Herbert A. SIMON
Professeur
Richard King Mellon University
PITTS BURG USA

THE MANY SHAPES OF KNOWLEDGE

Mots-clés : Acquisition de connaissance, attention limitée, brevets et copyrights, connaissance organisationnelle, connaissance en économie.

Key words : Acquisition (of knowledge), Attention (limits on), Patents and copyrights, Organizational Knowledge, Knowledge in economics

When one undertakes to put down on paper thoughts about the economics of knowledge, a baffling array of alternative paths stretches out to the horizon. In any human society, even the most primitive one, knowledge is everywhere and is essential to everything and everyone. A native hunter in New Guinea is as truly an expert system as is a medical diagnosis program for a computer. If you doubt that, try to survive in the hunter's jungle with only his artifacts, but lacking his knowledge and skills.

Knowledge is linked with economics in a myriad of ways. Economics enters into the creation of knowledge : the rewards that motivate discovery, the processes for discovering new knowledge, the problems of uncertainty in the acquisition of knowledge, and the good and bad consequences that new knowledge can bring about.

Not only is economics concerned with the costs of acquiring knowledge, but equally with the costs of storing it, of retrieving it when it becomes relevant, and the opportunity costs of attending to one piece of knowledge at the sacrifice of ignoring others.

Then there are the issues that surround property in knowledge : rewards for knowledge creation by vesting property rights in the creator, negative externalities produced by privatizing a cheaply reproducible good, costs of maintaining secrecy or enforcing property rights.

Finally (although I am sure we could find many other items to add to our list), there are important economic implications of the ways in which societies and organizations in these societies allocate knowledge among their members, arrange for the retrieval and application of the knowledge wherever relevant, and provide for the preservation of the knowledge with the turnover of members through the life cycle or mobility.

All of these aspects of knowledge – its creation, its storage and retrieval, its treatment as property, its role in the functioning of societies and organizations – can be (and have been) analyzed with the tools of economics. Knowledge has a price and a cost of production ; there are markets for knowledge, with their supply and demand curves, and marginal rates of substitution between one form of knowledge and another. Externalities create problems for us in reaching Pareto optima in the production and distribution of knowledge. Knowledge is simply one among the many commodities in which our economy trades, albeit one of large and rapidly growing importance. It requires special treatment only because of its special properties, several of which I have alluded to above.

But there is yet another, perhaps even more intimate, way in which economics and knowledge are linked. At the core of economic science lie human decisions : to buy, to sell, to make, to save. It is an important task of the science to discover the goals at which these decisions are aimed, and the choices that will be more or less successful in reaching the goals as a function of the circumstances under which the decisions take place. It is a central problem for economic science to discover the ways in which human actors go about making these decisions, and the limits on their ability to make them well. Viewed in this way, knowledge is not just something to be studied with the aid of the tools of economics ; knowledge itself is at the heart of these tools. We need a theory of knowledge, and of the ways in which human beings process knowledge as the foundation for our theory of economics – nothing less than a rather comprehensive theory of human cognition, motivation, and social interaction.

A treatment of any one of the five topics I have listed could easily fill a volume, and within the pages of a chapter there is only room for brief, and perhaps rather superficial, comments on them. So on the first four topics, I will limit myself to some observations about properties of knowledge that make it a rather special commodity, calling for its own treatment - surely a chapter or more – in any contemporary or future treatise on economics. To make up for brevity, I will try to be generous with references to writings where these issues are treated more extensively.

Then I will consider a little more fully the fifth topic : how a theory of knowledge can provide a suitable basis for an empirically sound economic science. In this context, by a theory of knowledge I do not mean an epistemology, in the classical philosophical sense, but a psychological theory of how information about the world is created, processed, communicated and used in solving problems and making decisions in general, and in making economic decisions in particular. Although I will present my two themes – the economics of knowledge and the knowledge foundations of economic science – successively rather than simultaneously, I hope that before I am done the reader will be able to hear them as the two voices of a fugue.

Creating knowledge

Motives for creating knowledge. There are two obvious reasons why human beings might be motivated to seek new knowledge : for its value in use, and for reasons of curiosity. Of course curiosity is not always « idle ». From an evolutionary viewpoint, curiosity could be expected to arise in any creature with wits enough to store knowledge ; for knowledge enormously enhances an organism's abilities to cope with its environment. As we can see from the epidemic of industrialized Homo sapiens that has spread over our globe today, stressing its ecosystem severely, knowledge acquisition is perhaps the most powerful of all fitness-enhancers, more powerful, at least for multi-celled organisms, than the slow processes of mutation and genetic crossover. That is why our species is today a threat to almost all others, except possibly the bacteria and viruses, the speed of whose reproductive cycle and mutation-plus-crossover rate competes favorably with the speed of our knowledge acquisition, as evidenced, for example, by the rapid adaptation of bacteria to new antibiotics.

The evolutionary force of knowledge acquisition depends not only on the speed and costs of creating new knowledge, but also on the rate at which knowledge can be transferred from one person or computer to another. In the case of humans, the transfer (learning) rates are slow enough to keep most of us in school for a sizable fraction of our lives, and are costly for families and society. In the case of the computer, knowledge, once created, can be transferred from one computer to another almost instantly and costlessly – at least in principle. As we are all painfully aware, the chaos of software, which we have not yet learned to reduce to order, puts serious limits on such transfer today. If we think that enhancement of human fitness is an important goal, then lowering the costs of transferring knowledge, between computers, between people, and between computers and people, must be activities with very high priority.

Processes for creating knowledge. As the processes for creating knowledge have been the object of much study in recent years (see, e.g., Langley, et al., 1987), I will not try to survey that topic here, beyond making a few brief comments. Scientific and technical *discovery* has been shown, by comparing historical and laboratory human data with the traces of simulations by appropriately programmed computers, to use the same basic processes that are central to all human problem solving. Two central processes are rather well understood : *recognition* of familiar patterns in problem situations, and highly *selective search* along alternative problem-solving paths. Both recognition and search are based on knowledge. Recognition is based on knowledge of significant patterns in a domain of expertise, together with knowledge, associated with those patterns in memory, of information about their significance, and how to act on them. Effective search is based on knowledge that enhances selectivity, distinguishing promising paths from fruitless ones with at least moderate reliability.

In addition to processes of recognition and search, we also understand a great deal about *learning* processes for building up knowledge bases for use in recognition and search. Theories in the form of computer models that learn to discriminate among large collections of patterns (in some cases, more than a quarter million of them), have shown how human experts can acquire the large indexes they need in memory to recognize situations in the domain of expertise in order to evoke the information they have acquired for dealing with these situations. Other theories, also in the form of computer models, have shown how the human information processing system can acquire, by examining worked-out examples, new problem solving skills in the form of productions (IF-THEN rules) that connect situations to appropriate actions. The programs for acquiring new discriminations and new productions are not blue-sky speculations, but are grounded in a large body of experimental data about human performance in real-life learning tasks (for example, learning to solve algebra equations and to prove geometry theorems, or learning a new language by simultaneous exposure to pictures and sentences describing them). Introductions to these theories and the evidence supporting them can be found in Simon (1996) and Richman et al., (1996).

Two other processes for creating knowledge have been explored, but less thoroughly than those just described. To work on a problem, the problem must be represented in some particular way, and the way it is represented may have a large effect on ease of solution. Most of the problems that we call « insight problems », and whose solution often follows a sudden « aha ! », are solved, not by searching directly for a solution, but by searching for a *problem representation* that will make the solution obvious, or nearly so. Thus, by inventing a new representation, the calculus, Newton turned the theory of gravitation from an intractable problem to one solvable by college students.

For scientific problem solving, *instrumentation* is a fourth key variable. A new instrument that allows us to see things that we previously had to guess about can make enormous contributions to progress. The microscope is an early modern example, chromatography a more recent one, both essential to the revolutions that have occurred in biology. The Geiger counter and cloud and bubble chambers are powerful means in physics for making the invisible visible. In the human sciences, we badly need instruments that can increase the temporal resolution of our observations of thought processes down at least to the level of tens of milliseconds. Developments over the past generation in the use of eye-movement records and verbal think-aloud protocols have brought us down to perhaps the half-second level, but an important gap remains.

Uncertainty. Uncertainty lies at the core of knowledge creation, and I will have a good deal more to say about uncertainty in the final section of this chapter. New knowledge means knowledge that we did not have before, nor do we search for it with any assurance that we will find it, or any notion of when it will appear. Socrates, in the *Meno*, even puzzled over how, when we found an answer to a question, we could recognize it as the answer – how could we do that without having already known it ? This particular misgiving of Socrates proves to be unfounded, but explaining why would take us away from our main path (Simon, 1976).

Beyond the question of whether we will find new knowledge is the question of how we should assess its value : what we might call the Prometheus/Pandora question : will the knowledge we find present a boon to humankind or a bundle of mischiefs, or both ? Scientists working in an applied field have some basis for approaching that question. They can ask, for example, whether humankind would be better off or worse off with knowledge of how to build a nuclear weapon. (Even the answer to this question may be complex, for the question could be stated slightly differently : « Would humankind be better off or worse off if the American and British learned how to build a nuclear weapon before Hitler's Germany did ? »).

But the scientist doing basic research is in a much poorer position than the applied scientist to deal with the Prometheus/Pandora question. When Einstein discovered, as a derivation from his theory of special relativity, that $E = mc^2$, the physical basis for nuclear energy, what consequences should he have anticipated, and what value should he have placed on the discovery ? We can ask the same question about Babbage, or Turing, or von Neumann and the computer. Some basic researchers would argue (and have argued) that the ineradicable uncertainty about the social effects of basic discoveries exempts them from the obligation to concern themselves with the potential human consequences of their research. If they do concern themselves with it, they should not expect thanks when their predictions run counter to public beliefs and fears. Allen Newell and I have been loudly criticized and even ridiculed over the years for predictions about the technical future and social implications of computers that we made in 1957 and that have turned out to be surprisingly accurate (Simon and Newell, 1958).

To the scientist's obligations to assess the social consequences of new knowledge, I would propose a different response than unconcern. I suppose that my response could be dismissed as a rationalization or subterfuge, but let me state it anyway. Although we cannot predict the consequences of new knowledge, we can look back in history and ask the question : On average and on the whole, has the creation of knowledge benefited humankind or harmed it ? We could even take a less anthropocentric view and ask whether the creation of knowledge has benefited the world's ecosystem or harmed it. My own answer (an anthropocentric one) is that I prefer having lived in the 20th Century to the prospect (or retrospect) of living in the 18th, the 12th or the 3rd. Readers may or may not have the same answer. But whichever answer we give, we cannot avoid recognizing the question as the central question of the economics of creating new knowledge.

In both basic and applied research, if it is often virtually impossible to arrive at estimates of the value of prospective discoveries, we can sometimes do better in estimating the values and potential costs of discoveries that have just been made. For example, just after the first atomic bombs had been exploded, and when only the scantiest unclassified information about the technology was available, it nevertheless proved possible to make reasonable estimates – setting lower and upper bounds – of the economic feasibility of nuclear power and the economic benefits that might be obtained from it (Schurr and Marschak, 1950). These estimates were a great deal more sober, and realistic as it turned out, than common predictions at that time of « free energy ». Had they paid more attention to the estimates, several large corporations, including General Atomics, could have avoided losses of some billions of dollars.

THE COSTS OF HAVING USABLE KNOWLEDGE

In costing out a knowledge system, we need to place a price tag on the acquisition of knowledge, and its retention and storage, on its retrieval when it becomes relevant and is wanted, and on the opportunity costs of attending to some knowledge while we ignore others. (On the topics of this section see Simon (1997a) and the references cited there.)

Acquisition costs. I have already discussed acquisition costs, at least for the case of discovery, and have pointed out that costs of transferring knowledge from one memory (human or computer) to another are as important as the costs of obtaining the knowledge in the first place. One way of putting this matter into perspective is to compare total expenditures for research and development (approximately 2.7 % of GDP for the United States) with total expenditures for education (7.2 % for the U. S.). Thus, we see that nearly three times as much is spent for the transfer of knowledge as for its creation. These numbers, especially that for knowledge transfer, are substantial underestimates, for they do not take into account the « labor costs » of the school population (about 10 % of gross domestic product, calculated at the minimum wage), or of on-the-job learning and updating of knowledge, which must account for a large fraction of the human effort that goes into knowledge acquisition.

Retention costs. Retention costs, for individuals, are the costs of renewing knowledge as it gradually deteriorates over time. For organizations and other groups of individuals, including societies, retention is closely related to the movement of people. Knowledge must be provided to new employees as they move into organizations and to new occupants of particular organizational locations ; and the loss of institutional knowledge through death, retirement and resignation imposes an enormous cost of reproduction on most organizations. Who bears the cost of acquiring this knowledge depends, of course, on the extent to which it is specific to particular jobs and companies, or, on the contrary, transferable from one job to another.

When an organization tries to adopt a new technology, or to maintain a body of knowledge and doctrine that differs from that of other organizations of its kind in the society, then it must assume additional costs, which can be quite substantial. In the former case, they are the costs of instructing its members in the new technology (or replacing employees with others who have already learned it) ; in the latter case, they are the costs of indoctrinating new members, who bring with them into the organization the « foreign » exterior culture which will gradually destroy the local deviant culture unless measures are taken to preserve the latter.

Established firms in the electrical industry faced the former problem when they had to convert their products, over a relatively short span of years, from electrical to electronic form. Business schools that wished to introduce such new developments as operations research and behavioral theories of the firm in the period just after WWII faced the latter problem, because traditional business schools and economics and psychology departments, initially largely innocent of these developments, were their main sources of new faculty.

Retrieval costs. The spectacular growth of the World Wide Web has attracted the attention of all of us to the fact that a store of knowledge is of no use unless the pieces of information that are relevant to any specific problem can be retrieved rapidly and cheaply when a person or organization is dealing with that problem. As we shall see in a moment, *not* retrieving information that is *not* relevant is as important as retrieving that which *is* relevant.

Retrieving the relevant information is not a new problem ; organizations face it all the time : finding the expert in the organization who has critical information to contribute to a particular problem, or even more difficult, finding the appropriate external information source. Sometimes retrieval is so difficult that it is cheaper to re-create the new information than to retrieve it : for example, to run a new experiment rather than to search the literature.

Filtering out irrelevant information is also not a new problem, but its urgency is multiplied many times as we create ever larger and more numerous data banks, usually accessible by computer. The larger the data banks, the less efficient are brute force searches, whether by people or computers, and the more powerful are the filters that we require for screening out irrelevancies if the search is to be worth while. That brings us to the question of limits on human ability to attend to information.

Limits on attention : The opportunity cost of looking at the wrong thing. In a world in which information is scarce, we design information systems to ferret out as much information as they can : more is better. In our present world, which is flooded with information (if not always with the information we want), information is simply not the scarce factor – the factor that is to be hoarded in the design of effective information systems. A person who can read at the rate of 30 pages an hour - probably an overestimate for serious technical material – and who spends half of his or her working hours in reading, can read in a year about 30,000 pages, which amount to about 75 books of 400 pages each.

Seventy-five books may seem to be quite a lot, but that is less than one quarter of one per cent of the number that are published in the United States in a year – to say nothing of professional and trade journals, newspapers, office reports and memoranda, and correspondence – including the ever-burgeoning e-mail. Moreover, hardly anyone spends half of his or her working hours in reading ; oral conversation usually takes up much more time, and writing (and perhaps even thinking) a significant amount.

The consequence is that if you are reading one book (or report or letter) there are thousands of others that you are not reading. The information retrieval problem is simply not a problem of retrieving more information - you are already exposed to as much as your time and the limits of your attention will admit. It is a problem of filtering out less relevant information so that you can use that scarce attention to deal with more relevant information. The cost of attending to information is an opportunity cost : you attend to it at the price of neglecting something else.

The problem of information retrieval, then, is a problem of attention management and information filtering, with *relative* relevance as the criterion. But to decide that a particular item of information lies above the critical level of relevance requires intelligence, whether human or computerized. But to use human intelligence for filtering information simply means, again, to reduce the amount of time available for attending to the information that is judged relevant. Hence, one of the key research problems for computer science over the next generation is to enhance the intelligence of computer programs that can perform the filtering function.

PROPERTY IN KNOWLEDGE

Anything of value can, in principle, become property, with rights of exclusive use vested in owners. Monopolies - that is to say, property rights - in information, and in the devices and books that are its products, have been granted by governments at least since early modern times, and patents and copyrights are their contemporary descendants. The existence of such property rights raises the question of negative externalities through too restricted use of the new knowledge. As the cost of transferring new knowledge is usually far less than the cost of discovering it, it would be advantageous to society if innovative discoveries were available to everyone and not restricted by the property rights attached to them.

Patents and copyrights as economic rewards. The argument for granting these rights, of course, is that such grants constitute a mechanism for motivating discovery and (ultimately) public access to it by rewarding it. It is far from clear, however, that there is any simple way to regulate the amount of reward and the amount of monopoly in such a way as to optimize the social return (or even to assure a satisfactory return). It has been forcefully argued that, for this reason, there will always be underinvestment in discovery if it is left to the private sector and rewarded by patents and copyrights ; and that, therefore, governments should subsidize research and require public access to its products. The argument has special force in the case of basic research, for the rewards from patents are much more likely to derive from applications than from the basic discoveries on which the applications are based. In fact, this argument was used successfully to maintain and even increase the level of Federal support for basic research during the ideologically conservative and market-oriented Reagan administration.

Enforcement of property rights in knowledge. Unless use is open to public inspection, it is extremely difficult to enforce the property rights defined by patents and copyrights. For example, in the American paper pulp industry it was (and perhaps still is) very common for companies that made innovations in their manufacturing processes to keep these innovations secret rather than patenting them. In this industry it is easy to keep outsiders from gaining information about factory processes, hence both easy to enforce secrecy and difficult or impossible to enforce patent rights because of the difficulty of detecting infringements. Book copyrights, on the other hand, have been both highly desirable and relatively easy to enforce, for there is no profit in producing a book unless it can be distributed publicly – which bars secrecy for either copyright owner or prospective infringer.

The scope of patents and copyrights. The owner of a piece of real estate has merely to build a fence around it in order to establish the scope and limits of the property. Matters are much less simple in defining the boundaries of coverage of patents and copyrights. In patents, what constitutes a « copy », hence an infringement, depends on an ill-defined notion of two devices or processes being based on the same principles, and that notion is the subject of constant litigation. In copyright law, the ambiguity is much the same, but the criterion is different. Here, it is not the idea that is protected, but « the form of words ». Clearly, there may be infringement without word-for-word copying, but again, the boundary has to be determined on a case-by-case basis.

Further ambiguity, especially in copyright law, is associated with the distinction between « legitimate use » and infringement. A scholar is not forbidden to hand-copy materials from publications, and within certain limits, to quote them in other publications. As soon as high quality and efficient mechanical copying machinery became available it became economically attractive under many circumstances to reproduce numerous copies of materials – for distribution to college classes, for example. But this clearly reduced the value of the property rights that were supposed to be protected. The computer has created even greater difficulties, for material stored in a computer memory can be copied at different times by many users in single copies – which can aggregate to totals of any size.

Today, copyright law, and to a lesser degree patent law, is in a state of considerable confusion, and it is no longer clear (if it ever was) to what extent awarding property rights to innovators of devices, processes, or written texts is an effective means for rewarding innovation without an unacceptable amount of restriction on information transmission. On the other hand, in the absence of a reasonable system of rewards to inventors and writers, we are in real danger of drying up the motives for producing innovations that would be valuable to our society.

Public investment in education and research. I have already stated the case for public funding of research as an antidote to underinvestment by the private sector because property rights in knowledge cannot usually capture its full marginal value in the market. The same case can be made for public funding of education. It is almost surely the case that much of what we learn, we learn by interaction with other people who already have the knowledge. If the people with whom we interact have lots of knowledge, we are likely to learn much from them ; if the level of their knowledge is low, we are likely to learn far less. It is very difficult, perhaps impossible, for the educated person not to behave altruistically - that is, to contribute knowledge to others (but without losing it, for there is no conservation law for knowledge).

I have frequently heard the argument for public funding of research ; I have seldom heard this particular argument applied to public funding for education. Perhaps it is too obvious to need saying. Perhaps one would hear it more often in a developing country where it is very obvious how much each person would gain if the others around him or her had more knowledge. The argument appears not to be hard to understand, for most people seem to be aware of the advantages of attending a college where the other students, as well as the faculty, are bright and well informed, and therefore capable of contributing to one's education.

THE KNOWLEDGE OF SOCIETIES AND ORGANIZATIONS

Knowledge, of course, is held in individual human heads, and individual books, and individual computer memories. Often, however, it is useful to view knowledge in more aggregative form : as held by the collectivities of human beings called societies or organizations, or by libraries, or by data banks and computer networks, and to view it as transmitted among all of these by networks of conversation, correspondence, telephone, publication, and web surfing. Then we can look both at the economic costs and the economic values of organizing knowledge in these ways.

Knowledge in developing societies. In the enthusiasm, shortly after WWII, for aiding third world countries to increase their productivity, economic development was at first closely identified with industrialization. Third world countries, it was argued, were plagued by hidden unemployment in the form of excess rural populations larger than needed to carry on agriculture efficiently on the land available. Moved to the cities, and provided with an adequate supply of industrial capital, this underemployed population could, it was claimed, become more productive. The key was to provide the capital required for their employment, and to build the factories and the required infrastructure of roads, railroads, and power networks.

After considerable experience had been gained, and considerable painful disappointment suffered, it was gradually recognized that in a great many developing countries, including the poorest, the « hidden unemployment » in rural areas was largely mythical (and often quickly replaced by total unemployment of migrants to the cities), and that the most feasible route to higher productivity was the Green Revolution of new plant varieties combined with manufactured fertilizers and some mechanization of farm equipment. It was further realized that the form of capital that most needed expansion was the intellectual capital of the population : the knowledge and skills that would enable and facilitate the creation of a modern economy, in both its agricultural and manufacturing sectors. Literacy programs, education abroad for the advanced cadre of « technocrats », strengthening of electronic communications systems all became increasingly important components of the development process.

Finally, in the most advanced of the developing countries, and those that made the most rapid progress, it began to be recognized that the future lay, not in focusing investment to reach self-sufficiency or build export markets in such traditional heavy industries as steel, shipbuilding, and coal mining, but in moving into the knowledge industries themselves, especially electronics and, more recently, computer software. Software is, of course, the quintessential knowledge industry, for its very product is knowledge and methods of transmitting and processing and storing knowledge.

Organizations and knowledge. In classical economics, the canonical organization is the business firm, which is described by its production function : i.e., by the relation between quantities of the factors of production consumed and quantities of product produced. However, the production function is less a description of the organization than of its interface with its market environment. To understand and describe the business firm, or any other organization, it is necessary to understand and describe the decisions that govern the transformation of inputs into outputs, and the knowledge and communications that inform these decisions (Simon 1997b). The invention of the computer and the rapidly expanding use of computers, not only for routine record-keeping purposes but as an integral part of the decision process, has brought about a growing understanding of the organization as a processor of knowledge, although the traditional production-function view still prevails in the economics textbooks.

The new decision-making point of view had its first practical impact through operations research, which provided new formal tools for improving the quality and economic effectiveness of decisions, especially at middle-management levels. The focus of these tools on formal mathematics and optimization delayed their extension, however, to decisions of a more qualitative and « judgmental » character that involved the processing of non-numerical information. With the recent expansion of expert systems based on the methods of artificial intelligence, these barriers to the application of the information sciences throughout the organizations and to decisions at the highest executive levels, are gradually being removed. Today, the respective roles of computers and human employees in decision making, the organization of data banks and of methods for retrieving information from them, the distribution of information within the organization, and the channels of communications among information sources and decision points have become central issues in the design of organizations.

THE ECONOMICS OF KNOWLEDGE TODAY AND TOMORROW

I have only sampled, and that in a not very systematic way, the burgeoning of our awareness, in almost every sphere of human activity, of the central role of knowledge, its creation, preservation, and communication, in economic processes.

We are increasingly concerned, on the one hand, with knowledge as a factor of production, having costs and values that we must try to estimate in order to make correct decisions for the conduct of business and government. We have discovered that applying an economic calculus to knowledge and communication is often, because of the intangibility of the product, far more difficult than costing and valuing wheat, coal, or even production machinery. Yet, whether we can make the measurements accurately or not, it is precisely these costs and values that determine the efficiency and profitability of our activities, and the social consequences that attach to them as well.

We are increasingly concerned, on the other hand, with new sciences of knowledge and communication, originating within the past half century, that provide us with novel and more powerful ways of managing our organizations and the decision processes in them – of carrying out the economic analysis, and the other problem solving required to operate them. I have mentioned operations research as one of these new sciences and artificial intelligence as another.

But it is no longer very useful to identify these sciences as isolated in particular disciplines. In carrying out all of our intellectual activities – whether they be decision-making in law or business, or scientific research, or administering governmental activities, or medical practice, or teaching – there now are tools for analyzing and rationalizing the task itself in terms of its information requirements and in terms of the tools that can help carry out its problem solving and decision making processes. Increasingly, each of these professions is finding it necessary to education its students in the uses of information and communications technologies in its domain.

Tools for dealing with knowledge have become as ubiquitous as paper and pencil. Today there is « computational physics », « computational law », « computational medicine », each with its specialists and professional organizations. The domain of information science, and the economics of information as part of it, has become too big for information scientists ; it is today, and must become to an even greater extent tomorrow, a part of the knowledge and skill that all of us possess and use daily.

THE KNOWLEDGE FOUNDATIONS FOR ECONOMICS

The new attention to information processing in the conduct of business brings us to the other part of our story : the role of knowledge and of its processing by human beings as the basic raw material for the production of an empirically veridical economics. One way of thinking about economics – perhaps the usual way – is in terms of the wheat and steel and baseball games that we produce and exchange. Our textbooks are filled with discussions of the quantities of these commodities that are traded and the prices that are paid for them. They form the core of the intermediate price theory course, the forge in which the minds of professional economists are hammered into shape. And in our theories, the very money that pays for these commodities becomes just another commodity, the numeraire or common denominator for the others.

Economics as a Science of Thought Processes

There is another way of thinking about economics. Economic behavior is a set of processes that take place inside human heads and in the acoustical and electrical waves that are used to communicate between heads. (Electronic banking has disabused us of the fantasy that money consists of tangible pieces of paper.) In this way of thinking, the physical commodities have no independent importance or even significance ; their significance derives from their utilities, also a set of intracranial states and processes within the same human heads. In particular, money is just a state of mind – a set of expectations about the kinds of interpersonal transactions it enables, where these transactions are exchanges of the utility of commodities for the utility of money. To call money « a state of mind » does not devalue it or demonetarize it, but perhaps « de-demonizes » it, giving it a true home (surely it does not live in banks) among the thoughts that reside within the human head.

Almost everything that has happened in economic theory in the past half century points, however gradually and haltingly, to this alternative view of economics. Matters began to move that way when we undertook in a serious way to incorporate uncertainty in our economic theories, for no one, at least since Frank Knight, has described uncertainty as anything other than a state of mind. At first that uncertainty was mainly uncertainty about our future environments, but then, heralded by the 1944 publication of *The Theory of Games and Economic Behavior* (and anticipated a century earlier by Cournot), it became uncertainty in our minds about the workings of other minds – a far more profound uncertainty, and one immensely more difficult to deal with in an organized social system. As events have proved, concepts like « rational expectations », or for that matter, Keynes'« animal spirits », were simply more limited and metaphorical expressions of the same realization.

As one consequence of these developments, we have been left, in a very wide range of game situations, without a generally acceptable, much less a universally agreed-upon, definition of what it means to behave rationally. As a further consequence, we have been left without a macroeconomics, for, depending upon our assumptions about how human minds deal with uncertainty, anything can happen. Without an empirically valid definition of rationality under uncertainty (i.e., a definition consistent with the ways in which human beings in fact behave under uncertainty), the traditional escape route, the assumption of perfect rationality, is no longer open. If money is a state of mind, then inflation is a state of mind – as it surely is. So, for that matter, is unemployment, which does not necessarily make it either « natural » or « voluntary ». Because all of these may be, and inevitably become, not just individual states of mind but collective states of mind, they have real effects, often of the greatest magnitude, in the real world.

The way out of what is now widely recognized in the profession as a dilemma is not to deny it but to proceed with the task of discovering how the minds of the economic actors actually work. A great deal of progress in that direction has already been made within the past half century, both by economists and by psychologists, and as a result, we are by no means wholly ignorant of human decision processes, even under uncertainty. As a rapidly growing chorus of voices, including mine, have given accounts on a number of occasions of some of the things we know about these processes, and of how our knowledge differs from the assumptions of neoclassical economic theory, I will not make the case again here. Whether agreeing or disagreeing with the argument, no one any longer has reason to plead ignorance of it. (For a recent restatement of the argument, see Simon, 1997c.) An excellent recent microeconomics textbook representing the modern view is Earl (1995).

Economic Science as a Historical Science

If the laws of economic behavior rest on beliefs residing in human heads, then we should not expect these laws to remain invariant over the centuries during which human beings acquire new knowledge about the world and about each other (Simon, 1998). Economic science is a historical science, along with history, geology and cosmology.

On a very concrete level, new financial instruments have been invented in a steady stream at least since the Renaissance and culminating in the « derivatives » of the present time. It would be very hard to argue that the availability of these instruments did not alter the workings of the economy, in particular its dynamic stability or instability. Similar statements can be made about changes in public institutions (e.g., the Federal Reserve System and the SEC). There is reason to believe that knowledge about the information-free characteristics of effective markets and about uncertainty in multi-person games change the behavior of investors. The invention of modern manufacturing methods and the large corporation produced enormous changes in the operation of industrial economies, and the computer and electronic commercial communications is beginning to produce others.

One might respond to all of this by saying that these developments do not change the basic laws of economics but only the boundary or environmental conditions under which the laws operate. The debate is really about terminology, not about the temporal variability of economic behavior. Perhaps some clarity can be gained by looking at the analogy between economics and geology. Geology is generally viewed as a historical science, with widely different successive Eras. A geologist would not operate very effectively in his science if he did not know what the Paleozoic was, or the Ice Age. Yet the Earth and the Heavens that affect it are governed by (presumably) invariant physical laws, and here again it is the continual change in the state of the system that accounts for the changes in its behavior.

In any real science, however, geology and economics included, we are concerned not only with the most general laws, but also with the special laws that govern system behavior in particular domains of interest to us and under certain ranges of conditions. Continental drift is presumably derivable (at least « in principle ») from the invariant laws of mechanics and thermodynamics. But the theory and history of Continental drift is an important component of geology. The special laws that govern it are as much a part of the science (independently of whether we can derive them with any precision from more general laws), and as important for the understanding of phenomena, as the time-invariant laws. To argue that only the most basic and general regularities are the laws of a science is to argue that as we have quarks we have no need for chemistry. (I have heard some particle physicists argue exactly that, but I don't think many of us are convinced by their argument).

In economics, the boundary between wholly invariant laws (unless we narrow these further by consideration of long-term evolutionary change of our species) and time-dependent laws suffers from the additional complication that a large part of what exists in the human brain and controls human behavior is not determined by its built-in genetics, but consists of knowledge and skills that have been gained through processes of learning and that are different in every generation. The economic actor whose behavior we are striving to model is the actor who possesses some part of that collection of knowledge and skills as were available in his or her lifetime. That knowledge and those skills are major determinants of how the actor will make economic decisions, and thereby, how the economy will behave and develop. We humans are as much afloat on the sea of history as the continents are on the skin of the Earth. Economics is an historical science that must learn to use, more than it does, the methods of historical analysis.

What is to be Done ?

What has been discussed less than the current deficiencies of economics science, and what perhaps deserves mention as I conclude my comments, is the question of what is to be done about them. I have a short agenda to propose, but not one that can be carried out without a large expenditure of the profession's effort. We now return again to questions of the economics of knowledge : in this case how the economics profession is to acquire and transmit the knowledge that is needed to carry out economics research, and to transfer that research into practice, in a world where economics is a state of mind - or rather, of many minds. Bringing that about is going to require a substantial investment.

The first prescription is that economists must learn some modern cognitive psychology, and economists in training must receive training in psychology, not as liberal education, but as an essential prerequisite to the study of economics. They will find in psychology rather rigorously stated theories (often in the form of computer programs, that is to say, systems of difference equations), strongly supported by evidence, that give realistic accounts of human problem solving and decision making processes in complex situations. I should add, however, that the theories'treatments of the various forms of uncertainty is not their greatest strength, although they have a great deal to say about the processes of design, in which uncertainty plays a major role, as well as the processes of discovery.

The second prescription is that economists must extend their tools for making empirical observations. The recent rapid growth of experimental economics is an important first step in that direction, but it provides only a partial solution. An equally important part calls for improving the tools for making real-world observations, including historical observations, of behavior in general and decisions in particular, within industries and individual firms, and for analysing such observations more rigorously than is usually done with « case studies. » It is especially important that economics learn to build evidence from the pooling of collections of cases that include highly qualitative as well as quantitative data. Again, some help can be obtained from psychology, but existing methodologies are far from adequate. (A quite recent excellent example of the application of a range of these methodologies to the theory of business firm size and growth will be found in Sutton, 1998.)

The third prescription is that economics must change its standards of publication so that the currency of economic knowledge will not continue to be debased by floods of hypothetical models whose empirical assumptions are built on little or no solid empirical evidence. In a world of information glut, the broadcasting of unbridled speculation creates enormous costly externalities.

I am sure that I have not exhausted the list of remedies, but liberal application of these three may be enough to occupy the energies of the profession for some time.

REFERENCES

Earl, P. (1995) « Micro-economics for business and marketing ». Cambridge, U.K. : Edward Elgar.

Langley, P., Simon, H. A.; Bradshaw, G. L., & Zytkow, J. M. (1987) « Scientific discovery : Computational explorations of the creative process ». Cambridge MA : The MIT Press.

Richman, H. B., Gobet, F., Staszewski, J. J., & Simon, H. A. (1996) « Perceptual and memory processes in the acquisition of expert performance : The EPAM model ». In K. A. Ericsson (Ed.), « The road to excellence : The acquisition of expert performance in the arts and sciences, sports and games » (pp. 167-188). Mahwah, NJ : Erlbaum.

Schurr, S. H. & Marschak, J. (Eds.) (1950) « Economic aspects of atomic power ». Princeton, NJ : Princeton University Press.

Simon, H. A. (1976) Bradie on Polanyi on the Meno paradox. « Philosophy of Science », 43, 147-150.

Simon, H. A. (1996) « The Sciences of the artificial » (3rd ed.). Cambridge, MA : The MIT Press.

Simon, H. A. (1997a) « The future of information systems ». *Annals of Operations Research*, 71, 3-14.

Simon, H. A. (1997b) « Administrative behavior » (4th ed.). New York, NY : The Free Press.

Simon, H.A. (1997c) « An empirically based microeconomics ». Cambridge, UK : Cambridge University Press. [The 1993 Mattioli Lectures, Universita Commerciale Luigi Bocconi, Milano.]

Simon, H. A. (1998) « Economics as a historical science ». *Theoria*, 13, 241-260.

Simon, H. A. & Newell, A. (1958) « Heuristic problem solving : The next advance in operations research ». *Operations Research*, 6, 1-10.

Sutton, J. (1998) « Technology and market structure ». Cambridge, MA : The MIT Press.

Von Neumann, J. & Morgenstern, O. (1944) « The Theory of games and economic behavior ». Princeton, NJ : Princeton University Press.