

A SPECIAL FEATURE¹

Guest Editors: Yuji Ijiri and Shyam Sunder

Information Technologies and Organizations

Herbert A. Simon

Carnegie Mellon University

I. Information Overload and Organization Hierarchy

Q. What has been the impact of information technologies, especially of artificial intelligence, on organizations. In *The Sciences of the Artificial*, you described systems that are adapted to the goals and purposes imposed by their environment. Have information technologies enabled business and governmental organizations to grow bigger, become centralized, or change in other ways?

A. Information technologies have had a great impact on the way in which information is gathered and disseminated, but their impact on the structure of the organizations has been rather limited. Certainly, there has been a scale change in organizations. Business organizations of the size of General Motors, AT&T, or IBM were inconceivable a hundred years ago. But, was that because people didn't know how to create and run them, or because there wasn't any need for them? These are two different questions.

The political process has seen large changes due to mass communication. Television allows the President of the United States today to speak directly to most citizens in their living rooms. George Washington could not have done that 200 years ago. Faster communication has made it possible for us to create more unified organizations spread

¹ **Editor's Note:** This article is based on an interview with Professor Simon conducted by Yuji Ijiri and Shyam Sunder on July 3, 1989. Professors Ijiri and Sunder, guest editors of *The Accounting Review* for this article, edited the transcript with the approval of Professor Simon.

Herbert A. Simon is the Richard King Mellon University Professor of Computer Science and Psychology at Carnegie Mellon University, where he has taught since 1949. For the past 30 years, he has been studying decision-making and problem-solving processes, using computers to simulate human thinking. He has published over 700 articles and 20 books and monographs that span the disciplines of economics, psychology, computer science, econometrics, management, organizational behavior, accounting, and administrative sciences.

Educated at the University of Chicago (Ph.D., 1943), he has been recognized by honorary degrees from a number of universities. He was elected to the National Academy of Sciences in 1967 and has received awards for his research from the American Psychological Association, the Association for Computing Machinery, the American Political Science Association, the American Economic Association, and the Institute of Electrical and Electronic Engineers. He received the Alfred Nobel Memorial Prize in Economics in 1978 and the National Medal of Science in 1986. He has been Chairman of the Board of Directors of the Social Science Research Council and of the Behavioral Science Division of the National Research Council, and has served as a member of the President's Science Advisory Committee.

over vast distances. When Wellington was sent to Portugal to fight the French, he was in charge. A message might have taken a week or more to get to and from England. Today, he would be in constant communication with London. The Vietnam War was fought from Washington. There never was a unified field command. Thus, the responsibility and authority domains of Westmoreland and Wellington were drastically different.

The effects of this change have been disadvantageous as often as advantageous. Organizations have begun to relearn the lessons of decentralization: divisionalization allows large, multiproduct corporations to reduce the amount of entanglement and coordination. With modern communications, particularly the computer, we have changed the balance between the number of messages that can be produced and that can be received. At first we think of producing more information as a great thing, without thinking of what happens at the receiving end. Gradually, as we get more sophisticated, some of these differences may be reduced.

People have to be trained not to take in information just because it's there. It is a big social training job. Perhaps we ought to start in the first grade and give children practice in turning off the television, even turning it off the minute the news broadcast starts. From time to time I give lectures to executives about why it's dangerous to read the daily newspaper; they must realize that they don't have to pay attention to information just because it is available. No matter how much information becomes available, if you decide not to read it, this growth of information becomes irrelevant.

Why do we follow our everyday information strategies? We must know whether we are processing information because we expect to make a decision or because we enjoy doing so, like getting hooked on reading a serial or a detective story. Once we understand why we do what we do, we would make better decisions about what to read. We would realize that not knowing something is a great way of making our friends feel good telling us about it.

Let me illustrate why humans have to process information selectively. Since I decided to go to China this spring, I paid more attention to Chinese news than I normally do because I was intending to make some decisions on that basis. It turned out I made the wrong one in spite of the fact that I had the best available information Friday, June 2, 1989, directly by phone from Beijing. Alas, it wasn't Saturday's information. I arrived Monday, June 5, right after the Tiananmen Square massacre, though, fortunately, I got out of Beijing two days later.

So you can see where we must switch our priorities. Communication technology has certainly affected our organizations and society by increasing the amount of information explosively. But I don't think it is a large feature. I think the large feature is the limited ability of people to absorb information; the scarce resource is human attention, not information. That puts a tight constraint on how much information you can pour into a system. This was the main reason why management information systems were such a disappointment; they were designed to provide information without ever asking how much information could be absorbed, or what information was relevant.

Once we understand where the binding constraint is, we can understand why information technologies have had so little impact on the structure of organizations. In the *Sciences of the Artificial*, I tried to describe complex systems and explain why they are hierarchical and why we should not think of designing organizations so that everybody speaks to everybody. There are good reasons why we use nesting-block designs of units within units of organizations.

idea is that this theory seems to look almost entirely at economic inducements. This theory also seems to assume that leisure is such a desirable good that people are intrinsically shirkers and that they will only do what can be enforced.

There is a tremendous amount of psychological evidence that contradicts this: human beings are not only capable of acquiring strong loyalties to organizations or organizational units, they are incapable of not acquiring them. We must look at loyalty structures to learn what ties people to organizations. (I prefer to call them “identification structures” because they have a strong cognitive component as well as a motivational component.)

People do all sorts of things for which they receive no reward. They don’t just do the minimum they can do without getting caught. We often find people who are enthused for an organization. I’d feel happier about the agency literature if the people who are writing it understood the centrality of the *identification mechanism*.

This relates to your question about greed, because it suggests that there is a big nongreed component in human behavior. If you were interested in evolutionary theory, you may ask “why don’t the greedy ones outbreed the nongreedy ones?” There is a very good answer to that.

In social species like ours, a person born with nothing but greedy genes would be unable to receive all the benefits that society can bestow in the way of influence and instruction. In contrast, a person born with a few docility genes—that is, with the willingness to accept instruction, influence, information, and persuasion from the social environment—has a great advantage. If you build a mathematical model of human behavior, you will find it very difficult to reconcile fitness with altruism without something like the docility mechanism.

Society can use docility, within limits, to influence the person to be nongreedy. Being loyal to an organization is a form of being nongreedy. In a social animal like the human being, you would expect to see a strong component of docility, by which I mean acceptance of social influence. You would also expect the society to gradually learn how to use this trait to produce people with very strong organizational loyalties. In fact, it takes more than contracts and economic rewards to make an organization work. There is no better way of having a work stoppage in an organization than for everybody to do exactly what their contract calls for, no less and no more. Work-to-rule is a euphemism for work slowdown.

Maybe there is a reason why people other than economists use the word greed for the narrow subset of motives that economists analyze. Noneconomists can see that they and others respond to other motivations. Therefore, they may think that a person is a kind of a monster if he responds only to the motivations that economists talk about.

III. Privacy and Control

Q. Computers have changed the activities of accountants and auditors considerably over the past two decades. What is your assessment of the impact of information technologies on accounting and control? In particular, are we heading toward the state of “overcontrol?” Has there been a change in the functions of controllers? Would artificial intelligence play an important role in accounting in the future?

A. I share a concern about any data bases, computerized or not, recording private information, and how those are managed, who manages them, what controls are placed

if we can get rid of our compulsion to have everything scanned by the human senses. Going beyond just highlighting exceptional items, computers can also do analyses.

However, I don't think that the basic framework of organizations will be altered by the introduction of artificial intelligence. None of the technological developments I've seen really strike at the heart of organizational structure, or require organizations to be turned upside down.

IV. Research and Instructional Technologies

Q. Would information technologies impact research and education critically? If so, in what way; if not, why not? In particular, we wish to have your comments on the availability of computers and databases on research, computer-aided instructions, and on a possible shift in emphasis away from skill training.

A. Information technologies have certainly affected the way in which research is conducted. All sciences are influenced by what instruments are available. Physicists wouldn't be doing the things they do without accelerators. Scientists can, however, ask themselves what kinds of instruments they would like to have. Publicly available sources of data may shift researchers' attention toward macro-level issues that can be addressed using such data. It may further postpone giving detailed attention to areas in economics where I think looking needs to be done—the details of decision-making processes in organizations. Such data are not as readily available. One has to go into organizations and conduct intensive studies.

Availability of public data banks tends to pull people away from facing up to that ultimate necessity of looking closely at decision-making processes. But beyond that, I always prefer theories that are based on data over theories that are not; we have enough of the latter in economics and allied fields.

I don't know what has been done in experimental accounting, but if it is related to experimental economics, it is a very useful development. Of course, you could run it, like any other research method, into the ground. Life is more comfortable for most of us if there is a procedure to tell us what to do next. Every science is subject to the ever present danger of turning into mindless, automatic science—compute a *t*-statistic and a chi-square statistic, and then you're done!

I don't know whether the availability of computers and databases has increased that danger. People who were determined to compute chi-squares before with desk calculators, or by hand, ended up doing it anyway. There is no doubt that computers are used for data dredging. The key question is whether a larger fraction of the scientific community's time is spent on doing this. Some scientists used to spend a large fraction of their day doing it because manual computation took so long. Now they may spend a large fraction of their day doing it because computers make it so easy to try out many different computations. What you have to fight for continually is that at least some people spend some of their time thinking.

Computers certainly changed the way in which research problems are formulated and solved. We can learn things about the behavior of complex systems by computation without closed form analytic solutions. We also know its limits: computation generally yields the solution for a particular set of boundary conditions. If you want to discover more general results, you have to run many different conditions. Therefore, whenever it is feasible, we still prefer to do things analytically.

for examples B, C, D, E, F, G, and H. How do we teach integral calculus now? I assume we spend less time getting people to memorize or guess the integrals of various sets of functions. But we probably still want to sample that. Pick some area in accounting, for example, have students understand it, and spend the rest of the time learning the underlying principles.

I hope someday we'll learn that curriculum building is a process of sampling. We sometimes get so wound up thinking that there are things that have to be covered. There are a million things that have to be covered and there is no way that more than a hundred of them are going to be covered. So we might just as well recognize that we are sampling.

V. Artificial Intelligence and Artificial Ethics

Q. How might the goals of educational institutions change because of technology? The need for teaching technical matters seems to be on a declining trend now that such matters can be better handled by computers. Would all of this change the philosophy of what education is all about? Do you think that there will still be the core of knowledge about human beings, such as ethics and the emotive side of human beings, that must be taught by human instructors? Should educational institutions focus on such matters only, delegating technical instructions to various training machinery?

A. I still think that students are going to be taught by human instructors. Anybody with tenacity can educate himself or herself in almost any subject, but it isn't necessarily efficient. Most realizable artificial intelligence in the near future will be focused on fairly concrete things. Expert systems, for example, do best when there are a zillion specific pieces of knowledge for which they can act as a big recognition memory. That's why legal retrieval and medical diagnosis systems have been effective. Expert systems are still not very good at reasoning about fuzzy things. That's not a permanent limit, but it will guide the division of labor between machines and people for a long time to come.

There are some interesting questions about whether you can teach ethics at all. Let's suppose that you can. A large part of the teaching of ethics consists in taking situations or actions and showing that they have many unanticipated consequences and side effects that people don't immediately think about. Frequently, awareness of such effects is the locus of the ethical component in the decision. An effective course in ethics would make simple choices in life seem more complicated.

The aim of a course in ethics is not to change peoples' values. A course on ethics has to make students think about the multitudinous frames of reference and consequences of decisions before they decide something. Take abortion—a good issue to teach ethics to undergraduates in our country today. Some people say right-to-choose and others say right-to-life, as though either of these slogans can settle the issue. The first step in teaching ethics is to show that there are many ways to look at that issue and most of them are quite complicated. Whether thinking of these complexities leads them to favor one side of the issue or the other is another question.

What should be the goals of courses in ethics? To say that the goal is to produce more ethical behavior is to beg the question. By their nature, ethical problems involve conflicts of values. Unless you know what the right values are, how can you judge whether people become more ethical after taking the course? Instead, we may judge an

way to state it, since we know that the real problems in the world today are ourselves. I guess we are going to solve that problem only if we know ourselves better. If we are that voracious species on the earth that is stamping out everything else, we better understand that species.

A Selected List of Books of Herbert A. Simon

- Administrative Behavior* (3d ed.). 1976. New York: Free Press.
- Centralization vs. Decentralization in Organizing the Controller's Department*. With C. Kozmetsky, H. Guetzkow, G. Tyndall, 1976. Republished by Birmingham, MI: Scholars Book Co.
- Organizations*. With James G. March, 1958. New York: Wiley.
- Human Problem Solving*. With Allen Newell, 1972. Englewood Cliffs, NJ: Prentice-Hall.
- The New Science of Management Decision* (rev. ed.). 1977. Englewood Cliffs, NJ: Prentice-Hall.
- Models of Discovery: And Other Topics in the Methods of Science*. 1977. Boston, MA: D. Reidel Publishing.
- Skew Distributions and Sizes of Business Firms*. With Yuji Ijiri, 1977. Amsterdam: North Holland Publishing.
- Models of Thought*. 1979. New Haven, CT: Yale University Press.
- The Sciences of the Artificial* (2d. ed.). 1981. Cambridge, MA: MIT Press.
- Models of Bounded Rationality* (Vols. I and II). 1982. Cambridge, MA: MIT Press.
- Reason in Human Affairs*. 1983. Stanford, CA: Stanford University Press.
- Protocol Analysis: Verbal Reports as Data*. With K. A. Ericsson, 1984. Cambridge, MA: Bradford Books/MIT Press.
- Scientific Discovery: Computational Explorations of the Creative Processes*. With P. Langley, G. Bradshaw, and J. M. Zytkow, 1987. Cambridge, MA: MIT Press.
- Models of Thought* (Vol. II). 1989. New Haven, CT: Yale University Press.