ECONOMIC RATIONALITY:
Adaptive Artifice

Herbert A. Simon
Carnegie-Mellon University

C.I.P. Working Paper #414

File: SABE.MSS This paper is a draft of Chapter Two in the forthcoming edition of THE SCIENCES OF THE ARTIFICIAL.

Chapter 2

ECONOMIC RATIONALITY:

Adaptive Artifice

Scarcity is a central fact of life. Because resources -- land, money, fuel, time, attention -- all are scarce in relation to our uses for them, it is a task of rationality to allocate them. The discipline of economics has taken the performance of that task as its focal concern.

Among all the social sciences, economics exhibits in purest form the artificial component in human behavior, and does so at three or more levels: the level of the individual actor (economic man or business firm), the level of markets, and the level of an entire economy. At all these levels, the outer environment is defined by available technologies and by the behavior of other economic actors, other markets, or other economies. The inner environment is defined by the system's goals and by its capabilities for rational, adaptive behavior.

Economics will serve well to illustrate how outer and inner environment interact, and in particular, how an intellligent system's adjustment to its outer environment (its substantive rationality) is conditioned by its ability to discover appropriate adaptive behavior (its procedural rationality).

The Economic Actor

In the theory of the business firm as it appears in elementary textbooks, a firm managed by an "entrepreneur" aims at maximizing its profit, but in such simple circumstances that its computational capabilities are not in question. Faced with a cost curve relating dollar expenditures to numbers of widgets produced, and a revenue curve relating dollar receipts to numbers of widgets sold, a widget company can control how much it produces (and sells). The goal (profit maximization) fully defines the firm's inner environment; the cost and revenue curves define the outer environment to which it must adapt.[1] We readily deduce that the rational entrepreneur will select the output quantity that yields the greatest positive difference between total revenue and total cost. Given the cost and revenue curves, anyone schooled in the elements of the differential calculus can find this optimal quantity by taking a simple derivative, setting it equal to zero, and solving for quantity as the dependent variable.

Here are all the elements of what, in the last chapter, we were calling an artificial system. The system adapts to an outer environment, subject only to the goal defined by its inner environment. To predict its behavior, we need information about the outer environment and the goal, but we need no information about the process used to compute the optimal output quantity. Above, I called the concept of rationality that applies to this kind of situation *substantive rationality*, in contrast to the *procedural rationality* required when the adaptation process is itself problematic.[2]

We can amplify this bare-bones theory of the firm in many directions, interpreting it either in a positive or a normative sense -- either as a description of how business firms behave, or as advice to them on how to maximize profits. In fact, it is taught in both senses to students in business schools and universities. Once, when I expressed some doubts about the theory's veridicality as a description of actual business behavior, a colleague remarked, "Well, if they don't act this way now, they will after they have graduated from our school." A self-confirming theory.

The Normative Theory

But let us follow the normative path for a moment. How does one use the theory to give advice? It is doubtful that any business firm, already motivated to make profit, would need to be advised to choose the production quantity that will maximize the difference between revenue and cost. The normative problem becomes interesting when we ask in more detail

[1]Notice that I am drawing the line between outer and inner environment not at the boundary of the firm, but at the skin of the directing entrepreneur. In this conception, the factory is part of the external technology, the brain -- assisted perhaps by computers -- is the internal.

[2]H. A. Simon, "Rationality as Process and as Product of Thought," *American Economic Review*, *68*(2):1-16 (1978).

how the firm goes about discovering that maximizing quantity. While historical cost accounting records of factory operations may produce an estimate of the cost curve, determining the price at which a specific output can be sold is likely to be more difficult. At best, only a probability distribution can be estimated, and the business firm might ask whether it should maximize the expected value of its profit or should seek a compromise between maximizing profit and minimizing risk (responding to a shadowy "utility function" that lurks somewhere in the recesses of the entrepreneur's mind).

In real life, the business firm must choose not only the quantity of its product but the quality as well -- or the assortment of different kinds of widgets that can be manufactured with common factory facilities. Now it must decide how to allocate its factory machinery to produce the most profitable combination of widgets. So we proceed step by step from the very simple model of the firm that serves as the textbook example to the complexities of real firms in the real world of business. And as we take these steps toward realism, the normative problem gradually changes from finding the right course of action (substantive rationality) to finding a way of calculating what that course of action is (procedural rationality).

Procedural Rationality

If modeling the real world faithfully introduces exuberant complexity into the business firm's outer environment of cost and revenue relations, it introduces corresponding complexity into the inner environment. For the inner constraints on adaptation include uncertainty about the outer environment -- the actual levels of revenue and cost curves -- as well as limits on the calculation capabilities available for solving the optimization problem with those facts. The normative theory of the firm becomes a theory of estimation under uncertainty and a theory of computation -- decidedly non-trivial theories as the obscurities and complications of information and computation increase.

Today several new branches of applied science assist the firm to achieve procedural rationality.[3] One of them is operations research (alias management science); another is

[3]For a brief survey of these developments, see H. A. Simon, "On How to Decide What to Do," *The Bell Journal of Economics*, 9:494-507 (1978). For an estimate of their impact on management, see H. A. Simon, *The New Science of Management Decision*, Revised Edition (Englewood Cliffs, N.J.: Prentice-Hall, 1977), chapters 2 and 4.

organization theory. Operations research provides algorithms for handling difficult multivariate decision problems, sometimes involving uncertainty. The simplex method for linear programming is an important algorithm of this kind; queuing theory is another; linear decision rules for inventory control and production smoothing are a third. These and similar computational procedures are widely used in business decision-making today.

These powerful algorithms impose a sufficiently strong mathematical structure on the decision problem to permit solutions to be found -- even when there are hundreds or thousands of variables -- with only an amount of search that human beings aided by computers can accomplish with a tolerable expenditure of effort. In linear programming, for example, the key simplifying assumption is that the criterion function (utility function) as well as all the constraints and other relations among variables are linear. For linear decision rules, the key assumptions are that the criterion function is quadratic and that the dynamic relations in the system are represented by linear difference equations. Accepting these severe limits on the mathematical structure of the problems makes it possible to compute decisions that are "optimal" whenever the assumptions hold exactly.

Of course, the decision that is optimal in the simplified model will seldom be optimal in the real world. The decision maker has a choice between optimal decisions for an imaginary simplified world or decisions that are "good enough," that satisfice, for a world approximating the complex real one more closely. The technique of heuristic search makes much weaker demands on the problem structure than do linear programming or linear decision rules, but it can generally find only satisfactory solutions, not the optimal one. On the other side of the coin, heuristic search can handle combinatorial problems (e.g., factory scheduling problems) that are too large for exact solution even with the biggest computers. It is an especially powerful problem-solving and decision-making tool for people who are unassisted by any computers except their own minds, hence must make extensive simplifications in order to find even approximate solutions.

The mathematical tools of operations research have been applied mainly to business decisions at the middle levels of management: production scheduling and inventory control, choice of optimal product or input mix, location of warehouses, and the like. A vast range of top management decisions relating to the long-term future of the firm -- its general investment policies, its research and development programs, specialization and diversification, the recruitment, development, and retention of managerial talent -- are still mainly handled by traditional methods, that is to say, by relying on the "judgment" of experienced executives.

It is not that such top-level decisions cannot be formalized, but formalization, in actual fact, must discard essential facets of the real-world situations and, worse yet, must postulate parameters and variables that cannot be measured. Uncertainty, computational complexity, and lack of operationality have been the principal barriers to extending operations research techniques to the upper levels of management. Qualitative concerns often elude the classical OR models, since human thinking and decision-making do not depend on the presence of numbers in the way that OR techniques do. Heuristic search, which is as applicable to non-numerical as to numerical information, has a broader domain of applicability. It is, in fact, the principal engine for human problem-solving, as we shall see in the next two chapters.

Satisficing

What a person *cannot* do he *will not* do, no matter how much he wants to do it. Normative economics has shown that exact solutions to the larger optimization problems of the real world are simply not within reach or sight. In the face of this complexity the real-world business firm turns to procedures that find good enough answers to questions whose best answers are unknowable. Thus, normative microeconomics, by showing real-world optimization to be impossible, demonstrates that economic man is in fact a satisficer, a person who accepts "good enough" alternatives, not because he prefers less to more but because he has no choice.

This conclusion still leaves open the question of whether the distinction between satisficing and optimizing matters -- whether the attainable approximate solutions to economic decision problems are so far from the unattainable optimal solutions that we need to take account of the discrepancy between them in descriptive economics. In particular, it is often argued that, for purposes of analysing whole markets or the entire economy, the gap is of no account, and that we are justified in employing the optimizing model in spite of the acknowledged unreality of its assumptions. While I think it does matter, and matter a great deal, reviewing this old argument here would take me away from my main theme, which is to show how the behavior of an artificial system may be strongly influenced by the limits of its adaptive capacities.

Markets and the Economy

The center of interest in traditional economic theory does not lie in the behavior of the individual consumer or business firm. Economics has been concerned primarily with larger artificial systems: the economy and its major components, markets. The social function of markets is to coordinate the decisions and behavior of multitudes of individual economic actors -- to guarantee that the quantity of Brussels sprouts shipped to market bears some reasonable relation to the quantity that consumers will buy and eat, and that the price at which they can be sold bears a reasonable relation to the cost of producing them. Any society that is not a subsistence economy, but in which there is substantial specialization and division of labor, needs mechanisms to perform this coordinative function.

Markets are only one, however, in a spectrum of methods of coordination, and no society relies exclusively upon any single one of the mechanisms that are available. For some purposes, statistics provide an adequate basis for coordinating behavior patterns. Highway planning, for example, relies on estimates of road usage that depend, in turn, on statistically stable patterns of driving behavior.

For other purposes, a society may rely on bargaining and negotiating processes to coordinate individual behaviors -- for instance, to secure wage agreements between employers and unions, or to form legislative majorities. Where only two or a few parties are involved, bargaining processes may be unavoidable but cumbersome, because they do not usually provide a unique way of determining a settlement. They are generally employed for lack of a better alternative in situations where conflict of interest is prominent.

For still other coordinative functions, societies employ hierarchic organizations -- business, governmental, and educational -- with lines of formal authority running from top to bottom. Finally, as a way for making certain important decisions and for selecting persons to occupy positions of authority, some societies employ a wide variety of balloting procedures.

Although statistical averaging, markets, bargaining, hierarchy, and voting can all be found as components of the system of coordination in almost any society,[4] the mix of methods and their use for different tasks of social allocation vary tremendously from one nation or culture to another. For example, whereas markets figure most prominently in coordinating economic activities in capitalist countries, hierarchic organizations play the largest role in socialist countries. But that is too simple a formula to describe the realities, which always exhibit a blend of all the mechanisms of coordination. The economic units in capitalist societies are mostly business firms, which are themselves hierarchic organizations, some of enormous size, that make only a modest use of markets in their internal functioning. Conversely, socialist states use market prices to a growing extent to supplement hierarchic control in achieving inter-industry coordination. In both kinds of societies a great deal of bargaining, and even balloting, takes place, but usually much more visibly and formally in capitalist than in socialist countries.

The Invisible Hand

In examining the processes of social coordination, economics has given top billing -- sometimes almost exclusive billing -- to the market mechanism. It is indeed a remarkable mechanism which, under many circumstances, can bring it about that the producing, consuming, buying and selling behavior of enormous numbers of people, each one responding only to his or her own selfish interests, leads to an allocation of resources that clears markets -- that does in fact balance the production with the consumption of Brussels sprouts, and all the other commodities the economy uses.

Only relatively weak conditions need be satisfied to produce such an equilibrium. What is mainly required is that prices drop in the face of an excess supply and that quantities produced and offered decline when prices are lowered. Any number of dynamic systems can be formulated with these properties, and these systems will seek equilibrium and oscillate

---

[4] R.A. Dahl and C.E. Lindblom, *Politics, Economics, and Welfare* (New York: Harper & Brothers, 1953).

stably around it over some wide range of conditions.[5]

In contemporary neoclassical economics, much stronger claims are made for the price mechanism than merely that it clears markets. If certain rather strong assumptions are granted (essentially, assumptions of perfect competition and of profit or utility maximization by the economic actors), it can be shown rigorously that the equilibrium produced by the market will be optimal in the sense that it could not be shifted so as to make everyone simultaneously better off. These are the familiar and famous propositions of the existence and Pareto optimality of competitive equilibrium that have been formalized so elegantly by Arrow, Debreu, Hurwicz, and others.[6]

The Pareto optimality theorems stretch credibility, so far as the markets of the real world are concerned, because they require substantive rationality of the kinds we found implausible in the theory of the firm. Markets populated by consumers and producers who satisfice instead of optimizing do not meet the conditions on which the theorems rest. For this reason, I would like to focus on the market-clearing properties of economic systems rather than on claims that might be made for the optimality of their equilibrium positions. As was said of the dancing dog, "The marvel is not that it dances well, but that it dances at all." One can marvel that the productive efforts and consumption activities of a great population can be brought into patterned order simply by allowing people to exchange goods at prices mutually agreed upon, without claiming that the pattern will have optimal properties.

We have become accustomed to the idea that a natural system like the human body or an ecosystem regulates itself. To explain the regulation, we look for feedback loops rather than a central planning and directing body. But somehow, our intuitions about self-regulation without central direction do not carry over to the artificial systems of human society. I retain vivid memories of the astonishment and disbelief always expressed by the architecture students to whom I taught urban land economics many years ago when I pointed to medieval

---

[5]Stability may require, for example, that demand respond more sensitively to price than supply. For a classic treatment, see P.A. Samuelson, *Foundations of Economic Analysis* (Cambridge, Mass.: Harvard University Press, 1948), Chapter 9.

[6]See Gerard Debreu, *Theory of Value: An Axiomatic Analysis of Economic Equilibrium* (New York: Wiley, 1959).

cities as marvelously patterned systems that had mostly just "grown" in response to myriads of individual human decisions. To my students, a pattern implied a planner in whose mind it had been conceived and by whose fiat it had been implemented. The idea that a city could acquire its pattern as "naturally" as a snowflake was foreign to them. They reacted to it as many Christian fundamentalists responded to Darwin: no design without a Designer!

Marxist fundamentalists reacted in a similar way when they shouldered the task of constructing the new socialist economies of Eastern Europe. It took them some thirty years to realize that markets and prices might play a constructive role in socialist economies and might even have important advantages over central planning as tools for the allocation of resources. Even today, the lesson has been learned only imperfectly. My sometime teacher, Oscar Lange, was one of the pioneers who carried this heretical notion to Poland after the Second World War, risking his career and his life for the idea.

Market processes commend themselves primarily because they avoid placing on a central planning mechanism a burden of calculation that such a mechanism, however well buttressed by the largest computers, could not sustain. They conserve information and calculation by making it possible to assign decisions to the actors who are most likely to possess the information (most of it local in origin) that is relevant to those decisions.

No one has characterized market mechanisms better than Friederich von Hayek who, in the decades after World War II, was their leading interpreter and defender. His defense did not rest primarily upon the supposed optimum attained by them, but rather upon the limits of the inner environment -- the computational limits of human beings. He states the case eloquently in his well-known article, "The Use of Knowledge in Society," published in 1945[7] (pp. 519-20):

> What is the problem we wish to solve when we try to construct a rational economic order?
>
> On certain familiar assumptions the answer is simple enough. *If* we possess all the relevant information, *if* we can start out from a given system of preferences and *if* we command complete knowledge of available means, the problem which remains is purely one of logic. . . . the marginal rates of substitution between any two commodities or factors must be the same in all their different uses.

---

[7] *The American Economic Review*, 35:519-30 (September, 1945).

> This, however, is emphatically *not* the economic problem which society faces. And the economic calculus which we have developed to solve this logical problem, though an important step toward the solution of the economic problem of society, does not yet provide an answer to it. The reason for this is that the "data" from which the economic calculus starts are never for the whole society "given" to a single mind which could work out the implications, and can never be so given.
>
> . . . the knowledge of the circumstances of which we must make use never exists in concentrated or integrated form, but solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess. The economic problem of society is thus not merely a problem of how to allocate "given" resources -- if "given" is taken to mean given to a single mind which deliberately solves the problem set by these "data." It is rather a problem of how to secure the best use of resources known to any of the members of society, for ends whose relative importance only these individuals know. Or, to put it briefly, it is a problem of the utilization of knowledge not given to anyone in its totality.

Of the price system as solution to this problem, he says (pp. 526-7):

> We must look at the price system as such a mechanism for communicating information if we want to understand its real function . . . The most significant fact about this system is the economy of knowledge with which it operates, or how little the individual participants need to know in order to be able to take the right action. In abbreviated form, by a kind of symbol, only the most essential information is passed on, and passed on only to those concerned. It is more than a metaphor to describe the price system as a kind of machinery for registering change, or a system of telecommunications which enables individual producers to watch merely the movement of a few pointers, as an engineer might watch the hands of a few dials, in order to adjust their activities to changes of which they may never know more than is reflected in the price movement. . . .

Much has been written about the market mechanism -- in explanation, in criticism, and in defense. Two-thirds of the world's population, not only primitive cultures but traditional subsistence agriculture as well, make minimal use of it in their daily lives, while at the same time, it is the central organizing device of the modern industrial world. As von Hayek points out, its most striking characteristic is the way it reduces and localizes informational and computational requirements.

In the literature of modern economics, therefore, there is not one market mechanism; there are two. The ideal market mechanism of general equilibrium theory is a dazzling piece of machinery that combines the optimizing choices of a host of substantively rational economic actors into a collective decision that is Pareto optimal for the society. The pragmatic market mechanism described by von Hayek is a much more modest (and believable) piece of equipment that strives for a measure of procedural rationality by tailoring decision-making tasks to computational capabilities and localized information. It makes no promises of optimization.

It is no part of my purpose here to evaluate the social utility of markets in either of these two interpretations. Clearly, it is the pragmatic version that describes the markets of the real world. Rules of substantive rationality that are not backed by executable algorithms are a worthless currency.

Uncertainty and Expectations

However great the simplifications introduced by satisficing at the individual level and by markets at the social, we should not suppose that these devices make choice a trivial intellectual exercise. Some of the largest difficulties remain, including the uncertainty of external events and of mutual expectations.

Uncertainty. Since the consequences of many actions extend well into the future, correct prediction is essential for objectively rational choice. Some things we need to know about the future relate to the natural environment: the weather that will affect next year's harvest. Some relate to social and political environments beyond the economic: an Iranian revolution. Some relate to the behaviors of other economic actors -- customers, competitors, suppliers -- which may be influenced, in turn, by our own behaviors.

In simple cases, uncertainty arising from exogenous events can be handled by estimating the probabilities of these events, as insurance companies do -- but usually at a severe cost in terms of computational complexity and information requirements. An alternative or supplementary measure is to use feedback to correct for unexpected or incorrectly predicted events. Even if the anticipation of events is imperfect and the response to them less than accurate, adaptive systems may remain stable in the face of sizable jolts, their feedback controls continually bringing them back on course after each shock that displaces them. Although uncertainty does not, therefore, make intelligent choice impossible, it places a premium on robust adaptive procedures instead of strategies that work well only when finely tuned to precisely known environments.

Expectations. In general, a system can be steered more accurately if it employs feedforward, based on predictions of the future, in combination with feedback, to correct the errors of the past. However, the formation of expectations to deal with uncertainty creates it own problems. Feedforward in a control system can have unfortunate destabilizing effects, for the attempt of the system to look ahead may cause it to become overreactive and to go into unstable oscillation. Feedforward in a market system can become especially destabilizing when each actor is trying to anticipate the actions of the others (and hence their expectations).

The standard economic example of destabilizing expectations is the speculative bubble. Of course, not all speculation blows bubbles. Under many circumstances, market speculation stabilizes a system causing its fluctuations to become smaller, for the speculator attempts to notice when particular prices are above or below their "normal" or equilibrium levels in order to sell in the former case and buy in the latter. Such actions push the prices closer to equilibrium. Sometimes, however, a rising price creates the expectation that it will go higher yet, hence induces buying rather than selling. There ensues a game of economic "chicken," all the players assuming that they can get out just before the crash occurs. What is destabilizing in this case is the absence of a belief that the actual price will very soon return to its long-run equilibrium. There is general consensus in economics that destabilizing expectations play an important role in monetary hyperinflation and in the business cycle. There is less consensus as to *whose* expectations are the first movers in the chain of reactions, or what to do about it.

But the difficulties raised by mutual expectations do not affect only the monetary system and the business cycle. They also appear wherever markets are not perfectly competitive. The market mechanism described in the previous section was the mechanism of perfect competition. Since each firm took market prices as "givens" that could not be affected by its actions, prices were as much a part of the external environment of decision as were the laws of the physical world. But in the world of imperfectly competitive markets, firms need not make this assumption. If, for example, there are only a few firms in an industry, each may try to outguess its competitors. If only one plays this game, there is no problem; but if it is played by more than one, even the definition of rationality comes into question.

More than a century ago Augustin Cournot[8] undertook to construct a theory of rational choice in imperfect markets involving two firms. He did it by assuming *limited* cleverness. Specifically, he assumed that each firm formed an expectation of the reaction of its competitor to its actions but that each carried the analysis only one move deep. But what if one of the firms, or both, try to take into account the reactions to the reactions? If we follow Cournot's path, we may be led into an infinite regress of outguessing.

A major step toward formulating this problem correctly was taken in 1944 -- a century after Cournot -- when von Neumann and Morgenstern published *The Theory of Games and Economic Behavior*.[9] But far from solving the problem, the theory of games demonstrated how intractible the task is to prescribe rational action in a multi-person situation where interests are opposed. The difficulty of arriving at such a definition exhibits itself in starkest form in the so-called Prisoners' Dilemma game.[10] In the Prisoners' Dilemma, each player has a choice between two moves one cooperative and one aggressive. If both choose the cooperative move, both receive a moderate reward. If one chooses the cooperative move but the other the aggressive move, the aggressor increases his reward significantly but the cooperator is punished severely. If both choose the aggressive move, both incur punishments somewhat less severe than in the second case. There is no dominant strategy. Each player will gain from cooperation provided that his partner does not aggress, but each will gain even more from aggression if he can count on his partner to cooperate. Treachery pays, unless it is met with aggression. The mutually beneficial strategies are unstable.

Are matters improved by playing the game repetitively? Even in this case, cleverly timed treachery pays off. Interestingly enough, Roy Radner has recently shown (personal communication) that if players are striving for a *satisfactory* rather than an *optimal* payoff, the cooperative solution may be stable even for finite repetition of the game. Insofar as this result can be generalized, bounded rationality appears to produce better outcomes than unbounded

---

[8] *Researches into the Mathematical Principles of the Theory of Wealth* (New York: Augustus M. Kelley, 1960), first published in 1838.

[9] Princeton: Princeton University Press, 1944.

[10] R. D. Luce and H. Raiffa, *Games and Decisions* (New York: Wiley, 1957), pp. 94-102.

rationality in this kind of competitive situation.

The Prisoners' Dilemma game, which has obvious applications to international relations and the policy of massive deterrence, is only one of an unlimited number of games that illustrate the paradoxes of rationality in situations where goals conflict totally or partially. Classical economic theory avoided these paradoxes by selecting for primary analysis the two situations (monopoly and perfect competition) where mutual expectations play no role. What appears to make market institutions workable (though not optimal) well beyond that range of situations is that human abilities to compute possible scenarios of complex interaction are sufficiently limited to avoid the infinite regress of mutual outguessing.

Rational expectations. Quite recently, the view has become popular in economics that the problem of mutual outguessing should be solved by assuming that economic actors form their expectations "rationally".[11] This is interpreted to mean that they know the laws that govern the economic system, and that their predictions of the future position of the system are unbiased estimates of the actual equilibrium. These assumptions have very strong consequences, ruling out, for example, most possibilities that speculative behavior will be destabilizing.

The difficulty with the assumptions underlying the rational expectations hypothesis is that, although they are empirical assumptions, almost no empirical evidence supports them. Indeed, our knowledge of the very narrow limits of human rationality must dispose us to doubt that business firms, investors, or consumers possess either the knowledge or the computational ability that would be required to carry out the rational expectations strategy.

In sum, our present understanding of the dynamics of real economic systems is imperfect. We know that, in principle, the ability of human beings to form expectations about future events and about the behavior of their fellows is a potential source of instability for economic equilibria. Since we have little empirical knowledge about how people do, in fact,

---

[11] The idea and the phrase "rational expectations" originated with J.F. Muth, "Rational Expectations and the Theory of Price Movements," *Econometrica*, *29*: 315-335 (1961). The notion was picked up, developed, and applied systematically to macroeconomics by R.E. Lucas, Jr., E.C. Prescott, T. Sargent and others.

form expectations about the future, it is hard to choose at present among the models that are currently proposed by competing economic theories to account for the cyclical behavior of the economy.

Markets and Organizations

Markets are only one of the mechanisms that people use to achieve rational behavior above the individual level. The chief competitor of the market for this purpose is the hierarchic organization. In a society like ours, organizations are, with one exception, the smaller raisins embedded in the larger cake of the market structure. The exception is the central government, which is a very large raisin indeed. In socialist countries the government is commonly conceived to be *the* central organizer, without which rational, coordinated behavior would be impossible. We have already seen the fallacy of this view. In democracies, the central government's role is mixed, but it is thought to be more rule-giver and umpire -- a definer of the game -- than a central planner.

Decentralization in Organizations. Putting governments aside for the moment, let us focus our attention on business firms as examples of what formal organizations are like. In the economics literature, organizations are often treated as the antithesis of markets: as highly centralized structures in which all the important decisions are made at the center. This is a misunderstanding. Organizations operating in a highly centralized way would again exceed the limits of procedural rationality and would lack many of the advantages that are attainable from the use of hierarchic authority. Real-world organizations behave quite differently.[12]

In an organizational context, a person exercises authority when he provides some of the premises for the decisions (and hence the behavior) of one or more other persons and enforces the use of those premises. Since a single decision may be influenced by a large number of factual premises and criteria of choice, the fact that some fraction of these premises are specified by superiors does not imply complete centralization. In fact, the main

[12] J. G. March and H. A. Simon, *Organizations* (New York: Wiley, 1958).

advantage to be gained from hierarchic authority is identical with that gained from using prices as communicators: matters of fact can be determined at the particular loci in an organization that are best equipped by skill and information to determine them, and they can then be communicated to "collecting points," where all the facts relevant to a specific issue can be put together and a decision reached. Only a small part of the source knowledge and information and expertise need be present at the collecting points, and these points can themselves be numerous and dispersed through the organization.

Thus, business organizations, like market economies, are vast, distributed computers whose final choice processes are substantially decentralized. If we look closely at the top levels of some of our large corporations, particularly corporations that are subdivided into specialized product groups, we find these top levels largely preoccupied with a few functions: (1) an "investment banking" function of allocating funds for capital projects to the divisions; (2) the selection of executive personnel for the divisions and evaluation of their performance, and (3) long-range planning for capital-funds and for possible new activities outside the scope of existing divisions. It is often remarked that the separate divisions of these large corporations operate much in the manner of independent business firms, even to the point of having their own balance sheets and "bottom line," and sometimes having the right to purchase from outside the corporation commodities that could as readily be provided by other corporate divisions.

Without making exaggerated claims for the independence of typical corporate product divisions, we can see that formal hierarchic organization is not at all synonymous with complete centralization of decision-making. Thus, there are at least two distinct mechanisms for distributing computational functions through a social system: the mechanism of markets and the mechanism of hierarchy. The two mechanisms are not fully equivalent in their effects. For example, none of the theorems of optimality in allocation of resources that can be proved for ideal competitive markets can be proved for hierarchic decision-making processes. As a result, hierarchies sometimes resort to internal markets or schemes of "shadow prices" in order to come closer to allocative optimality. For example, large business firms make a limited use of internal pricing to govern transactions among divisions or smaller components.

Similarly, the socialist states have been depending increasingly on pricing schemes to regulate transactions among their economic organizations. Both cases provide evidence for the limits of centralized planning.

Externalities. What is the case for organizations as against markets? Economists would tend to state the matter in terms of *externalities*, but we need to pay some attention to uncertainty as well. The question of externality arises because the price mechanism works as advertised only when all of the inputs and outputs of an activity are subject to market pricing. To take a traditional example, if a factory is allowed to spew smoke from its stacks without compensating the surrounding homeowners, then the price mechanism will not secure a level of manufacturing activity that is optimal from a social standpoint -- the product will be priced too low, hence will be overused.

The economist's preferred remedy for externalities is to bring the unaccounted-for consequence within the calculus of the price system -- tax the emission of smoke, for example. Since there is no obvious market for smoke (or rather, for its absence), this raises the question of how the price is to be set. Although there are possible answers to this question in the techniques of cost-benefit analysis, they are administrative answers and not answers given by an automatic market mechanism.

Similar questions of externalities of divisional operations make large corporations less than fully willing to allow transactions among their component divisions and departments to be governed wholly by internal prices. In the absence of perfect competition, it is not clear how to determine the correct prices when there exists no independent market to set them.

Uncertainty Absorption. Uncertainty creates a second set of problems that sometimes makes it attractive for social systems to use hierarchy rather than markets in making certain decisions. It is not entirely reasonable to allow the production department and the marketing department in the widget company to estimate next year's demand for widgets if the production department is to make the widgets that the marketing department is to sell. It may be more important, and more profitable, for both departments to operate on the *same* estimate -- even if incorrect -- than for each to operate on a more nearly correct, but different, estimate. For facing uncertainty, standardization may be more effective than prediction.

The assertion that uncertainty leads to the substitution of hierarchy for markets seems, in a way, paradoxical, since uncertainty would seem to call for the greatest degree of flexibility, and flexibility would appear to be more attainable with the decentralized decision-making of markets than with the centralized decision-making of hierarchies. But that is a superficial analysis. All depends on the sources of the uncertainty. If what is uncertain is a multitude of facts about conditions in individual markets, then decentralized pricing will appear attractive; if the uncertainty is global, infusing major events that will affect many parts of the organization in the same direction, then it may be advantageous to centralize the making of assumptions about the future and to instruct the decentralized units to use these assumptions in their decisions.

Clearly, this brief discussion does not exhaust the topic of social mechanisms for expanding the limits of rationality. Enough has been said, however, to show that, in a world of bounded rationality, there are several ways to magnify the computing capabilities of individual human beings and to enhance the possibilities of their collective survival and prosperity. With the combined use of markets and administrative hierarchies, the human species has enormously increased its capabilities for specialization and division of work. It would be too much to attribute the vast growth and spread of human populations to such mechanisms alone -- modern medicine and modern technology have had something to do with it too -- but the dominance of our species over the globe today is witness to the augmentation of human reason made possible by these social artifacts.

The Evolutionary Model

One way to create an artifact is to let it spring from the brain of a creator. Another is to let it evolve in response to some kind of selective force. The simplest scheme of evolution is one that depends on two processes, a generator and a test. The task of the generator is to produce variety, new forms that have not existed previously, whereas the task of the test is to cull out the newly generated forms so that only those that are well fitted to the environment will survive. In modern biological Darwinism, genetic mutation is the generator, natural selection the test.

The Alternative Theory of Economic Man

No one supposes that a modern market economy is the product of deliberate design. Surely it evolved from earlier subsistence economies, shaped by myriads of decisions made by hosts of economic actors over hundreds of years. By contrast, the accounts just given of the behavior of individual firms -- both the optimizing and the satisficing versions -- assume that individual actors deliberately select actions appropriate to their goals, within the context of a given market environment. Adaptation comes about through selection *by* rational actors, not through natural selection *of* actors whose behavior happens to be adaptive.

Sometimes when the classical optimizing model of the business firm is attacked as unrealistic and factually wrong, an alternative theory of the firm is put forth which rests on evolutionary premises. It is argued that it does not matter whether people (and business firms) maximize or satisfice, for in a world of competitive markets only those who make decisions *as if* they were maximizing will survive. If this very important argument were correct, we would not have to concern ourselves with whether economic man uses maximizing *processes*, for the outcome would be about the same whether he did or not. We must, therefore, examine carefully whether the evolutionary argument implies optimization.[13]

Our discussion will have implications that extend beyond economics into biology, for the literature of evolutionary biology uses the language and concepts of optimality quite freely, and in recent years it has even borrowed linear programming and some of the other optimizing algorithms of operations research to describe the outcomes of natural selection in biological systems. This is legitimate only if the processes of deliberate rational adaptation would lead reliably to the same system equilibria as would the processes of natural selection.

[13]The idea of an evolutionary rationale for optimization was introduced by A. A. Alchian, "Uncertainty, Evolution, and Economic Theory," *Journal of Political Economy*, *58*:211-22 (1950), and embraced by Milton Friedman in "The Methodology of Positive Economics," Chapter 1 in *Essays in Positive Economics* (Chicago: University of Chicago Press, 1953). The identification of selection with optimization is challenged by S. G. Winter, for example in his "Economic Natural Selection and the Theory of the Firm," *Yale Economic Essays*, *4*:225-272 (1964), and in a number of later papers. Relations between economics and sociobiology are surveyed by J. Hirshleifer in "Economics from a Biological Viewpoint, *Journal of Law and Economics*, *20*:1-52 (1977).

Local and Global Maxima

In the mathematics of optimization, the difference between local maxima and the global maximum is crucial. In the landscape of California, every tiny hill is a local maximum of altitude, but only Mt. Whitney is a global maximum. For many purposes, it makes a difference whether one finds oneself standing on Nob Hill or on Mt. Whitney. Finding a local maximum is usually a rather easy task: simply walk uphill until there is no place to walk. Finding the global maximum, on the other hand, is usually exceedingly difficult unless the terrain has very special properties. Linear programming owes its computational feasibility and its popularity to the fact that its postulates are strong enough to guarantee that a local maximum will be the global maximum.[14] Moreover, the simplex algorithm for solving linear programming problems is a quite efficient procedure for rapid hill-climbing toward this maximum.

There is no reason to suppose the real world (either biological or economic) has the simplicity of a linear programming problem. Especially in a world of independently evolving subsystems (species, business firms), it is quite easy to imagine equilibria in which each subsystem is optimally adapted to the other subsystems it finds around it, but in which the equilibrium is only local. If any of the subsystems were displaced sufficiently far from such an equilibrium, they might then evolve toward an entirely different one.

The Myopia of Evolution

Darwinian evolution is completely myopic. It proceeds by incremental step-by-step improvement from one situation to another. At each step the evolving organism becomes *fitter* (relative to its current environment of other organisms), but there is no reason to suppose that the process is leading it to some global maximum of fitness, much less a maximum for the system as a whole -- if, indeed, we can even speak of fitness in an aggregate sense.

Many mathematical models of the evolutionary process are misleading in this respect,

[14]The space of reachable points defined by the equations and inequalities of a linear programming problem is convex everywhere to the origin, and the criterion function is a set of parallel linear hyperplanes. One need only visualize the two-dimensional case to see that, under these assumptions, a persistent hill climber cannot be trapped in a local maximizing cul-de-sac.

for they hypothesize two or more varieties of creatures compete. Under the simplest assumptions, the fittest variety survives, the others vanish -- unless they can find a different, specialized, niche in which they hold a comparative advantage. The theory can make no claim that all niches are occupied or that the system of niches is invariant -- which it almost surely is not. It can certainly make no claim of optimality for the entire system of niches.

If a system can be in equilibrium only at a global optimum, then (unless it is our task to find the optimum), we may be more or less indifferent as to how it reached that position. We may be satisfied to postulate a mythical process like the one that permits economic man to maximize his profit or his utility, or we may postulate an equally mythical optimization by natural selection. But if we are considering a system that lives in an environment having a multitude of local maxima, we cannot understand the system or predict its behavior unless we know something of the method and history of its evolution. Nor can we judge whether there is any reasonable sense in which such a system can be regarded as "fittest."

This is not just an in-principle objection to confounding hill climbing with optimization. In a system that develops mostly by myopic hill-climbing, it may be difficult, even in moments of far-sightedness, to move from a local maximum that has been achieved to another that is now in view, but across a deep valley. The movement from the English system of measures to the metric system is a case in point. If a society were to start from scratch in selecting a system of measures, and if it were familiar with both systems, it would surely prefer the metric to the English system. On the other hand, if future benefits are discounted at some rate of interest, it might never be economical to switch from the one system to the other once the original system had become thoroughly embedded in the artifacts and technical practices of the society. If it were not for the sizable advantages to the United States of conforming to international standards, it is very doubtful whether a cost-benefit analysis would show the shift to the metric system to be economically justified.

Hence, while it does seem realistic to view an economy composed of business firms or a biological ecosystem as a system that is evolving by virtue of some kind of generator-and-test process, -- that is, some kind of heuristic search, -- one cannot, from this premise, draw the conclusion that the position the system has reached or is likely to reach bears any resemblence to the equilibria that are found in the pure theory of perfect competition. Moreover, there is no fixed environment to which the system is adapting. Each business firm is adapting to an environment of business firms and consumers that is changing and evolving at the same time. Each species in an ecosystem is adapting to an environment of other species that is evolving simultaneously with it. The evolution and future of such systems can only be understood in the light of their histories.

The Mechanisms of Economic Evolution

If we conclude that the adaptation of both the business firm and the biological species to their respective environments are instances of heuristic search -- hence of hill-climbing and local optimization or satisficing, we still have to account for the mechanisms that bring the adaptation about. In contemporary biology the mechanism is located in the genes, which demonstrate their fitness by their success in reproducing themselves. It is less than obvious what component of the business firm is the counterpart of the gene.

Winter and Nelson, who have written thoughtfully on this topic,[15] suggest that business firms accomplish most of their work through standard operating procedures -- algorithms for making daily decisions that become routinized and that are handed down from one generation of executives and employers to the next. The evolutionary generator consists of all the processes that produce innovation and change in these algorithms. The evolutionary test is the consequent profitability and rate of growth of the firm. Profitable firms are fit and grow by the reinvestment of their profits and their attractiveness for new investment.

Winter and Nelson observe that in economic evolution, in contrast to biological evolution, successful algorithms may be borrowed by one firm from another. Thus, the hypothesized system is Lamarkian, both because any new idea can be incorporated in operating procedures as soon as its success is observed, and because successful mutations can be transfered from one firm to another.

[15]See for example their "Forces Generating and Limiting Concentration under Schumpeterian Competition," *Bell Journal of Economics* 9:524-548 (1978) and references cited there.

Economics and Psychology

The major thesis of this chapter has been that to understand the artificial system we call the economy we must understand its inner environment -- in particular, the limits on the capacities of business firms and consumers for gathering information, drawing inferences, and making complex computations. Economics is a theory of human rationality that must be as concerned with procedural rationality -- the ways in which decisions are made -- as with substantive rationality -- the content of those decisions.

At the individual level, this concern with procedure leads us to examine the techniques for satisficing, and these, in turn, lead us to cognitive psychology, topics that will occupy us in the next two chapters. At the level of markets, concern with the limits of rationality causes us to view the price system mainly as an institution that reduces the amount of non-local information the actors must possess in order to make reasonable (i.e., satisficing) decisions. At the level of the economy, it focuses our attention particularly on the ways in which expectations are formed, and on the stabilizing and destabilizing effects of the actors' attempts at outguessing each others' intentions.

The Utility Function

Little has been said in this chapter about economic goals -- the components of the utility function. It is unrealistic to suppose that utility functions are given and fixed, once and for all. New experiences produce new tastes. Some attempts have been made to save the classical theory along this dimension by replacing tangible goods and services as the arguments of the utility function with more basic "wants" -- for example, pleasure from music listened to, rather than number of hours of listening. Thus, Becker and Stigler speak of investing in musical experience in order to increase the pleasure, per unit of time, in listening to music.[16]

It may be doubted whether anything is gained by trying to rescue the traditional view of utility with such heroic measures. If, to continue the example, we do not wish to speak of a

[16] G. J. Stigler and G. S Becker, "De Gustibus non est Disputandum," *American Economic Review*, 67:76-90 (1977).

change in utility function as the result of listening to music, then we must postulate within the human head a production function (itself changeable by experience) that manufactures musical pleasure from musical listening. We have merely relocated "taste" from the utility function to that hypothetical new production function. It would seem more parsimonious simply to regard the utility function as an evolving structure.

Levels of Aspiration

In modern economic theory, the amounts of utility an actor attributes to the choices available to him are to be inferred from his behavior -- by observing which alternatives he prefers to which others. A utility function can be imputed to him if his choices are consistent and transitive. Although there is a considerable body of evidence today that human choices are often *not* consistent in the way that would be required for the existence of a utility function,[17] I wish to disregard that difficulty for the moment, and to raise, instead, a different set of issues -- issues about the relation between utility, as defined in economics, and pleasure or happiness, as a psychologist might view them.

As far as I am aware, no one has constructed in the psychological laboratory a psychophysical scale of happiness, to take its place alongside the scales of perceived light intensity or perceived weight. Some courageous public opinion pollsters have asked respondents about their levels of happiness or job satisfaction, but they have been content to report how many respondents chose each item on the nominal scale, and I do not believe that they have undertaken a formal scaling exercise. It might be an interesting task.

We can speculate about some of the properties such a thermometer of happiness might have. First, unlike the utility function, it would not be limited to positive values. It would have a zero point (of minimal contentment) above which degrees of happiness could be measured, and below, degrees of unhappiness. Second, if periodic readings were taken of people in relatively stable life circumstances, we would only occasionally find temperatures very far

---

[17] See for example D. Kahneman and A. Tversky, "On the Psychology of Prediction," *Psychological Review*, *80*:237-51 (1973), and H. Kunreuther et al., *Disaster Insurance Protection* (New York: McGraw-Hill, 1978).

from zero, in either direction; and the divergent measurements would tend to regress over time back toward the zero mark. Many, if not most, people would consistently register either slightly below zero (mildly discontented) or a little above (just satisfied).

To deal with phenomena like these, psychology has introduced the concept of *aspiration level*. Aspirations have many dimensions: one can have aspirations for pleasant work, love, good food, or travel. On each dimension, expectations of the attainable determine an aspiration level with which the current level of experience is comparied. If experience exceeds the aspiration level, satisfaction is recorded as positive; if the aspiration level is higher, dissatisfaction. There is no simple mechanism for comparison *between* dimensions. In general, a large increment along one dimension is required to compensate for a small decrement along another -- hence the system's satisfactions are history-dependent, and decisions involving compensatory offsets are difficult to make.

In another place,[18], I have argued that aspiration level mechanisms are precisely the kinds of choice mechanisms we would expect evolution to have produced in an organism that must live in an uncertain and fluctuating environment.

Man in Society

Economics is the science that celebrates human rationality in all the ways it exhibits itself in human behavior and the operation of human societies. It has been unfairly labeled a "gloomy" science, largely because, in its Ricardian form, it did not hold out much hope for human progress. The label is unfair. Economics in fact draws a romantic, almost an heroic, picture of the human mind. Classical economics depicts mankind, individually and collectively, as solving immensely complex problems of optimizing the allocation of resources. The artifice of economic man enables him to make the very best adaptation, in the environment in which he finds himself, to his wants and needs. In this chapter, while keeping the adaptive capabilities of mind in the center of things, I have tried to suggest a more

[18] H. A. Simon, "Rational Choice and the Structure of the Environment," *Psychological Review*, 63:129-38 (1956), reprinted in *Models of Thought* (New Haven: Yale University Press, 1979), Chapter 1.2.

complex state of affairs. A verdical picture of economic man and of economic institutions must incorporate in the account of his strategies of adaptation an account of the limits -- information-processing limits -- set by his inner environment.

The picture must also accommodate both the conscious rationality of economic actors and the unplanned but adaptive evolutionary processes that have molded economic institutions.

Operations research and artificial intelligence have introduced new techniques that enhance the procedural rationality of economic actors -- help them make better decisions. On a large scale, markets and hierarchic organizations are social schemes that facilitate coordinated behavior, at the same time conserving the critical scarce resource, human ability to handle complexity and great masses of information. In this chapter, I have not tried to evaluate these forms of individual and social organization (nor the equally important bargaining and voting procedures) but simply to describe them as commonly used solutions to a central human problem.

The analysis shows that a deeper understanding of the tools of procedural rationality requires a closer examination of how the human mind works, of the limits on human rationality. The next two chapters will describe what has been learned in the past quarter century about human information processing. Chapter 3 will carry the story up to about 1970; Chapter 4 will focus on the past decade.