

In M. Greenberger (Ed.),
Computers, communications, and
the public interest. Balti-
more, MD: The Johns Hopkins
Press, 1971.

2

**DESIGNING ORGANIZATIONS
FOR AN
INFORMATION-RICH WORLD**

Speaker **HERBERT A. SIMON**

*Richard King Mellon Professor of
Computer Science and Psychology
Carnegie-Mellon University*

Discussants **KARL W. DEUTSCH**

*Professor of Government
Harvard University*

MARTIN SHUBIK

*Professor of the Economics of Organization
Yale University*

Moderator **EMILIO Q. DADDARIO**

*Chairman of Subcommittee
on Science, Research, and Development
U.S. House of Representatives*

DADDARIO. I was attracted to participate in this series by the appropriateness of the theme, *Computers, Communications, and the Public Interest*. Perhaps the title of the present session should be "Designing Organizations for an Information-Rich, *Communications-Poor, Problem-Overwhelmed* World." If anything characterizes the current age, it is the complex problems of our technological civilization and the unpleasant physical and mental trauma they induce. John W. Gardner and other social critics warn us that a nation can perish from internal strains: indifference, unwillingness to face problems, incapacity to respond to human suffering, failure to adapt to new conditions, and the waning energy of old age. Gardner speaks of the "waxwork of anachronisms" in government and the "impenetrable web of vested interests" in unions, professions, universities, and corporations. He argues for "a society (and institutions) capable of continuous change, continuous renewal, and continuous responsiveness."¹

I see no room for complacency by the growing community devoted to communications and information processing in the face of the obvious needs of our society. Today we exchange a growing proportion of knowledge in new ways: via magnetic tapes, remote consoles wired to computers, national and international information networks, and large data banks. Expansion is so rapid, it is hard to document what is happening.

What concerns many of us is what I expect our speaker and discussants will be addressing in part. The creation of powerful computerized information systems, unless we take precautionary steps, may spawn new systems in Parkinsonian abandon, leading to quality-poor scientific and technical information. Furthermore, science can only flourish when it is untrammelled and open-ended. We must be careful not to institutionalize our information systems in such a way that they inhibit or interfere with this necessary freedom.

Herbert Simon is a member of the distinguished Panel on Technology Assessment of the National Academy of Sciences. I owe the panel a personal debt of gratitude for an outstanding report it recently completed on technology assessment.²

SIMON. If men do not pour new wine into old bottles, they do something almost as bad: they invest old words with new meanings. "Work" and "energy" are venerable English words, but since the Industrial and Scientific Revolutions they have acquired entirely new definitions. They have become more abstract and divorced from directly sensed qualities of human activity; and they have become more precise, finding expression in quantitative units of measurement (foot-pound, erg) and exact scientific laws (Conservation of Energy). The word "energy" uttered in a contemporary setting may represent quite different concepts and thought processes from the word "energy" uttered in the eighteenth century.

Old word meanings do not disappear; they tend to persist alongside

the new. This is perhaps the most insidious part of what C. P. Snow has dubbed the problem of the two cultures. To know what a speaker means by "energy" it is not enough to know what century he is speaking in, but also whether his talk belongs to the common culture or the scientific culture. If the former, his words should not be credited with the quantitative precision that belongs to the latter; and if the latter, his words should not be interpreted vaguely or metaphorically.

Old Words in New Meanings

All of this is preliminary to raising a difficulty I must hurdle to communicate. I intend to use familiar words like "information," "thinking," and "organization," but not with the meanings that the common culture has attached to them over the centuries. During the past twenty-five years these words have begun to acquire new, increasingly precise and quantitative meanings. Words associated with the generation and conversion of information are today undergoing a change of meaning as drastic as that experienced by words associated with the generation and conversion of energy in the eighteenth and nineteenth centuries.

Within the common culture, one cannot carry on a twentieth-century conversation about energy with a physicist or engineer. Similarly, it is increasingly difficult to carry on a twentieth-century conversation about information with a social scientist who belongs to the humanistic rather than scientific subculture of his discipline. The difficulty does not stem from jargon but from a complete disparity of meanings hidden behind a superficially common language.

What do I mean when I say: "Machines think"? The word "machine" seems obvious enough: a modern electronic digital computer. But "machine" has all sorts of unintended humanistic overlays. A machine, in the common culture, moves repetitively and monotonously. It requires direction from outside. It is inflexible. With the slightest component failure or mismanagement it degenerates into senseless or random behavior.

A computer may exhibit none of these mechanical properties. While retaining the word "machine" in the scientific culture as a label for a computer, I have revised drastically the associations stored with the word in my memory. When I say "Machines think," I am *not* referring to devices that behave repetitively and inflexibly, require outside guidance, and often become random.

The word "think" itself is even more troublesome. In the common culture it denotes an unanalyzed, partly intuitive, partly subconscious and unconscious, sometimes creative set of mental processes that sometimes allows humans to solve problems, make decisions, or design something. What do these mental processes have in common with the processes computers follow when they execute their programs?

The common culture finds almost nothing in common between them. One reason is that human thinking has never been described, only labeled. Certain contemporary psychological research, however, has been producing computer programs that duplicate the human information processing called thinking in considerable detail.³ When a psychologist who has been steeped in this new scientific culture says "Machines think," he has in mind the behavior of computers governed by such programs. He means something quite definite and precise that has no satisfactory translation into the language of the common culture. If you wish to converse with him (which you well may not!) you will have to follow him into the scientific culture.

As the science of information processing continues to develop, it will not be as easy to sequester it from the main stream of managerial activity (or human social activity) as it was to isolate the physical sciences and their associated technologies. Information processing is at the heart of executive activity, indeed at the heart of all social interaction. More and more we are finding occasion to use terms like "information," "thinking," "memory," and "decision making" with twentieth-century scientific precision. The language of the scientific culture occupies more and more of the domain previously reserved to the common culture.

Make no mistake about the significance of this change in language. It is a change in thought and concepts. It is a change of the most fundamental kind in man's thinking about his own processes—about himself.

The Scarcity of Attention

My title speaks of "an information-rich world." How long has the world been rich in information? What are the consequences of its prosperity, if that is what it is?

Last Easter, my neighbors bought their daughter a pair of rabbits. Whether by intent or accident, one was male, one female, and we now live in a rabbit-rich world. Persons less fond than I am of rabbits might even describe it as a rabbit-overpopulated world. Whether a world is rich or poor in rabbits is a relative matter. Since food is essential for biological populations, we might judge the world as rabbit-rich or rabbit-poor by relating the number of rabbits to the amount of lettuce and grass (and garden flowers) available for rabbits to eat. A rabbit-rich world is a lettuce-poor world, and vice versa.

The obverse of a population problem is a scarcity problem, hence a resource-allocation problem. There is only so much lettuce to go around, and it will have to be allocated somehow among the rabbits. Similarly, in an information-rich world, the wealth of information means a dearth of something else: a scarcity of whatever it is that information consumes. What information consumes is rather obvious: it consumes the attention of its recipients. Hence a wealth of information creates a poverty of attention

and a need to allocate that attention efficiently among the overabundance of information sources that might consume it.

To formulate an allocation problem properly, ways must be found to measure the quantities of the scarce resource; and these quantities must not be expandable at will. By now, all of us have heard of the *bit*, a unit of information introduced by Shannon in connection with problems in the design of communication systems.⁴ Can we use the bit as a measure of an information-processing system's capacity for attention?

Unfortunately, it is not the right unit. Roughly, the trouble is that the bit capacity of any device (or person) for receiving information depends entirely upon how the information is encoded. Bit capacity is not an invariant, hence is an unsuitable measure of the scarcity of attention.

A relatively straightforward way of measuring how much scarce resource a message consumes is by noting how much time the recipient spends on it. Human beings, like contemporary computers, are essentially serial devices. They can attend to only one thing at a time. This is just another way of saying that attention is scarce. Even the modern time-sharing systems which John Kemeny described are really only doing one thing at a time, although they seem able to attend to one hundred things at once.⁵ They achieve this illusion by sharing their time and attention among these hundred things. The attention-capacity measure I am proposing for human beings applies as well to time-sharing systems and also to an organization employing many people, which can be viewed as a time-sharing system.

Scarcity of attention in an information-rich world can be measured in terms of a human executive's time. If we wish to be precise, we can define a standard executive (IQ of 120, bachelor's degree, and so on) and ask Director Lewis Branscomb to embalm him at the National Bureau of Standards. Further, we can work out a rough conversion between the attention units of human executives and various kinds of computers.

In an information-rich world, most of the cost of information is the cost incurred by the recipient. It is not enough to know how much it costs to produce and transmit information; we must also know how much it costs, in terms of scarce attention, to receive it. I have tried bringing this argument home to my friends by suggesting that they recalculate how much the *New York Times* (or *Washington Post*) costs them, including the cost of reading it. Making the calculation usually causes them some alarm, but not enough for them to cancel their subscriptions. Perhaps the benefits still outweigh the costs.

Having explained what I mean by an information-rich world, I am now ready to tackle the main question. How can we design organizations, business firms, and government agencies to operate effectively in such a world? How can we arrange to conserve and effectively allocate their scarce attention?

I shall proceed with the help of three examples, each illustrating a

major aspect of the problem of organizational design. I make no attempt to cover all significant problem areas, and any fancied resemblance of my hypothetical organizations to real organizations, living or dead, in the city of Washington, are illusory, fortuitous, and the product of the purest happenstance.

Information Overload

Many proposals for eliminating *information overload* (another phrase to describe life in an information-rich world) call for a new computing system. There is good precedent for this. The Hollerith punched card is a creative product of the Census Bureau's first bout with information overload, and a series of crises in the central exchanges of the phone company led to the invention of automatic switching systems.

Today, some argue that the postal service is doomed to collapse from information overload unless means are found to automate the sorting operations. This cannot be so. There is no reason why mail-sorting costs should increase more than proportionally with the volume of mail, nor why unit costs should rise with volume. A major cause of the problem is that certain information-processing services are almost free, resulting in an explosive demand for them. The Post Office is not really prepared to provide this implicit subsidy and reneges by performing the services badly, with insufficient resources. The crisis in the Post Office does not call for computers; it calls for a thoroughgoing application of price and market mechanisms.

This is not to argue that any particular manual Post Office operation, such as sorting, cannot be made more economical by computer. This kind of technical question is settled by cost-benefit analysis within reasonable limits of error and debate. But there is no magic in automation that allows it to resolve dilemmas posed by an organization's unwillingness or inability to allocate and price scarce information-processing resources, whether the resources are sorting clerks or electronic devices. Free or underpriced resources are always in desperately short supply. What is sometimes alleged to be technological lag in the Post Office is really failure of nerve.

A computer is an information-processing system of quite general capability. It can receive information, store it, operate on it in a variety of ways, and transmit it to other systems. Whether a computer will contribute to the solution of an information-overload problem, or instead compound it, depends on the distribution of its own attention among four classes of activities: listening, storing, thinking, and speaking. A general design principle can be put as follows:

An information-processing subsystem (a computer or new organization unit) will reduce the net demand on the rest of the organization's attention only if it absorbs more information previously received by others than it produces—that is, if it listens and thinks more than it speaks.

To be an attention conserver for an organization, an information-processing system (abbreviated IPS) must be an information condenser. It is conventional to begin designing an IPS by considering the information it will *supply*. In an information-rich world, however, this is doing things backwards. The crucial question is how much information it will allow to be *withheld* from the attention of other parts of the system.

Basically, an IPS can perform an attention-conserving function in two ways: (1) it can receive and store information that would otherwise have to be received by other systems, and (2) it can transform or *filter* input information into output that demands fewer hours of attention than the input.

To illustrate these two modes of attention conservation, let me talk about some of the information needs of a nation's Foreign Office. (Since the United States has a State Department and not a Foreign Office, I am obviously talking about some other country.) The bulk of information that enters a system from the environment is irrelevant to action at the time of entry. Much of it will never be relevant, but we cannot be sure in advance which part this is.

One way to conserve Foreign Office attention is to interpose an IPS (human, automated, or both) between environment and organization to index and store information on receipt. A second way is to have an IPS analyze, draw inferences from, and summarize the information received, then index and store the products of its analyses for use by the rest of the system.

This proposal has a familiar ring about it. I have simply described in unconventional language the conventional functions of a conventional intelligence unit. Moreover, I have solved the information-overload problem simply by adding information processors. I eliminated scarcity by increasing the supply of scarce resources. Any fool with money can do that.

But the very banality of my solution carries an important lesson. The functional design an IPS must have to conserve attention is largely independent of specific hardware, automated or human. Hardware becomes a concern only later in economic considerations.

My proposal, however, is actually far less conventional than it sounds. If the IPS is to be even partly automated, we must provide precise descriptions (in the language of the scientific culture) of the processes denoted by vague terms like "analyze" and "summarize." Even if we do *not* intend to automate the process, the new information-processing technology still will permit us to formulate the programs of human analysts and summarizers with precision so that we can predict reliably the relation between inputs and outputs. Looking more closely at the structure and operation of the IPS, we see it really will not resemble a traditional intelligence unit very closely at all. (My thinking on this problem has benefited greatly from acquaintance with the analyses that have been made over the past several years of information-processing requirements in the U.S. State Department.

These planning activities have been laudably free from premature obsession with automated hardware.)

The purpose of the intelligence IPS I have proposed is not to *supply* the Foreign Office with information but to *buffer* it from the overrich environment of information in which it swims. Information does not have to be attended to (*now*) just because it exists in the environment. Designing an intelligence system means deciding: when to gather information (much of it will be preserved indefinitely in the environment if we do not want to harvest it now); where and in what form to store it; how to rework and condense it; how to index and give access to it; and when and on whose initiative to communicate it to others.

The design principle that attention is scarce and must be preserved is very different from a principle of "the more information the better." The aforementioned Foreign Office thought it had a communications crisis a few years ago. When events in the world were lively, the teletypes carrying incoming dispatches frequently fell behind. The solution: replace the teletypes with line printers of much greater capacity. No one apparently asked whether the IPS's (including the Foreign Minister) that received and processed messages from the teletypes would be ready, willing, and able to process the much larger volume of messages from the line printers.

Everything I have said about intelligence systems in particular applies to management information systems in general. The proper aim of a management information system is not to bring the manager all the information he needs, but to reorganize the manager's environment of information so as to reduce the amount of time he must devote to receiving it. Restating the problem this way leads to a very different system design.

The Need to Know

That brings me to the question of *the need to know*. How do we go about deciding where information should be stored in an information-rich world and who should learn about it?

Those of us who were raised during the Great Depression sometimes do not find it easy to adapt to an affluent society. When we ate potatoes, we always ate the peel (which my mother insisted was the best part of the potato). Nonreturnable containers seem to us symbols of intolerable waste.

Our attitudes toward information reflect the culture of poverty. We were brought up on Abe Lincoln walking miles to borrow (and return!) a book and reading it by firelight. Most of us are constitutionally unable to throw a bound volume into the wastebasket. We have trouble enough disposing of magazines and newspapers. Some of us are so obsessed with the need to know that we feel compelled to read everything that falls into our hands, although the burgeoning of the mails is helping to cure us of this obsession.

If these attitudes were highly functional in the world of clay tablets, scribes, and human memory; if they were at least tolerable in the world of the printing press and the cable; they are completely maladapted to the world of broadcast systems and Xerox machines.

The change in information-processing technology demands a fundamental change in the meaning attached to the familiar verb "to know." In the common culture, "to know" meant to have stored in one's memory in a way that facilitates recall when appropriate. By metaphoric extension, "knowing" might include having access to a file or book containing information, with the skill necessary for using it.

In the scientific culture, the whole emphasis in "knowing" shifts from the storage or actual physical possession of information to the process of using or having access to it. It is possible to have information stored without having access to it (the name on the tip of the tongue, the lost letter in the file, the unindexed book, the uncatalogued library); and it is possible to have access to information without having it stored (a computer program for calculating values of the sine function, a thermometer for taking a patient's temperature).

If a library holds two copies of the same book, one of them can be destroyed or exchanged without the system's losing information. In the language of Shannon's information theory, multiple copies make the library *redundant*. But copies are only one of three important forms of redundancy in information. Even if a library has only one copy of each book, it still has a high degree of informational overlap. If half the titles in the Library of Congress were destroyed at random, little of the world's knowledge would disappear.

The most important and subtle form of redundancy derives from the world's being highly lawful. Facts are random if no part of them can be predicted from any other part—that is, if they are independent of each other. Facts are lawful if certain of them *can* be predicted from certain others. We need store only the fraction needed to predict the rest.

This is exactly what science is: the process of replacing unordered masses of brute fact with tidy statements of orderly relations from which these facts can be inferred. The progress of science, far from cluttering up the world with new information, enormously increases the redundancy of libraries by discovering the orderliness of the information already stored. With each important advance in scientific theory, we can reduce the volume of explicitly stored knowledge without losing any information whatsoever. That we make so little use of this opportunity does not deny that the opportunity exists.

Let me recite an anecdote that illustrates the point very well. We are all aware that there is a DDT problem. DDT is one of technology's mixed blessings. It is very lethal to noxious insects, but uncomfortably persistent and cumulatively harmful to eagles, game fish, and possibly ourselves. The

practical problem is how to enjoy the agricultural and medical benefits afforded by the toxicity of DDT without suffering the consequences of its persistence.

A distinguished chemist of my acquaintance, who is a specialist neither in insecticides nor biochemistry, asked himself this question. He was able to write down the approximate chemical structure of DDT by decoding its name. He could recognize from general theoretical principles the component radicals in the structural formula that account for its toxicity. The formula also told him on theoretical grounds why the substance is persistent and why the molecule does not decompose readily or rapidly. He asked, again on theoretical grounds, what compound would have the toxicity of DDT but decompose readily. He was able to write down its formula and saw no theoretical reason why it could not easily be produced. (All of this took ten minutes.)

A phone call to an expert in the field confirmed all his conjectures. The new compound he had "invented" was a well-known insecticide, which had been available commercially before DDT. It is not as lethal as DDT over as broad a band of organisms but is nearly so, and it decomposes fairly readily. I do not know if the new-old chemical "solves" the DDT problem. The durability of DDT was intended by its inventors to avoid frequent respraying and reduce the costs of treatment. There may be other economic issues, and even chemical and biological ones.

What the story illustrates is that good problem-solving capacities combined with powerful (but compact) theories (and an occasional telephone call) may take the place of shelves of reference books. It may often be more efficient to leave information in the library of nature, to be extracted by experiment or observation when needed, than to mine and stockpile it in man's libraries, where retrieval costs may be as high as the costs of recreating information from new experiments or deriving it from theory.

These considerations temper my enthusiasm for using new technology to store and retrieve larger and larger bodies of data. I do not mean to express a blanket disapproval of all proposals to improve the world's stores of information. But I do believe we must design IPS's with data-analysis capabilities able to keep up with our propensities to store vast bodies of data.

Today's computers are moronic robots, and they will continue to be so as long as programming remains in its present primitive state. Moronic robots can sop up, store, and spew out vast quantities of information. They do not and cannot exercise due respect for the scarce attention of the recipients of this information. Computers must be taught to behave at a higher level of intelligence. This will take a large, vigorous research and development effort.

In a knowledge-rich world, progress does not lie in the direction of reading and writing information faster or storing more of it. Progress lies in

the direction of extracting and exploiting the patterns of the world so that far less information needs to be read, written, or stored. Progress depends on our ability to devise better and more powerful thinking programs for man and machine.

Technology Assessment

Attention is *generally* scarce in organizations, *particularly* scarce at the tops, and *desperately* scarce at the top of the organization called the United States Government. There is only one President. Although he is assisted by the Budget Bureau, the Office of Science and Technology, and other elements of the Executive Office, a frightening array of matters converges on this single, serial, human information-processing system.

There is only one Congress of the United States. It can operate in parallel through committees, but every important matter must occupy the attention of many Congressmen. Highly important matters may claim the time and attention of all.

There is only one body of citizens in the United States. Large public problems such as the Vietnam War, civil rights, student unrest, the cities, and environmental quality (to mention five near the top of the current agenda) periodically require a synchrony of public attention. This is more than enough to crowd the agenda to the point of unworkability or inaction.

Congressman Daddario has devoted a great deal of thought in recent years to improving the procedures in society and government for dealing with the new technology we produce so prodigiously. At the request of his House Subcommittee on Science, Research, and Development, a panel of the National Academy of Sciences on which I served recently prepared the report on technology assessment to which he referred.

Technology assessment is not just a matter of determining the likely good and bad effects of new technological developments. Even less is it a matter of making sure, before new technology is licensed, that it will have no undesirable effects. The dream of thinking everything out before we act, of making certain we have all the facts and know all the consequences, is a sick Hamlet's dream. It is the dream of someone with no appreciation of the seamless web of causation, the limits of human thinking, or the scarcity of human attention.

The world outside is itself the greatest storehouse of knowledge. Human reason, drawing upon the pattern and redundancy of nature, can predict some of the consequences of human action. But the world will always remain the largest laboratory, the largest information store, from which we will learn the outcomes, good and bad, of what we have done. Of course it is costly to learn from experience; but it is also costly, and frequently much less reliable, to try through research and analysis to anticipate experience.

Technology assessment is an intelligence function. If it operated perfectly, which it is certain not to, it would do two things for us. First, it would warn us before our taking action of the really dangerous (especially the irreversibly dangerous) consequences possible from proposed innovations. Second, it would give us early warning of unanticipated consequences of innovations as they became visible, before major irreversible damage had been done. In performing both of these functions, technology assessment would be mindful of the precious scarcity of attention. It would put on the agenda only items needing attention and action (including the action of gathering information to evaluate the need for further attention).

A phrase like "technology assessment" conjures up a picture of scientific competence and objectivity, deliberateness and thoughtfulness, concern for the long run, and a systems view that considers all aspects and consequences. But these desirable qualities of a decision-making system cannot be imposed without considering the organizational and political environment of the system.

As our scientific and engineering knowledge grows, so does the power of our actions. They have consequences ramifying over vast reaches of space and time. The growth of knowledge allows us to recognize consequences we would have been ignorant of or ignored before. We are able to make bigger waves and at the same time have more sensitive instruments to detect the rocking of the boat. Today we sterilize and quarantine everything that travels between earth and moon. Less than five hundred years ago we diffused tuberculosis, smallpox, and syphilis throughout the Americas in happy ignorance.

The injunction to take account of *all* effects conjures up the picture of an integral stretching out through space and time without ever converging. We must assume, as mankind has always assumed, that a reasonable allocation of our limited attention and powers of thought will solve the crucial problems facing us at least as fast as new ones arise. If that assumption is wrong, there is no help for us. If it is right, then technology assessment becomes part and parcel of the task of setting an agenda for society and government.

To bring the notion of technology assessment out of the realm of abstraction, let me go back to the example of DDT. Although I have not researched the history of DDT, I believe it was introduced on a large scale without thorough (or at least adequate) study of its potential cumulative danger in the atmosphere and in organisms (especially predators). It was hailed for its agricultural and medical benefits as one of technology's miracles. Now, some decades later, we learn that the miracle has a flaw.

The possible adverse effects of DDT have been known to specialists for some time. They were probably even known, but ignored, at the time DDT was introduced. If so, this would underscore my fundamental theme of the scarcity of attention.

Suppose the dangers of DDT were not known beforehand but were discovered only in the laboratory of nature. Then, with apologies to eagle lovers, I am not sure that we (or even the eagles) have suffered unconscionable or irreversible loss by letting actual use tell us about DDT rather than trying to anticipate this experience in advance. Technology assessment has been (and is being) made by the environment. We are getting signals from the environment calling attention to some of its findings, and these signals are strong enough to deserve and get our attention. The DDT issue has been claiming attention intermittently for some months, with the loudest environmental signal being the detection of DDT in Great Lakes game fish. The issue is now high enough on the agenda of newspapers, courts, and committees to bring action.

I know this sounds complacent, and I really do not feel complacent. But it serves no useful social purpose to treat with anguish and hand-wringing every public problem which by hindsight might have been avoided if we had been able to afford the luxury of more foresight. Now that we *know* the problems, we should address them rather than hold inquests about who should have seen the problems earlier.

Our information about the effects of DDT and of long-continued diffuse contamination is in many respects unsatisfactory. (So is our information about almost any issue of public policy.) But this does not mean we could improve the situation by massive collection of data. On the contrary, we mainly need carefully aimed, high-quality biological investigations of the cause and effect mechanisms underlying the diffusion and metabolism of DDT. After we understand better the chemistry and biology of the problem we might make sense of masses of data, but then we probably would not need as much.

First-rate biologists and chemists capable of doing the required research are in as short supply as most other high-quality information-processing systems. Their attention is an exceedingly scarce commodity, and we are unlikely to capture much of it soon. The practical question, as always, is how to deal with the situation given the scrappy, inadequate data we now have.

We begin to ask questions like these: Assuming the worst possible case for the harmful effects of DDT, what is the magnitude of the effects in human, economic, and ecological terms, and to what extent are these effects irreversible? In the same terms, what would it cost us to do without DDT? What is the next best alternative?

These are common-sense questions. We do not have to know anything about the technology to ask them, although we might learn something about it from the answers. The most effective IPS for getting answers consists of a telephone, a Xerox machine (to copy documents the telephone correspondents suggest), and some very bright professionals (not necessarily specialists) who do know something about the technology. With this retrieval

system, just about anything in the world now known on the problem can be extracted in a few man-weeks of work. (The time required goes up considerably if hearings and briefings are held or a research project is organized.)

There are numerous locations inside and outside the federal government where the questions may be asked. They may be asked by the Office of Science and Technology, the National Academy of Sciences, the National Academy of Engineering, the RAND Corporation, Resources for the Future, or a Congressional committee. (An excellent example of the last is the recent series of reports on steam-powered automobiles.)

The location of the investigating group is significant from only one standpoint, which may be crucial. The location of the group can determine the attention it commands and the legitimacy accorded its findings. These are interdependent but by no means identical matters.

Legitimacy may sometimes be achieved (and even attention secured) by the usual credentials of science: the right degrees, professional posts, and reputations. But many an impeccable report is ignored, and many a report without proper credentials gains a high place on the agenda. The Ralph Naders of the world demonstrate that writing and speaking forcefully, understanding the mass media, and being usually right about the facts can compensate for missing union cards and lack of access to organizational channels. Rachel Carson showed that even literary excellence is sometimes enough to turn the trick.

I agree with Congressman Daddario that we can and should strengthen and make more effective the processes of technology assessment in our country. We shall still need the world itself as a major laboratory, but we may be able to substitute foresight for hindsight to a modest extent. Did we have to wait until all Los Angeles wept before doing anything about automobile exhausts? Well-financed institutions for technology assessment should be spending a hundred million dollars a year instead of ten million to find out whether the steam automobile offers a long-term solution to the smog problem. Our current measures are temporary expedients at best.

Strengthening technology assessment means improving our procedures for setting the public agenda. It does not mean pressing more information and problems on an already burdened President, Congress, and public. In an information-rich world, there is no special virtue in prematurely early warnings. Let the world store information for us until we can focus attention and thought on it.

Assessing Information-Processing Technology

The final issue I should like to address is itself a problem in technology assessment. The science and technology of information processing is only a quarter-century old, and we have merely the faintest glimmerings of what it will be like after another quarter-century. How shall we assess it and make sure it develops in socially beneficial ways?

The most visible and superficially spectacular part of the technology is its hardware: computers, typewriter consoles, cathode-ray tubes, and associated gadgets. These devices give us powerful new ways for recording, storing, processing, and writing information to improve and replace the human IPS's with which we had to make do throughout man's history.

By itself, the hardware does not solve any organizational problems, including the problems of attention scarcity. The hardware boxes will begin to make inroads on these problems only as we begin to understand information-processing systems well enough to conceive sophisticated programs for them—programs that will permit them to think at least as well as man does.

Each step we take toward increasing our sophistication and scientific knowledge about the automated IPS also increases our sophistication and scientific knowledge about the human IPS, about man's thought processes. What we are acquiring with the new technology is something of deep significance—a science of human thinking and organization.⁶

The armchair is no more effective a scientific instrument for understanding this new technology than it was for previous technologies. If we are to understand information processing, we must study it in the laboratory of nature. We must construct, program, and operate many kinds of information-processing systems to see what they do and how they perform.

Our first systems have performed and will perform in all sorts of unexpected ways (most of them stupid), and by hindsight they seem incredibly crude. They will never pass a cost-effectiveness test on their operating performance, and we shall have to write them off as research and development efforts. From their behavior, we may learn that the new technology contains dangers as well as promises. There already is considerable concern about threats to privacy that the new technology might create. Such concerns will be mere armchair speculations until they are tested against a broad base of experience.

Very early in the computer era, I advised several business firms not to acquire computers until they knew exactly how to use them and pay for them. I soon realized this was bad advice. Computers initially pay their way by educating large numbers of people about computers. They are the principal forces for replacement of the vague, inadequate common-culture meanings of words in the information-processing vocabulary by the sharp, rich, scientific meanings these words must have in the future.

I think this points to a clear public policy for understanding and assessing the new technology. We need greatly increased public support for research and development efforts of as varied a nature as possible. They should certainly include network experiments of the sort John Kemeny envisages. They should include data-bank experiments. Above all, they should include experiments in robotics, large-scale memory organization, and artificial intelligence, leading to a basic foundation for a science of information processing.

Past experience suggests that a program pursued in the experimental spirit I have indicated will have valuable by-products. List processing is an esoteric development of computer-programming languages that was motivated initially about fifteen years ago by pure research interests in artificial intelligence. Today, its concepts are deeply imbedded in the design of large programming and operating systems regularly used in accounting and engineering computation.

The exploration of the moon is a great adventure. After the moon, there are objects still farther out in space. But man's inner space, his mind, is less well known than the space of the planets. It is time we establish a national policy to explore this inner space systematically, with goals, timetables, and budgets. Will you think me whimsical or impractical if I propose that one of these goals be a world-champion chess-playing computer program by 1975; and another, an order-of-magnitude increase by 1980 in the speed with which a human being can learn a difficult school subject, such as a foreign language or arithmetic?

If we are willing to dedicate ourselves to national goals of this kind (if you do not like my two, substitute your own), set deadlines for them, and commit resources to them (as we have committed resources to exploration of outer space), I think we soon shall have an understanding of both the information processors we call computers and those we call man. This understanding will enable us to build organizations far more effectively in the future than has ever been possible before.

PANEL DISCUSSION

DEUTSCH. The speaker has given us an example of a good information-processing program by his own definition. His remarks represent a program which does not immediately produce all the answers, but sets into motion a process which eventually may well generate them. He has done this by giving us some highly simplified versions of his important and profound thought. But they seem overly simple on two points.

Simon has suggested that information overload could be relieved by better pricing policies. I think he implied that when anything is offered free, the demands for it become infinite. I do not know what goes on at the water cooler in his office, but I shudder at the vision of money being charged for drinking water, and licenses for breathing oxygen being marketed among the more affluent citizens of Los Angeles! Social scientists, including Simon, know cases well where the price mechanism is neither the only nor the best way to distribute a good. At zero cost, if the supply of a good is well ahead of the demand, demand does not automatically grow to infinity. It may

increase, but the rate of increase is a matter of empirical fact, and it need not be very fast.

The second notion that disturbed me even more for professional reasons was the splendid and provocative way Simon put before us a general theory of knowledge of a "try and see" type. He was obviously thinking of computers when he advised us not to be quite so fearful of the possible damages but to experiment and see what happens. If things go a little wrong, we will learn from the damage. After all, even DDT may have done more good than harm during the decades of its first application.

I do not know enough about DDT or ecology to have an opinion on this point. But one of my special interests is the study of international politics, and I tremble when I translate Simon's advice into this field. There we developed many interests in distant countries and strange peoples, together with supposedly adequate methods for gathering and evaluating information about them. We thought we had an adequate political intelligence system in the broadest sense, military and civilian, perhaps because of our fascination with a technology which permits faster transmission of mistaken estimates and faster transportation of misinformed officials. Our government acted on what it thought was information about Southeast Asia. We sent half a million men there, and back came some more information; also forty thousand bodies in plastic bags.

In foreign policy, at least, I should warn of the old-fashioned approach to pragmatism that says, "Let's go in and try and see what happens." I remember too readily the story of the man who fell off the top of the Empire State Building and on passing the second floor said, "On pragmatic grounds, so far things have been going well."

To make the case more general, if relatively little power is involved, the amount of damage possible from neglecting early warnings or not making forecasts is limited. The more powerful an operation, the more foresight becomes necessary to avoid fatal damage. The faster an automobile, the farther ahead must its headlights cast their beam.

As modern technology becomes more powerful, we need what my colleague John Platt calls anticipatory or crisis research:⁷ determining at which point current problems are likely to reach critical mass. I think Simon himself is one of the people who is very likely to contribute to this work, but under the pressures of speaking quickly and popularly he is in danger of giving an impression of cheerful pragmatism, telling all the "nervous Nellys" not to worry. Political scientists would not agree.

Having been critical, let me say how delighted I am with the important intellectual contributions which I think are made in Simon's paper. I believe there are four major contributions:

1. the economy of attention is itself put into the focus of the analysis;
2. attention economy is proposed as a design criterion for large organizations;

3. knowledge is defined as access to information, or the procedure of its retrieval, rather than its storage;
4. emphasis is put on producing information-recreating programs, programs whose economically small number of units and operating rules can regenerate information more cheaply and perhaps more quickly than storage and retrieval of the information itself.

Simon has put before us the problem of attention overload. The population of the world doubles every thirty-five to forty years. Since most of these people want to be heard from, information will increase proportionately to the population. Simon says this is all it will increase. I beg respectfully to differ.

The per capita income of people in the world goes up 2 percent a year, and the richer they get, the more they spend on services involving information. The total information load, therefore, goes up not 2 percent a year with population, but 4 to 5 percent a year with income. That is, it doubles about every eighteen years. In addition, almost 1 percent of the work force in the United States in recent years has shifted from production-line labor or material handling into knowledge-handling or symbol-manipulating occupations.

I have recently seen a curve telling the same story for Czechoslovakia from 1929 to 1965. The curve shows that the change-over to a Communist regime did not seem to make the slightest difference in this trend. The shift from proletarians to paper shufflers goes on under Communism as much as anywhere else! The tendency toward knowledge handling increases all the time.

Finally, about every four or five years in the last quarter-century a critical threshold has been broken, demanding new thinking. Examples include nuclear energy, the mass culture of television, interplanetary transportation, footprints on the moon, the change of universities from elite to mass institutions, and many others.

For these reasons, there will be much more information overload in the future than there has been in the past. The information pressure will at the very least, in my estimation, double every twelve to fifteen years.

There are, to be sure, various ways of dealing with overload:

1. skipping;
2. delay or forming waiting lines;
3. chunking, or doing things in large batches;
4. doing things in a shoddy way, or quickly, and accepting a large measure of error;
5. (the method which Simon prefers) filtering or reducing the information.

With filtering, once again in my own field I am frightened of the results. Lyndon B. Johnson was one of the best-filtered chief executives this country has ever had. He had secretaries, White House staff, and an envi-

ronment that very carefully gave him largely one kind of information. When a man like George Ball sometimes got the thankless task of serving as the devil's advocate, the devil got relatively little prime time and very little backing!

We run large risks of failure in coping with the problems of our cities, the poor, the races, international politics, arms control, veterans' systems, and all the rest, if our main preoccupation is to save executives from looking at too much uncongenial information. The uncongenial information may be the vital information.

This leads to the question of filter design. How do you make a filter intelligent enough to understand what to let through and what to screen out? The only approach I know is to analyze the operations involved in "attention." There are six:

1. *recognize* loosely what it is one should pay attention to (the target), such as things unfamiliar, strangers, or things that do not fit;
2. *track* the object of attention, and keep attention on it;
3. *interpret* the object and ask what it resembles;
4. *decide which response* to the object is most appropriate, and what should be done about it;
5. *carry out* the response;
6. *accept feedback*, and learn from the results of the response whether it was the right one and how future responses should be corrected.

All six are needed, but this means we cannot build good filters without delegating to them part of our *memories*. Filters without memories are primitive filters. Delegating memories means decentralizing decision-making capabilities and accepting some redundancy by building more capabilities, thereby increasing the total stock. We must ask whether the cost of errors of filtering our response is larger or smaller than the capital cost of building more response capabilities. The information-processing patterns of the future may not be the star or the wheel with the single hub, but rather the network where loads can be referred from one nodal point to another.

I would like to end with three small pieces of outlook.

1. The future of mankind in the advanced countries will become more intellectualized, not less, because of the shift toward knowledge-producing occupations. In the United States early in the century an eighth of the population was in knowledge-producing jobs. It is now a third, and by the end of the century it will be more than a half. The same goes for national income, given the present trends. It, too, will be allocated in an increasing proportion to the processing of information.
2. We will need much more innovation. It is impossible to carry on a great many of our past practices in either industry or politics without running into diminishing returns. The answer is to change the production function and find new ways that get around old, critical bottle-

necks. Creativity will be a matter of life and death for civilization in the coming decades.

3. There will be a great need for decentralization, shorter feedbacks, faster turnaround times, and better fits between decentralized memory and nature.

I think an intellectualized, innovative, more decentralized future is coming. We shall vitally require the contribution of computers, and even more, the contribution of creative human scholars like Simon.

SIMON. Professor Deutsch can disagree with me, but that does not force me to disagree with him! And there are really very few things I want to disagree with him about.

I will plead guilty to hyperbole. I will cheerfully plead guilty to cheerful pragmatism. Frankly, I think we could use some more of it in this country at the present time.

Without enlarging on my reasons, I will take exception to Deutsch's example. However one diagnoses our failures in Vietnam, I find it hard to interpret them as failures of technology assessment. This seems to me a very peculiar diagnosis.

On filtering, I think we can move toward appropriate decentralization and operate partly in parallel, but we shall still require some serial and synchronous operation for which we do need to ask how we filter. How will the people who stand at the apexes of decision systems receive an appropriately filtered range of considerations bearing on the decisions they have to make? Nothing we can do will change the organization of the United States government so that there will be seven Presidents instead of one. We have to find ways to allow that one President to operate as intelligently as possible on the basis of the best information obtainable.

SHUBIK. In Lewis Carroll's *The Hunting of the Snark*, the bellman (who is the central character) says, "What I have said three times is true." Even though Herbert Simon may have said three times that man is a one-at-a-time, sequential, data-processing animal, I think the question is still wide open. As a matter of fact, I offer as a mild piece of counterevidence that at precisely the same time that I was listening to Simon discuss the meanings of the word "know" I noted two eminent Old Testament experts in the audience also discussing the meaning of the word "know."

My remarks, which are variations on some of the themes of Simon, might be subtitled "How to be data rich and information poor, or let's bury ourselves with the facts." Consider the following three quotes:

1. "Don't confuse me with the facts. My mind is made up."
2. "It isn't what you know. It's whom you know."
3. "Why should I read his paper? I can prove the results more quickly myself."

Should these statements be regarded pejoratively? Not necessarily. Since

data are expensive to handle, interpret, and turn into useful information, it is often good that somebody does not intend to gather further facts but has his mind made up.

A human society may be regarded as a data-processing system. Information is stored in institutions, books, brains, statistical tables, computer libraries, and elsewhere. The knowledge of who knows what, or where the information is, may frequently be more valuable than deep, narrow knowledge of a few subjects themselves.

In accordance with Simon's remark that various words used in everyday conversation sometimes take on completely different technical meanings, I wish to differentiate between *raw data* and *information*. When a communications engineer talks about information he refers to an abstract concept defined in terms of messages sent through a transmission system. Whether or not the messages make sense to the listener is of no concern so long as they are heard. A Shakespearean sonnet and a group of random noises can contain the same amount of information under this definition. The difference comes in the interpretation by the listener.

In human decision-making systems we distinguish between messages containing gibberish and messages containing significant information. A key unsolved problem is how to measure the contextual or semantic information content of a message in a social system. The technical definition of information does not help much.

In operating a communications network, the costs and techniques of data transmission are important. In a social system two added features are critical. These are the determination of the best way to use symbols or words to convey meaning and translation of the message into information of importance to decision making. The key word in coding and decoding messages in a social system is *evaluation*.

Many of us, despite Simon's "timely" warning, spend too much time reading the *New York Times*. I am reminded of the *Harvard Lampoon's* observation that the slogan of the *Times* should be, "All the News that Fits We Print." At least this sensory system provides us with some preselection. We live in a data-rich world, with the costs of transmission, storage, and reproduction declining rapidly. Wise men can now spend more time and resources decoding and evaluating more easily obtainable data. Less wise men run the danger of confusing their lack of understanding of the ever-mounting volume of data with being better informed.

An optimum information system is not necessarily the one which processes the most data. An optimum system for protecting the average stockholder does not supply him with full, detailed financial accounts. In fact, one can easily swindle the unwary by supplying them with financial details and footnotes they do not understand. It is now possible to bombard a generally uncomprehending public with myriad details on pollution, the pros and cons of insecticides, the value and dangers of irrigation schemes, on-the-spot

reports of rioting and looting, televised moon landings, suicides, murders, and historical prices of thousands of stocks and commodities.

The computer and modern communications technology are to the study of man and society as important as the telescope and the microscope were to astronomy and the biological sciences. But the lyrical descriptions of technologists often omit consideration of limited human time and capability. The data-rich world may appear to make the millenium around the corner. But unorganized data are not knowledge, and modern communications contain the potential to confuse as well as educate.

How much time can the man on the street devote to politics? As population grows and the world becomes more complex, how can society keep the individual supplied with the right information for making political decisions and preserving his freedom? The problem is not the speed of generation or transmission of bits of raw data per second. It lies at the far more fundamental level of interpretation and understanding. Within a few years it may be possible to have a virtually instant referendum on many political issues. This could represent a technical triumph—and a social disaster if instability resulted from instantaneous public reaction to incompletely understood affairs magnified by quick feedback.

Consider some of the possible dangers. What is the first great TV, time-sharing demagogue going to look like? How will he put to use such extra features of modern communications as virtually instantaneous feedback? When will a TV screen with the appropriate sensory feelings be able to portray the boss behind his mahogany desk (two thousand miles away) who fires or chastises his employee, and makes him feel just as small, and his palms just as clammy with sweat, as if he were in the room with him? When will the first time-shared riot occur? Orson Welles came close in the thirties with a fairly good radio panic. Current techniques for mob control require physical proximity. In the Brave New World, will we still regard a mob as a great number of closely packed people, or will isolated mobs interacting via TV consoles and operating over large areas be more efficient?

We are moving into an era of large data banks and mass data handling. There are still many technical problems to be solved, but they are minor in comparison to the sociotechnical problems of designing organizations for a data-rich world. If we are to avoid the organizational nightmares depicted in Borges' library or Kafka's courtrooms, we must improve our design and understanding of man-machine organizations. The answers lie more in educating man to process data into better and more relevant information than in shaving microseconds off processing times.

Simon has suggested that working on *prima facie* impractical problems such as a program to play championship chess may provide us with valuable insights and important applications. I agree, and would like to extend his remarks. We must know more about ourselves as data processors and

decision makers; but knowledge of ourselves is expensive. People do not have billions of man-hours to expend in controlled experiments. The electronics age, however, provides us with devices to observe ourselves as we function.

As one example, we could wire the chess boards at all chess clubs and sensitize the chess pieces so that a computer could immediately record each game played. Someone who knew the right questions to ask could extract information and then destroy the data, in case retrieval and storage costs exceeded the expected value of keeping the raw data.

As a second illustration, a way to charge people far more equitably than at present for automobile driving and parking would be to attach a small transmitter to each car and have a central computer bill the owner every month according to times spent in various zones. These data would also permit study of traffic movement and driving behavior, although many of us might argue against having Big Brother keep track of every move of our personal cars.

I think that Simon's interest in chess games and artificial intelligence is legitimate and terribly important, but it omits social intelligence, which may be extremely different from individual intelligence. It may be easier to build an artificial player in some social situations involving pure competition than an artificial player for chess. How do you construct a nice guy? How do you construct a stable individual? Artificial players with these properties in social situations do not have to be terribly bright to be reasonably effective.

I would like to suggest that our organizations and many of our activities be looked at as experiments. We must work out how to make on-line observations with computers and still preserve individual freedom and privacy. If we do not, we run the extreme risk of never learning enough about ourselves and our organizational abilities to be able to cope with the complexities of the world ahead.

The generation and transmission of raw data with speed and economy provides twentieth-century man with the necessary conditions for survival. The conditions are not sufficient, however, without man's studying himself. The new technology has provided us the needed medicines, but we must be our own doctors.

SIMON. I am very much in accord with Shubik's final suggestion that the big job ahead is to use the new information-processing technology to understand ourselves better.

On the question of whether a human being is a serial processing device, the lesson I drew from Shubik's counterevidence was that clearly I was not saying enough per minute to load his channel! By the way, I am obviously not a Biblical scholar or I would not have used the particular definition of "knowing" that I did.

I would like to express rather sharp disagreement with Shubik on the Orwellian tones of some of his remarks. I think he underestimates human

beings as receivers of information. Even before television, we lived in an environment of information mostly conveyed by our neighbors, including some pretty tall tales. We acquired a variety of techniques for dealing with information overload. We know there are people who can talk faster than we can and give us an argument on almost any topic. We listen patiently, because we cannot process information fast enough to refute them; that is, until the next day, when we find the hole in their argument. A relevant rule that my father taught me was, "Never sign in the presence of the salesman." By adopting such rules and their extensions, we allow ourselves the extra processing time needed to deal with the information overload.

I am not really worried about Big Brother booming over the air-waves. I think that at all levels of intelligence, human beings have common sense protecting them from the worst features of their information environment. If information overload ever really gets the best of me, my last resort is to follow the advice of Gertrude Stein in the opening pages of *The Autobiography of Alice B. Toklas*. "I like a view, but I like to sit with my back turned to it."

DADDARIO. During this interesting exchange, we have seen some harsh sniping by Shubik, and a charge by Deutsch that Simon has been a little too simple! We now come to the question period, if Martin Greenberger will take the podium.

GREENBERGER. In the interests of efficiency, I shall group the questions. The first set of three should have special interest to Congressman Daddario, as well as to Herbert Simon.

William Moore asks if restricting scientific research and thereby slowing down the process of change would not have the good effect of reducing disruption and social revolt, as alleged by historian Arnold Toynbee.

Quentin Ludgin, a laboratory chief at the Bureau of the Census, inquires whether Simon's opposition to premature early-warning methods recognizes the possibility that the ecology may be endangered by rapid and soon irreversible changes. Is not an early-warning network desirable?

Phil Hirsch, Washington editor of *Datamation*, wonders if technology does not carry an inherent probability or certainty of danger, given its present development and rapid advance. Should not information scientists and technicians, as citizens and taxpayers, apply their expertise to this problem, even if it means expressing controversial opinions?

DADDARIO. I would answer the first question in the negative. We should not restrict scientific research, or we would certainly have a sterile society which is unable to cope with its complex problems.

Combining the next two questions, could our technology become irreversible, does it carry dangers, and are we doing enough? I would say we are not doing enough. There is a risk that some of the tendencies becoming set in our environment are so subtle and complicated that they could in fact be

irreversible and difficult to determine until they are exceedingly dangerous. The answer is to do more than we are: establish within our society the capability to anticipate these dangers, develop mechanisms to cope with them, and make positive rather than negative the results of our technical and scientific society to meet the challenge.

SIMON. My answers are very nearly identical with Congressman Daddario's. I do not think we should or can restrict scientific research. A number of our problems are going to be solved only if we have more knowledge, not less.

Whether an early-warning network is desirable depends on what is meant by such a network. With the large number of people in our country knowledgeable about large numbers of things, we already have an early-warning network. The real problems are the bottlenecks of attention. How do you get knowledge to the appropriate points in the network? How do you decide which of two things is more alarming?

I think information technologists have the same responsibility that all knowledgeable people in a democracy have: to devote some time and energy to understanding and trying to explicate the social significance, importance, and possible consequences of things about which they are knowledgeable. A good many people with whom I am acquainted do devote a considerable part of their energy to this kind of activity. In fact, I believe that is one of the important things this whole series of lectures is about.

GREENBERGER. The next question is from Stephen J. Tauber, chief of the Information Sciences Section of what he calls "Branscomb's Executive Embalming Service," namely the National Bureau of Standards. He asks if the information supply and demand problem really is not one largely of an ever-increasing demand function and a time lag in the supply function before the new equilibrium can be established. Also, William M. Hornish, a manager at Western Union, inquires how information can be organized so that its use is meaningful to our individual experiences.

SHUBIK. It is again the problem of what we mean by "information." People who insist on mistaking raw data for information see the demand for information as getting bigger and bigger. Some retailers, for example, suggest putting a card on every dress to obtain immediate notice of stock-out. It is easy to obtain new stacks of data; but what is often forgotten is how to analyze the extra data.

I will tell a story that is germane. It is the story of the owl, who is regarded as the wisest beast in the forest, and the centipede, who comes to the owl with ninety-nine sore feet and asks, "What am I going to do?" The owl looks at the centipede and says, "It is simple. You walk an inch off the ground for the next two weeks." The centipede thinks about it and becomes convinced that sure enough, it would give its pads time to heal. Then the centipede says to the owl, "How?" and the owl replies, "I have solved your conceptual problem. Don't bother me with the technical details."

This is the fundamental dilemma of data generation in our society. We are too ready to request a big sample without worrying about how to analyze it. That is a technical detail.

GREENBERGER. Perhaps Mr. Hornish's question was asked in the spirit of the owl and centipede story. Simon spoke of the need for computers to organize, filter, and make information more meaningful and easier to digest. Mr. Hornish may be seeking further advice on what really can be done to help the busy executive with his scarcity of attention.

SIMON. I think that is at the very heart of the question. As long as we use the vast power of computers to spew out data in forms which human beings then have to process extensively, we obviously are not using that capacity very intelligently. Each step we take, on the other hand, toward making computers able to carry out the kind of processing, analysis, and condensing of data that is called "thinking" in people, the more people and computers will be able to work in fruitful symbiosis.

GREENBERGER. The next question is from Professor Eliezer Naddor of The Johns Hopkins University, and is really an assertion rather than a question. He states that computers should not be programmed to make decisions, but only to assist humans make better decisions. Hence, he disagrees with the 1975 goal to have a computer be world chess champion, preferring the goal of having a computer assist Herbert Simon to become champion.

SIMON. I should enjoy being world chess champion, but the joy would be tarnished if the computer had to stand at my back while I was playing. The national goal of developing a computer world chess champion would force us really to discover some fundamental things about how human beings think and solve certain classes of problems. The reason for setting the task is to permit the needed research and investigation.

The sweeping statement that computers should not make decisions baffles me very much indeed. Whenever I fly into an airport nowadays, I ask myself whether the plane is being landed by a pilot or computer. I hope it is being landed by a computer, but I suspect it is being landed by a pilot. When it affects my personal safety, I want computers to make all decisions they can make better than men.

GREENBERGER. The final question is from David Foster, who is a programming analyst with the General Electric Information Networks Department. He requests comment on Marshall McLuhan's thesis that the growth of information technology, as opposed to mechanical technology, is causing a trend away from narrow specialization toward more generality and synthesis.

DEUTSCH. My impression from the history of science is that periods of cross-disciplinary work tend to alternate with periods of advance in special disciplines. I do not think the coming of computers will change that. To

insist only on cross-disciplinary activity might require that every specialty slow down its advance to the speed of the rest of the regiment. This is the demand of the humanist who would have everything be intelligible to every well-educated man.

Within a specialization, on the other hand, the problems that can be solved with the intellectual resources of that field alone tend to get solved relatively soon, and then the field becomes stagnant until somebody brings in additional information from somewhere else. Then new questions get asked, new resources are applied, and there is another creative period, either in the old field that has somewhat changed or in a new field at its border.

The alternation between specialization and cross-disciplinary work reminds me of the process of breathing out and breathing in. They are part of the same long-term production cycle of knowledge, even though they may extend over generations of scientists.

DINNER DISCUSSION

GREENBERGER. A question that I know is bothering some people concerns the title of Simon's paper, "Designing Organizations for an Information-Rich World." Did you ever really get to the subject of designing organizations?

SIMON. Well, it is a little like Shubik's story of the owl and the centipede. I stated the general principles. Anybody who wants to, can apply them.

Seriously, we often ask what information the decision maker in an organization should have in order to make decisions, without considering his limitations as an information processor and the kind of system he is part of. In postulating attention as a scarce resource, I do think I was dealing with a central problem of organizational design.

PECHMAN. Lyndon Johnson, according to the newspapers, got himself into a position where he did not receive certain information on developments in Vietnam. How does an executive ensure that information he does not like to see will still filter through?

Participants: Andrew Aines, Office of Science and Technology; John Buckley, Office of Science and Technology; Karl Deutsch, Harvard University; Lincoln Gordon, The Johns Hopkins University; Nicholas Johnson, Federal Communications Commission; Anthony G. Oettinger, Harvard University; Joseph A. Pechman, The Brookings Institution; John Platt, University of Michigan; Martin Shubik, Yale University; Leonard Silk, New York Times (formerly at The Brookings Institution); Herbert A. Simon, Carnegie-Mellon University; Joseph Weizenbaum, Massachusetts Institute of Technology; Martin Greenberger, chairman, The Johns Hopkins University.

SIMON. I do not know any way we can get the President of the United States to accept information that he really does not want. But he can organize himself so that he has at least one information channel from each of the different points of view on a question. That does not eliminate the problem of getting information in usable form, given the very limited time at his disposal. He cannot make the day longer than twenty-four hours.

GORDON. I suspect Julius Caesar also may have suffered from listening to only what he wanted to hear. Eisenhower was a classic example of the screened President, but Lyndon Johnson was not. I saw something of his working methods firsthand and had a strong impression that a great deal of information got to him. Johnson used to take home a vast stack of reading material. He did not need much sleep, and would wake up at 5 A.M. for two consecutive, uninterrupted hours of reading. Whenever one of my memos got to him, he acted on it the same or the next day.

JOHNSON. In my own personal experience with Lyndon Johnson, I found him to read widely and seek advice wisely. He would deliberately set up debating societies before him (in effect) to hear all points of view. He also read outside normal channels. This is what I do in my own life, and I suspect most of us do. I sample a wide spectrum of material so as not to become a prisoner of my own screening system.

OETTINGER. By coincidence, Simon's paper and mine dovetail very well,⁹ although I happen to have a strong disagreement with him on the desirability of programming chess. Simon has offered three very deep, important, fundamental principles that shed light on things I had not perceived clearly:

1. attention is a scarce commodity;
2. information technology allows effort to be displaced from possession, storage, and accumulation of information to its processing, even if the information is located in the world itself rather than in the file;
3. filtering and organizing the environment for persons whose attention is scarce are critical.

It remains for others to apply these general principles to particular organizations and explore their political and economic implications.

WEIZENBAUM. There is a student at MIT currently working on what it means for a computer to "know." He took a short story about Mary and Jane being invited to a party. Mommy said they could go but their three-year-old sister could not. Mary and Jane are seven and nine years old. The little sister wants to go, and she cries. Mommy tries to comfort her. That is the story.

What the student is trying to do is get the computer to "understand" the story. He is trying to write down all of the knowledge that the four characters in the story must have to operate in their tiny little framework—knowledge about children, parties, and so on. This has already taken about seventy pages. It is very difficult.

A short time ago Joel Moses wrote a Ph.D. thesis at MIT on symbolic

integration.⁹ His program can integrate symbolic expressions probably better than any living mathematician. Now mathematicians from all over the world write to him with difficult integrations for his program; sometimes he finds a mistake in the standard tables. This program has more "knowledge" about symbolic integration than any single person and certainly more than its author. But its intelligence lies within an isolatable compartment of human knowledge. One can draw very strict boundaries around it and say, "So what? It knows all about integration, but what else does it know?" This boundary has to be broken; there are students working on breaking it; but there is a long way to go yet.

As far as organizations are concerned, I am impressed with a problem posed by a vice president of General Electric who complained that he was often put on the spot by not having information when he needed it. An admiral might call him about a steam turbine on an aircraft carrier that was six months late and ask what he was going to do about it. This would be the first time he had even heard of the contract. What he wanted was an early-warning system about everything going on in the General Electric Company that might get to such a point. Well, everything going on in the General Electric Company is far less than everything going on that the President of the United States should know about.

SIMON. That vice president did not really want to know about the problem. He wanted a way of dealing with it. And he thought the way of dealing with it was to know about it. He should have told the admiral that he would call him back.

PLATT. In the days of Socrates, the total number of volumes in the library at Alexandria was on the order of a hundred thousand. There is no evidence that Socrates ever read anything. He probably based his reasoning on a very small fraction of the totality of human knowledge.

The same is true of Aristotle, who really tried to know everything. The total number of book-length manuscripts that he could have read in his lifetime (or by the time he wrote his great encyclopedia) was not more than five thousand. So it is obviously possible to make enormous strides forward in human thinking and organization by methods of filtering known a long time ago.

SILK. In a certain sense, my fifteen years at *Business Week* were spent filtering. Simply defined, the editorial function is filtering, although it can be done badly. A typical Sunday *New York Times*, for example, weighs five and a half pounds and contains enough words to take the human eye four days to read, if it read every one; not a very good filtering performance.

Why does the press not filter better? I think some of its distortions result from the kind of institution it is. The press recruits very differently from universities or large corporations, as a hangover from an earlier period when it had a more localized function. The filtering it does is not based on what are the most important things to know but rather on what is the sexiest

thing to print, although the degree varies from page to page and department to department.

If you take a Jungian conception of personality types, the good editor knows that his public does not consist of just "thinking" types. He is writing for a broad spectrum of personality types and is appealing to them in a gut as well as rational way.

SIMON. I sense that the press recruits people with a deep urge to inform the public. Perhaps if publishers just relaxed and settled on "making a buck" some of the problems would disappear. There is a serious incongruity between trying to be a major source of public information and education and getting out a publication at periodic intervals (like every day). Most of what people should take in (if they really were going to inform themselves) does not occur on a daily basis. In fact, the world is terribly redundant. Much of what happened today could have been predicted from what happened yesterday and did not need to be published.

SHUBIK. I was in Chile at the time of Project Camelot, and I think one of the reasons it blew up was that the Communist organization in Chile has a daily newspaper which delights in good headlines. The story of Project Camelot supplied the Communist newspaper with headlines for about three weeks. Each day the newspaper slapped another section of the Camelot secret report in the center of the front page and wrote a little critique around it. It was a newspaperman's dream—three weeks of free front pages!

GORDON. Justice Holmes disposed of the daily paper nicely when he compared reading it to watching the second hand of a clock!

I want to divert the conversation for a moment to an aspect of Simon's paper which has not been mentioned. As I understood his presentation, one of the elements in designing organizations for an information-rich, communication-poor world is leaving large amounts of information either in nature or the minds of friends, to be drawn on when needed.

I have two questions on this point. First, if you really knew that some experiment in nature would produce the particular results needed, have you not in a sense already verified the experiment and gone a step beyond leaving information in nature?

My second question is about the index to information. I suppose a purpose of the educational process should be to instruct students in the development of their own indices. Obviously, it should not be to put vast amounts of information into their heads for retention there indefinitely. But how is the person in an organization sure that he has the right index? This seems to me to be a very important design problem.

SIMON. Your first question centers in a technical sense on a point of logic. I do not agree that designing a question which will elicit a "yes" or "no" answer is the same as knowing whether the answer is "yes" or "no." I may understand how to pose the question without yet knowing the answer.

There are many situations where we do not want to pose the question until we need to know the answer. The best example of this is in computer technology itself. Howard Aiken's Mark I computer at Harvard University was constructed in large part because people felt they needed better mathematical tables. But after thinking about it a while, they realized they did not want the tables at all. What they really wanted were subroutines and programming languages that allowed them to get entries in the tables on demand.

On the second question, we may not have to know what the best way is but just the change required from what we do now. We need to worry much less about storing facts in people and much more about storing indexes in them. I do not know the best way to do it, but I do know that we should give people better capabilities for moving around the world, acquiring information.

GREENBERGER. Karl Deutsch has been politely silent. May we have your thoughts?

DEUTSCH. Simon dealt with attention as a scarce commodity but did not explore its purpose or function. What *goals* does an organization seek? What state of affairs is it trying to preserve? What will get a President reelected? What will keep a university happy and working; what will make it blow up? The answers to such questions would help determine the objectively relevant information input.

Relevance depends on goals and needs—functional requirements. Goals and needs, in turn, depend on organizational structure. In order to interpret relevance, we must think about knowledge not only as access to information but as the entire cycle of obtaining, storing, and processing information.

We have to know some facts in order to derive general rules from which more facts can be derived. With too few facts, we cannot derive such rules; if we did derive them we would not know what we had. Our capacities for recognition depend upon our stock of memories. We need a basis of memories of facts to recognize new information and even our own ideas.

I have reservations about Simon's enthusiasm, in the name of simplification and economy of thought, for throwing out vast amounts of what universities now teach. Much of what we learn in social science used to be interpreted against our knowledge of history. If we throw out too much historical data, many of our abstractions may lose meaning. A critical design problem for education is to determine the amount of memories from the past needed for producing and interpreting new information.

If we could build general models of the *expected distribution* of outcomes, we could then pay special attention to events falling outside the distribution. We would not expect a single outcome from a process but (for example) results fitting a bell-shaped curve. If some results fall outside the curve, we might suspect our image of the process and turn up the magnification to examine the reason for the deviation.

This use of expected distributions could be applied to surveillance or early-warning systems. Students will always gripe, and there will always be some campus conflicts and even attempts at suicide. But if too many of these things happen (outside the expected distribution curve), we might become alarmed soon enough to make institutional or organizational changes. (We might also find that the frequency of tragedies or sufferings accepted as normal was, in fact, incompatible with our values.)

Simon warned against excessive fear of unforeseen consequences. He feels experience with DDT may have been the cheapest way of learning about its dangers. This leads to the problem of which warnings to take seriously. Among the many Cassandra calls, which ones are worth heeding?

Statistical background data can help decide such matters. With regard to population explosion, we can now find out the number of people and the countries involved. Do the increases really take place as predicted? What is happening to the food supply? How fast are human reproduction habits changing, and under what conditions? Where, when, and to what extent is there a real danger?

We may not have enough information, on the other hand, for assessing the danger to the atmosphere of the CO₂ or greenhouse effect, which allegedly could change the temperature balance of our planet. The urgency of this danger is therefore presumably less. A lack of knowledge increases the risk of error but does not make it impossible to judge.

When factual knowledge and predictions are unclear, we must fall back on ethics. I think ethics is essentially a set of rules on where the *burden of proof* belongs. If the evidence is incomplete or dubious in a criminal case, Anglo-American ethics says the defendant is innocent until proven guilty. Other legal systems, from France to Russia, treat him as guilty until he proves his innocence. If we must err, on which side would we rather err?

I agree with Simon on the need for cheerful experimentation whenever the value of new experience exceeds the risk of unexpected damage. If irreversible damage results, however, such as when people get killed, we need something better and safer than discovering the consequences by experience. To be sure, there are cases when nonaction can kill more people than action. In 1939 nonaction against Hitler killed more people than action would have killed. But in such cases the evidence should be very strong before irreversible action is taken. As Edmund Burke said, the statesman should be in nothing so economical as in the production of evil.

When we take action, can we make it self-correcting? Can we set up continuous feedbacks to correct our behavior again in the light of its results? Can we make sure that it is not disloyal to discover the action was wrong? We may have to shift Simon's priorities. Instead of going ahead now and learning from experience, instead of mainly seeking to conserve attention, it may be better to stress the continuous processing of information in self-correcting feedback systems.

PLATT. It is just possible that, in the course of correcting our present instability, we might freeze ourselves into a tightly integrated system where a small error will speedily propagate like an East Coast power failure. We need to be very careful to leave lots of looseness, diversity, and lax coupling in the system, even while trying to stabilize it.

AINES. Gentlemen, we actually live in a strangely different world from the one being discussed here. It is a world in which we observe the proliferation of information systems in science and technology. Some are manual, some are computerized, but all are growing.

Sometimes they appear not to be based upon user needs or demands, but on a desire of (for example) the American Chemical Society to refurbish its information programs and create new and more efficient ways of communication. Sometimes they grow because of international competition. There is a tremendous current of international ferment visible in the development of information systems.

In the United States we lack planning and policy entities to guide this development. Yet the proliferation continues resolutely, sometimes disregarding logic and evident need. There are already about five hundred systems in the fields of astronomy, behavioral and social sciences, biological sciences, chemistry, environmental science and related technology, electronics, electrical engineering, and medical and health sciences. These systems are beginning to function. Some are federal, some national, and some international.

What I hoped I would hear discussed (and this in no way faults what has been discussed) is that we must begin to look at some of the organizational problems related to these developments. I see no such activity, even though we are seeking very hard to stimulate interest. The information systems, though growing like crystals in a favorable solution, do not necessarily appear as an integrated or harmonious array. The resulting duplications and inefficiencies in the long run may be terribly costly to society.

BUCKLEY. I cannot see these information systems affecting any real decisions, although I look at the government from the same office that Aines does. But then, I deal with the President and a number of Cabinet officers specifically on environmental problems and have a very limited view of how decisions are made, except in a pragmatic way.

It seems to me that the screens for information turn out to be very largely human ones, carefully arrayed so that there is a diversity of inputs and backgrounds. The pesticide problem did not get the attention of the President and Cabinet, despite two agency reports and a staff input (with different screens applied), until it was covered by the *New York Times* and *Washington Post*. The President commented that there was no way to keep from the top level of government those things that appear in the daily paper. Given these realities, I am afraid what Aines talks about may well lead to a waste of funds.

AINES. Let me focus on the environmental quality area. Some of us feel that the many information activities that operate in this area do not communicate effectively with each other. There is no such thing as a data system for environmental quality. It seems to me that pollution cannot be controlled until we establish a data baseline so we know specifically what we are trying to improve.

The same comment applies to the urban area. People in urban-renewal research and other programs are active all over the country, and many of them do not know what the other chap is doing. The information systems at city, state, and federal levels are relics of a previous age. We need something better.

BUCKLEY. What I feel I need more than anything else is a link in the Bell Telephone System. I have a very good list of telephone numbers, both within the government and outside. I sincerely feel that in less than five phone calls I can get the best piece of information on any subject available any place in the world. I do not worry about not having all the data. I do not have the capacity to deal with it. But I do have the capacity to find it when I need it. Therein lies my utility to the President and to Dr. DuBridge.

GREENBERGER. Nicholas Johnson, may we have your thoughts?

JOHNSON. I will make four points. First, I think we should emphasize more in our thinking the absolutely crucial importance of television as an information medium in our society. Next to the hours they spend at work and in bed, American people spend most time watching television. The average man of sixty-five will someday have spent nine full years of his life, twenty-four hours a day, three hundred sixty-five days a year watching television.

As one example, *TV Guide* has the greatest circulation of any magazine in the United States (even when we read, we read about television). An article that I wrote for *TV Guide* produced fifty to seventy-five letters.¹⁰ In contrast, an appearance that I made on "Face the Nation" (which was nationally advertised to appear half an hour after the actual time of broadcast, thus assuring that those interested would miss it) produced twenty times as much mail.

More than 60 percent of the American people say they get most of their information and opinions from television. Our society does have an information system for adult and child education—television. It is not doing very well, but it is there. I think we cannot make a serious effort to address society's information problem without considering television and the totality of its impact.

The second point I want to make is about pricing. In setting telephone rates, my operating theory is that we ought to make communications as cheap as possible to give people potential access to a maximum of information. We need to develop better devices for selection, but if we can make a Xerox copy or a long-distance call cheaper than we can now, we should do it.

Deutsch's water-cooler example has its analogue in the local telephone

system where the incremental cost of placing a call is zero. We make local telephone calls not on the basis of their cost but rather on the basis of whom we want to talk to and what else we have to do with our time. I see no reason why long-distance service should not operate on the same principle. With domestic satellites all calls go forty-four thousand miles—twenty-two thousand up and twenty-two thousand back down. It makes no difference whether the two ground stations are a thousand miles or thirteen thousand miles apart. Why base price on distance in this kind of system?

Next, I think Aines' effort to speak in specific terms is constructive. Let me offer one example, which I call the personalized journal. Many government agencies now have a morning clip service, which is an effort to survey and select from a large number of magazines and newspapers for a particular specialty. The Federal Communications Commission has a very useful service on communications items which provides me with input from many sources on a regular basis.

Bell Laboratories, IBM, and other companies have selective dissemination systems which make selections according to the user's interest, whether it be by author, subject matter, or journal. Xerox copies of the relevant items are delivered to the user on a weekly basis. I think we are going to see a great deal more of this. There is no reason why it cannot be extended with added technology to closed-circuit television presentations for executives in the morning. It could include sections from books and short courses that the user wishes to view.

My last point concerns calling up people for information. I do it, too, and agree it is now the most efficient information retrieval system. But we must keep in mind the distinction between things we call others for and things for which they call us. This imposes an obligation on us to keep some information of our own on hand, at least if we happen to be at the working staff or executive level. Otherwise, suddenly one day everyone will be calling everyone else, and no one will know anything. That will be great only for the telephone company.

PLATT. I am curious if this procedure leads to closed loops: someone calling you in order to find out something about which you asked someone else.

SHUBIK. I am reminded of the old psychiatrist and the young psychiatrist who are going down in the elevator together. The young psychiatrist is all haggard after a dreadfully tough day, yet he knows that the old psychiatrist (who is completely composed) had four times as many patients. He says, "Doctor, how do you do it?" And the old psychiatrist replies, "Who listens?"

SIMON. This story almost makes my point. We have developed all kinds of information sources: the systems Aines was discussing, the free telephone that Johnson proposes, and so on. Now it is time for us to shift our attention to the people at the receiving end and ask how they really filter all of this information.

[*Added by Simon during editing.*] On rereading the discussion and questions, I am struck with how strongly they reflect the prevailing mood of distrust of technology and of panic in the face of contemporary problems. I cannot share that mood and must reaffirm the optimism of my paper, based on some premises that seem to me supportable by good empirical evidence and logic.

First, while technology demonstrably generates some problems, and these problems have to be dealt with (using that same technology!), technology is man's one best and only hope to escape from the curse of Adam. We need more technology, not less.

Second, the information overload is in the mind of the reader. Information does not have to be processed just because it is there. Filtering by intelligent programs is the main part of the answer.

Third, inaction is also action, and experimentation on the real world is not as risky as it sounds, at least no more risky than that form of experimentation which consists of doing nothing new or different until all the facts are in. Life requires us to balance risks; it does not permit us to avoid them altogether. Moreover, it is easy to exaggerate how irreversible our experiments on nature are. I find it hard to come by genuine examples of important irreversibility.

Fourth, most science fiction about Big Brother is science fiction precisely because it ignores Big Brother's information overload. Lack of information, real or manufactured, has never been the limiting factor on the operations of political police, and I see no reason to believe that the availability of television or computers tilts the balance in their favor.

REFERENCES

1. John W. Gardner, "What Kind of Society Do We Want?", *Reader's Digest*, September 1969.
2. National Academy of Sciences, *Technology: Processes of Assessment and Choice*, Report to the Committee on Science and Astronautics, U.S. House of Representatives, Government Printing Office, July 1969.
3. Edward A. Feigenbaum and Julian Feldman, *Computers and Thought*, McGraw-Hill, 1963.
4. Claude Shannon, *Mathematical Theory of Communication*, University of Illinois Press, 1949.
5. John G. Kemeny, "Large Time-Sharing Networks," this volume.
6. National Academy of Sciences, *Technology*; Herbert A. Simon, *The Shape of Automation*, Harper & Row, 1965; Herbert A. Simon, *The Sciences of the Artificial*, MIT Press, 1969.
7. John Platt, "What We Must Do," *Science*, 166 (1969), 1115-1121.
8. Anthony G. Oettinger, "Communications in the National Decision-Making Process," this volume.
9. Joel Moses, "Symbolic Integration," MAC-TR 47 (thesis), Project MAC, Massachusetts Institute of Technology, AD662666, December 1967.
10. Nicholas Johnson, "The Silent Screen," *TV Guide*, July 5, 1969, pp. 6-13.