In M. Greenberger (Ed.), Computers, communications, and the public interest. Baltimore, 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 untrammeled 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

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 CO_2 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, "Compunications 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.