing computer program. A great deal can still be learned about effective decision procedures by studying how humans make choices.

The human mind is programmable: it can acquire an enormous variety of different skills, behavior patterns, problem-solving repertoires, and perceptual habits. Which of these it will acquire in any particular case is a function of what it has been taught and what it has experienced. We can expect substantive rationality only in situations that are sufficiently simple as to be transparent to this mind. In all other situations, we must expect that the mind will use such imperfect information as it has, will simplify and represent the situation as it can, and will make such calculations as are within its powers. We cannot expect to predict what it will do in such situations unless we know what information it has, what forms of representation it prefers, and what algorithms are available to it.

There seems to be no escape. If economics is to deal with uncertainty,

it will have to understand how human beings in fact behave in the face of uncertainty, and by what limits of information and computability they are bound.

The Empirical Study of Decision Making

As I remarked earlier, since my own recent research has removed me from the study of decision making in organization settings, I am not in a position to comment on the current state of our empirical knowledge of organizational decision making.

In trying to understand procedural rationality as it relates to economics, we do not have to limit ourselves, however, to organizational studies. I have already commented upon the understanding we have gained, during the past twenty years, of human problem solving processes – mostly by study in the laboratory, using puzzle-like tasks. Most of these studies have used naive subjects performing tasks with which they had little or no previous experience. In one case, however – the research on chessplaying – an intensive investigation has been made of highly skilled, professional performance, and a body of theory constructed to explain that performance.

Chess may seem a rather esoteric domain, but perhaps business is no less esoteric to those who do not practice it. There is no reason to believe that the basic human faculties that a chess professional of twenty years' experience brings to bear upon his decisions are fundamentally different from the faculties used by an experienced professional businessman. In fact, to the extent that comparable studies of business decision making have been carried out, they give us reason to believe in the basic similarity of those faculties.

On the basis of the research on chessplayers, what appears to distinguish expert from novice is not only that the former has a great quantity and variety of information, but that his perceptual experience enables him to detect familiar patterns in the situations that confront him, and by recognizing these patterns, to retrieve speedily a considerable amount of relevant information from long-term memory.²³ It is this perceptual experience that permits the chessmaster to play, and usually win, many simultaneous games against weaker opponents, taking only a

23. A.D. de Groot (8, 1965); W.G. Chase and H.A. Simon (2, 1973).

few seconds for each move. It is very likely similar perceptual experience about the world of business that enables the executive to react "intuitively," without much awareness of his own cognitive processes, to business situations as they arise.

There is no reason to suppose that the theory of cognitive processes that will emerge from the empirical study of the chessmaster's or businessman's decision processes will be "neat" or "elegant," in the sense that the Laws of Motion or the axioms of classical utility theory are neat and elegant. If we are to draw an analogy with the natural sciences, we might expect the theory of procedural rationality to resemble molecular biology, with its rich taxonomy of mechanisms, more closely than either classical mechanics or classical economics. But as I suggested earlier, an empirical science cannot remake the world to its fancy: it can only describe and explain the world as it is.

A major source of complication in theories of professional decision making is the dependence of decisions upon large quantities of stored information and previously learned decision procedures. This is true not only at an individual psychological level, but also at a social and historical level. The play of two chessplayers differs as a result of differences in what they know about chess: no less do the decisions of two businessmen differ as a result of differences in what they know about business. Moreover Bobby Fisher, in 1972, played chess differently from Paul Morphy, in 1861. Much of that latter difference was the result of the knowledge of the game that has cumulated over the century through the collective experience of the whole society of professional chessplayers.

Economics, like chess, is inevitably culture-bound and history-bound. A business firm equipped with the tools of operations research does not make the same decisions as it did before it possessed those tools. The considerable secular decline over recent years of inventories held by American firms is probably due in considerable part to this enhancement of rationality by new theory and new computational tools.

Economics is one of the sciences of the artificial.²⁴ It is a description and explanation of human institutions, whose theory is no more likely to remain invariant over time than the theory of bridge design. Decision processes, like all other aspects of economic institutions, exist inside human heads. They are subject to change with every change in what human beings know, and with every change in their means of calculation.

24. H.A. Simon (18, 1969).

For this reason the attempt to predict and prescribe human economic behavior by deductive inference from a small set of unchallengeable premises must fail and has failed.

Economics will progress as we deepen our understanding of human thought processes; and economics will change as human individuals and human societies use progressively sharpened tools of thought in making their decisions and designing their institutions. A body of theory for procedural rationality is consistent with a world in which human beings continue to think and continue to invent; a theory of substantive rationality is not.

Conclusion

In this paper I have contrasted the concept of substantive rationality that has dominated classical economics with the concept of procedural rationality that has prevailed in psychology. I have described also some of the concerns of economics that have forced that discipline to begin to concern itself with procedural rationality – with the actual processes of cognition, and with the limits on the human organism that give those processes their peculiar character.

The shift from theories of substantive rationality to theories of procedural rationality requires a basic shift in style, from an emphasis on deductive reasoning from a tight system of axioms to an emphasis on detailed empirical exploration of complex algorithms of thought. Undoubtedly the uncongeniality of the latter style to economists has slowed the transition. As economics becomes more and more involved in the study of uncertainty, more and more concerned with the complex actuality of business decision making, the transition will become inevitable. Wider and wider areas of economics will replace the oversimplified assumptions of the omniscient decision maker with a realistic characterization of the limits on Man's rationality, and the consequences of those limits for his economic behavior.

References

1. Chase, W.G. and H.A. Simon, "Skill in Chess," *American Scientist*, LXI (1973), pp. 394-403.
2. Chase, W.G. and H.A. Simon, "Perception in Chess," *Cognitive Psychology*, IV (1973), pp. 55-81.

3. Clarkson, G.P.E., "A Model of the Trust Investment Process," in Feigenbaum and Feldman (eds.) *Computers and Thought*, New York, 1963.
4. Cyert, R.M., E.A. Feigenbaum and J.G. March, "Models in a Behavioral Theory of the Firm," *Behavioral Science*, IV (1959), pp. 81-95.
5. Dorfman, R., P.A. Samuelson and R.M. Solow, *Linear Programming and Economic Analysis*, New York, 1958.
6. Edwards, W., "Conservatism in Human Information Processing," in B. Kleinmuntz (ed.), *Formal Representation of Human Judgment*, New York, 1968, pp. 17-52.
7. Feldman, J., "Simulation of Behavior in the Binary Choice Experiment," in Feigenbaum and Feldman (eds.), *Computer and Thought*, New York, 1963.
8. Groot, A.D. de, *Thought and Choice in Chess*, Den Haag, 1965.
9. Gould, J. and W.L. Kolb (eds.), *A Dictionary of the Social Sciences*, Glencoe, 1964.
10. Hall, R.L. and C.J. Hitch, "Price Theory and Business Behavior," *Oxford Economic Papers*, II (1939), pp. 12-45.
11. Holt, C.C., F. Modigliani, J.F. Muth and H.A. Simon, *Planning Production, Inventories and Work Force*, Englewood Cliffs, N.J., 1960.
12. James, W., *Principles of Psychology*, 1890.
13. Jorgenson, D.W., "Capital Theory and Investment Behavior," *American Economic Review, Papers and Proceedings*, LIII (1963), pp. 247-259.
14. Kahneman, D. and A. Tversky, "On the Psychology of Prediction," *Psychological Review*, LXXX (1973), pp. 237-251.
15. Muth, J.F., "Rational Expectations and the Theory of Price Movements," *Econometrica*, XXIX (1961), pp. 315-335.
16. Newell, A. and H.A. Simon, *Human Problem Solving*, Englewood Cliffs, N.J., 1972.
17. Rapaport, A. and T.S. Wallsten, "Individual Decision Behavior," *Annual Review of Psychology*, XXIII (1972), pp. 131-176.
18. Simon, H.A., *The Sciences of the Artificial*, Cambridge, Mass., 1969.
19. Soelberg, P., *A Study of Decision Making: Job Choice*, Ph.D. Dissertation, Carnegie-Mellon University, 1967.
20. Theil, H., *Economic Forecasts and Policy*, Amsterdam, 1958.
21. Tinbergen, J., *On the Theory of Economic Policy*, Amsterdam, 1952.