California Management Review

HERBERT A. SIMON

Programs as Factors of Production

California Management Review, vol. 10, 2, Winter 1967

As a member of that rapidly growing happy band who spend our days trying to find ways of inducing computers to do interesting new things, I feel a good deal of responsibility for understanding the probable economic and social consequences of introducing these new devices into our society and widening the range of their applications. From an economic and social standpoint, are computers and automation something new under the sun or are they, to paraphrase Clausewitz, simply "a continuation of the Industrial Revolution by other means"? Do they call for a new chapter in the economic textbooks or are they merely details in the chapters on capital and distribution?

In a series of essays published under the title of *The Shape of Automation*,[1] I explored some of the macroeconomic aspects of these questions—in particular, the implications of automation for full employment, real wages, and the demand for unskilled labor. I will not repeat my conclusions here except to observe that I ended my investigation with more optimism than I began with. From the standpoints considered, automation does, indeed, seem a natural continuation of the Industrial Revolution, fully compatible with full employment, creating a high probability that labor will reap all or most of the gains of rising productivity.

The present essay is an inquiry in a different direction. It begins with the observation that what we generally call a "computer" includes both a hardware component and a software component—both a collection of electronic gear and a collection of programs and data that are stored in the hardware memory. During the first five years that computers were obtainable commercially, to buy or rent a computer meant to buy or rent the hardware. Since that time, during the last decade or so, the merchandise on the market has generally included a substantial software as well as hardware component. In fact, the software represents a steadily increasing part of the total purchase or rental cost, and no computer of any size could be marketed today without being accompanied by appropriate software. Categories of software include:

- Monitor programs and scheduling algorithms, to make the system available to a multitude of users and to allocate and schedule its facilities among them.
- Programming languages, such as assembly languages and user-oriented languages (e.g., FAP, FORTRAN, COBOL, SIMSCRIPT).
- Utility routines, such as linear programming algorithms, standard statistical packages, and programs for solving differential equations.

For the most part, computer users either employ these precooked programs, or write their programs in user-oriented languages. Few users program in machine language, and almost none in large installations run their programs outside monitor and scheduling systems. Some current suggestions that, because of economies of scale, large central computing hardware systems will take on the character of public utilities will, if realized, accentuate and hasten these trends.

If computers, regarded as a factor of production, are to be classified with capital, they are capital with a difference. To be sure, there are precursors, such as the Jacquard loom, which was as truly programmed as the most modern solid-state machine.

Mr. Simon is Richard King Mellon Professor of Computer Sciences and Psychology in the Graduate School of Industrial Administration, Carnegie Institute of Technology.

But the software component of modern computers is so prominent in comparison with anything that went before that we must treat the difference as having qualitative significance.

An alternative to regarding computers as a new form of capital is to regard them as **a new form of labor.** This, too, is a familiar idea used to denote their applicability to a widening range of humanoid tasks. I would like to introduce an allegory for exploring further the reasonableness and limitations of regarding computers as labor. Like most allegories, this one will simplify real life.[2]

In the time of Columbus, devices for ocean transportation incorporated both a hardware component —a sailing ship—and a software component—a navigator. A shift in the production function for ocean transportation could result from an improvement either in sailing ships or in the skills of navigators. Improvements in hardware were incorporated in the production function as new ships were launched. I suppose that the economics of the matter were handled by using Terborgh-like replacement formulas to determine when ships had become obsolete and should be replaced by improved ones.

Improved navigational technology, however, could be incorporated in the production function either by replacing navigators or by retraining the present ones. At a cost, an experienced navigator might be trained to use a magnetic compass to find north or, some centuries later, a chronometer to determine his longitude.

One important **difference,** then, between the hardware and software components in ocean navigation lay in the greater opportunities for revising and improving the latter without complete replacement. One important **similarity** was that, even in the absence of technological change, both hardware and software gradually wore out and had to be replaced anyway. In both cases, the replacement cost was not negligible. Ships were obviously costly to build, and navigators could only be produced by years of training and experience.

Let us now introduce an automated navigator into this technology, in the form of a programmable computer. Only one thing has changed in the economic structure of the situation, but a rather significant one thing: Many technological changes in the art of navigation can now be introduced, almost without cost, by replacing the present program in the automated navigator with a copy of a program incorporating the improved method. The automated technology is an example of a technology that can be copied almost without cost.[3] To understand the significance of the difference, we must consider the economics of copying.

The significance of cheap copying processes is that when they are available, the cost of developing improvements need be paid only once. Darwinian evolution is as much a matter of multiplication of the fittest as it is survival of the fittest. Genetic material, DNA and RNA, is organized as a copying mechanism, permitting improved organisms to be multiplied in number at no higher cost than would be required to produce the old, unimproved ones. Like improvements in the ship, however, and unlike improvements in the navigator, the superior technology cannot be introduced into existing organisms but must wait until they are removed by obsolescence and wearing out. As a matter of fact, copies continue to be made of the unimproved as well as of the improved organisms until competition gradually weeds out the former.

A second instructive example of a copying process is the one used by animals, but especially by man to transmit culture from one generation to the next by training, instructing, and educating progeny. This particular copying process can hardly be called "cheap," since in human cultures it commonly occupies a span of years nearly as long as the period during which the adult is fully productive. It is worthwhile—even necessary—precisely because man has reprogrammable software and can improve on the programs that he copies genetically from his parents. We might even say that these latter programs provide him with little more than a monitor system and an assembly language—that is, capabilities for acquiring performance programs.

The peculiarity of this particular copying technique and the reason why it is costly is that the program to be copied cannot be inserted directly into the human head, but must gradually be grown there by those poorly understood processes we call education and experience. Copying is by no means synonymous, by the way, with memorization: you can memorize a page of Hoyle to the point where you can recite it perfectly, without being able at all to play the game described there. As a result, also, of the indirectness of the copying process, an exact copy is almost certainly never produced. In general, we can expect degradation in copying, the quality of the program being restored only by new improvements after experience in use.[4]

In order for knowledge and skills to be transmitted from one generation to another, they have to be stored reliably by memory. Until about five hundred years ago, the two major storage depositories were human memory and man's artifacts. Although writing has been known, of course, for some thousands of years, it was used to only a very limited extent to store the information needed to transmit skills from one generation to the next. One reason, undoubtedly, was the high cost of providing children with the programs (i.e., reading skills) needed to retrieve information from this memory source. A second was lack of knowledge about how to communicate "how-to" information in words generally, and in writing in particular. A third, and the most obvious, was the high cost of producing copies so that the information would be widely available.

Artifacts were a more interesting memory device. Houses were not built to teach people how to build houses, but were built to shelter them. Nevertheless, the houses were always present and available for inspection as a source of information on how to build houses. Hence, the existence of houses greatly reduced the cost of copying houses (as compared, say, with what the cost would have been if the builders only had oral or written descriptions of houses, or even pictures).

Five hundred years ago, the invention of printing greatly reduced the cost of copying verbal information, as well as the cost of copying pictures and diagrams. It provided an important new cheap, reliable memory device. For a long time, however, it appeared to be used far more for remembering abstractions and intangibles than the concrete particulars of everyday life and its technology. I would conjecture that these concrete particulars were already recorded in artifacts better than they could be in the verbal expressions (oral or written) available at that time.

Thus, among the crucial events in human evolution have been the introduction of five important advances in the technique and copying and storing information: organismic reproduction with duplication of genetic material; indirect programming through learning; preservation of artifacts; writing; and printing. Each of these has its characteristic structure of costs. None of them allows "instant" reprogramming of existing hardware. The third and fourth are simply storage devices; the fifth allows cheap copying; but all these last three contribute to production only through the second—learning—hence do not avoid its costs.

Direct copying of computer software has characteristics, therefore, quite different from any of these earlier copying techniques. When an improved program has been invented for the automatic navigator, it can not only be installed in new navigators at no addition to cost, but it can also be provided to existing navigators instantly and substantially without cost. The comparative advantage of automatic navigators relative to human navigators will increase in any field where the technology is advancing rapidly, since improvements can be incorporated in the former sooner than in the latter. The economics of the matter are developed more fully in the Appendix.

In the case of any copyable technique, there is a problem of how the costs of developing improvements are to be recovered. In the absence of adequate opportunity for recovery, there will, of course, be underinvestment in research and development. In a competitive economy, the problem becomes the more severe the less expensive and the more rapid the copying process.

Patent and copyright laws are the usual modern procedures for returning rewards to the authors of technological advances that can be copied. Secrecy is another method still widely used, but not applicable when the improvement can be copied from artifacts embodying it. (Study of artifacts may permit copying not only the objects themselves, but even improved methods of manufacture, evidences of which are preserved in the manufactured object.) At an earlier point in history (and even today in the military sphere), governments intervened to prevent the export of technological improvements, whether in the form of machinery or of programs recorded in the memories of artisans.

Because of the cheapness of the copying process, and the potential value of even single exemplars, the protection of inventors' interests in improved computer programs is a matter of great technical difficulty. Here we must distinguish programs written to run on a particular type of machine, on the one hand, from programs written in higher-level languages that are easily transferred to different machines. The machine manufacturer can recover investments in software developments of the former kind, since they can only be used with his machines. On the other hand, the improvements are then not used everywhere they might be, and competing manufacturers must duplicate development investments, both sources of misallocation of resources.

With progress in software technology, programs have tended to become more independent of hardware. Hence, the problem from a social point of view appears to be to secure a sufficiently high rate of investment in software development.

Labor's contribution to production is achieved by the coordination of a system of sensory organs—eyes and ears—with a system of effectors—principally hand and mouth—by means of those stored programs we call "skills."[5] In our Columbian ocean transport technology there are both the skills of the navigator—whether human or automated—and the skills, or technical know-how, of the shipbuilders. Evidence from wartime destruction shows that an economy that has lost most of its physical capital but retained its pool of technology can restore previous levels of productivity relatively rapidly.

As was pointed out earlier, the stored skills in a pre-computer economy must be replaced each generation, even if there is no technological change, since these skills must exist in human brains in order to be useful for production. The replacement costs are by no means the same as the costs of formal education in the economy. In the first place, the entire time, not just the school time, of children prior to their entrance into the labor market should be charged as part of the replacement cost. Whether in school or out, children learn to speak, become acquainted with the common artifacts of their culture, and, at least in simple economies, learn one or more relevant production technologies.

In the second place, in societies with formal educational systems, a large part of what is taught and learned in the schools has no productive significance. This is certainly true of most of the curriculum of the contemporary American school system. Schooling is best regarded, under such circumstances, primarily as a consumption good that has as a small by-product the storage of a certain amount of production skill.[6] If I were given a contract, at a fixed price, to produce research scientists, I would certainly turn the finished product out of my educational institution at an age earlier than 25!

In the third place, much of the transmission of programs takes place through on-the-job training and experience. These training costs do not show up in the social accounts, but are hidden as direct costs of production.

In a peasant culture, the avoidable cost of replacing programs each generation is probably very small, because by the time children are physically capable of doing hard manual work they have already learned most of the skills they will use. On the other hand, the absence of mechanisms for cheap transmission and reliable storage of programs probably operates to slow technical progress or even cause the loss of discoveries, so that improvements need to be repeatedly re-invented.

We observe, for example, that there was only minor technological advance in peasant cultures from pre-Christian times to the Industrial Revolution. We may conjecture that the technologies of these cultures remained in a state of dynamic equilibrium—they were able to maintain just that level of technology at which the forgetting from one generation to the next was balanced by re-invention. Increases in the density of population and improvements in the security or economy of travel and transportation would allow increases in specialization, hence permit a larger stock of programs to be transmitted. Nomadism, on the other hand, through increasing the difficulty of retaining numerous physical possessions, would decrease the stock of the culture's artifacts, hence degrade an important store of technological information.

It is not my purpose here to rewrite cultural history in terms of the problems a society faces in maintaining and transmitting stored programs. The notion of a learning-forgetting equilibrium of technology is not relevant, however, only to peasant cultures. It is equally applicable, for example, to the problem that a university department faces in remembering its own policies and all of the subtle considerations that went into their formulation. As faculty come and go, secretaries marry and are replaced, files are lost, and conditions change, the subtleties vanish, and the policy tends to reduce to a few general (though not necessarily sound) principles, plus some specific regulations that happen to have been recorded in documents that continue to be referred to. Often, the documents themselves become inoperative because the "retrieval" programs that would cause them to be referred to on appropriate occasions are lost.

In sum, the costs of maintaining a store of human programs over periods during which personnel turnover is substantial are very large. Storing technology in the form of computer programs rather than human programs opens up new possibilities for greatly reducing such costs.

Learning is needed not only to transmit programs from generation to generation and to modify programs to incorporate new technology, but also simply to adapt programs to problems posed by a constantly changing environment. The economic gains to be realized from cheap copying will be inconsequential if new programs have to be devised ad hoc for each small change in circumstances.

The program of the navigator may make use of large amounts of information about tides, currents, winds, coastlines, and harbors in different parts of the world. The program, to be workable, must be factorable into two parts:

- A store of information that can be augmented readily by simple processes of memorization and by simple programs for consulting available reference sources.
- A general purpose program that can apply appropriate parts of this data store to any given specific navigational problem.

The economy of automated programs will depend on their having comparable features. They must be learning programs, at least in the sense that they can apply new information to new situations and probably also in the stronger sense that they are capable of some adaptive modification in their own structures. Without such features, each program would be applicable to only a narrow range of situations; hence little would be gained from the availability of cheap means for producing copies.

The ocean transportation technology of our allegory depends both on the programs stored in the navigator and on the programs for the manufacture of ships. All of our discussion of the costs of programming and improving the programs of human or automatic navigators applies quite as well to the programs of human or automated shipwrights.

Technological advance requires the invention of new techniques, but also the development and storage of the programs necessary to apply these techniques. For any extensive technological change, a whole series of "reprogramming" decisions have to be made—by managers, engineers, and workmen. In evaluating these decisions, numerous externalities will be encountered, because the effectiveness of the new technology in comparison with the old will depend on what programs have already been stored. Thus, the productivity of capital in the form of automobiles will depend on the commonness of driving and mechanics' programs in the population, as well as the presence or absence of such material artifacts as roads and gas stations.

One particular difficulty in the diffusion of new technologies is that the new programs have to be ingested, at least in considerable part, in order to evaluate them. Hence, much of the reprogramming cost must be borne before an accurate evaluation can be made—or acceptance of the new technology must be postponed until its advantages are obvious even to the untutored eye. To the extent that the programs of the new technology are computerized, the costs of developing the programs will have to be borne, but not borne anew in each application. We would expect more rapid diffusion of new technologies under these conditions.

Automation of the programs that constitute a technology will make explicit not only the problem of modifying programs to take advantage of advances in knowledge, but also the problem of using existing programs in relevant situations.

Specialization increases the repertory of programs that are available within the economy taken as a whole. It does not guarantee, however, that the sophisticated programs stored in the specialist's brain will be used whenever relevant. Someone, at the point of problem impact, must note the relevance and must have an effective procedure for locating the specialist. Even when he has been located, there may be difficult problems of compatibility between his programs and those of the persons consulting him—what we call now an "interface" problem. The specialist may fail to understand the problem properly, and those consulting him may fail to understand his solution or how to combine his knowledge with aspects of the problem that fall outside his specialty.

There is much talk today about the "knowledge explosion" and how this explosion makes it more difficult to locate relevant knowledge. Much of this alarm is ill-considered, for the advance of knowledge is not primarily an additive (or multiplicative) accumulation of knowledge. It is primarily the reorganization of knowledge to make it more parsimonious and more applicable. To become a research chemist should involve less learning today than it did fifty years ago, because physical chemistry and quantum mechanics have provided such powerful tools for organizing facts, and indeed making them derivable from theory.

In this age, as in any other, an important part of the programs that define the technology are programs for retrieving knowledge from its storage places. Among the important prospective consequences of automation are its consequences for retrieval techniques.

The cost of retrieving relevant programs depends, first, on finding them; second, on making them usable in the application situation. Finding costs depends, in turn, on the structure of available indexes and on the power of the available search programs. As illustration, consider the boxed problem.

In spite of the apparent simplicity of a scheme for retrieving specialized information, retrieval in our present technology is by no means a perfected art. A good example of the inadequacies is provided by the lag between the level of sophistication of the statistical techniques applied by data users and the level of sophistication among experts in statistics. The more sophisticated programs are not retrieved when they would be appropriate because

- The user is not aware of their potential relevance.
- His access to the existing knowledge through appropriate inquiry procedures has not been institutionalized. (Among other things, he may have no way to recompense the expert properly for his time and trouble on a problem that is only a matter of "application," hence not of direct professional interest to the technique-oriented expert.)

Automation of technologies will cause the problem of retrieving relevant programs from the stock of existing ones to become more explicit than it has been in the past. The computer technology—both hardware and software—will also provide new means for retrieval. It will also create interesting new problems for economists, relating to the design of efficient retrieval systems.

One question that will arise repeatedly is the question of how far information should be processed when stored, and to what extent, on the other hand, it should be processed on demand. Should executives, for example, have "instantly" available the answers to large numbers of questions they might conceivably ask, should they have available the programs that will seek out and compute the answers in a short time, or should they have available programming languages that will allow them to write programs that, in turn, will find the answers? The cost structures of automated information systems are so different from those of manual systems that all of these issues will have to be rethought as the new technology develops.

Conclusions. A technology exists largely in the minds of its labor force and in the future will be distributed between those minds and the memories of computers. If programs, stored in one or the other of these forms, constitute the core of a tech-

How Does the National Academy of Sciences Retrieve Information?

One of the functions of the National Academy of Sciences is to provide agencies of the national government with expert scientific advice when needed. What kind of processing system is required for the National Academy to perform this function?

We can dream up designs for elaborate indexes of specialists. One such index, the National Register, actually exists. However, there is a much simpler and fully adequate device available to the staff of the National Academy: the telephone. Any given inquiry can be roughly classified by the field of knowledge to which it belongs. The appropriate members of the Academy staff can carry around in memory indexes to the names of a few persons who are knowledgeable in each of these fields. Each of the knowledgeable respondents, in turn, has a more detailed map of specialties within the general area and an index of names of the corresponding specialists. These specialists will have their own indexes, and so on. A series of three or four phone calls can hardly fail to locate the best program in the United States to address itself to the problem at hand.

When we look at this retrieval scheme in detail to determine what would be involved in automating it, we see that there is nothing very complicated about the processing. Each of the memories employed contains, among other things, a taxonomy. We may visualize each choice point in the classification key as containing questions to be answered in order to make the choice (i.e., questions the specialized informant can ask in order to pin down the nature of the inquiry).

nology, then important consequences are likely to follow from the fact that automation greatly decreases the cost of making copies of such programs.

One of the obvious consequences of cheaper copying is that there will be underinvestment in program improvement unless steps are taken to reward inventors of programs or to subsidize invention. A second consequence is that the comparative advantage of automation will tend to be particularly great in situations where frequent and rapid program change is called for, and will tend to be relatively less in areas where only a few copies of a program can be used. Since human programs are at least modestly capable of on-the-job learning and adaptation to specific situations, the range of feasible automation will depend heavily on the extent to which similar learning and adaptive features can be incorporated in automated programs.

The concept of technology as consisting of stored programs gives us a somewhat novel framework for theory about the rate of technological progress and the rate of diffusion of new technology. The level of technology that a society can maintain will depend heavily—indeed, in the past may **have** depended heavily—on the costs of transmitting programs from each generation to the next. It will depend also on the possibility of economizing transmission costs through specialization.

To the extent that there is specialization of programs within an economy, retrieval programs for locating relevant knowledge and skill become an important element in productive capacity. When copying costs are high, locating a relevant specialist will be useless if his time is fully occupied. With techniques for copying programs cheaply, the numbers of specialist programs will respond flexibly to demands, hence retrieval programs will take on an even greater importance than they have at present.

The automation of programs will have many consequences beyond those I have identified. If there is anything we can say with confidence about a new technology, it is that we will not really understand its implications until we have lived with it for a few generations. Now that we have perhaps achieved some understanding of the First Industrial Revolution—the revolution of power—we are already in the midst of the Second—the revolution in the processing of information. It is important that we identify the salient characteristics of the new technology and their consequences for the economy.

APPENDIX

The argument in the text with respect to the effect of cheap copying on technological change can be made more rigorous by a simple mathematical model. We assume that invention is an autonomous activity, as a result of which there exists at any given time, t, a most efficient technology. This technology takes the form of programs (e.g., programs for navigators and shipwrights) and is implemented by copying these programs and installing them in place of existing programs. (It should be observed that, in contrast to the "learning by doing" theories of Arrow and others, we assume that application of the new technology does not increase the rate of invention. A more complete theory would combine the "learning by doing" mechanism with the copying mechanism discussed here.)

Navigational technology. Let $R(t)$ be the net revenue per year (exclusive of program-copying costs) produced by a ship that employs the best navigational technology available at time t. Let P_n be the cost, assumed constant, of a navigator's program. If navigators' programs are replaced, on the average, every T years, then the average age of the technology in use will be $T/2$. The net revenue will average:

$$\overline{R}(t) = \frac{1}{T}\int_{\tau=0}^{T} R(t-\tau)d\tau. \qquad (1)$$

Assume that invention produces a constant rate of increase in net revenue returnable by the best technology:

$$R(t) = A+Bt. \qquad (2)$$

Then, from (1) and (2):

$$\overline{R}(t) = \frac{1}{T}\int_{\tau=0}^{T}\Big[A+B(t-\tau)\Big]d\tau$$
$$= A+B\left(t-\frac{T}{2}\right) = R\left(t-\frac{T}{2}\right). \qquad (3)$$

The cost of replacing navigators' programs every T years will be, per ship per year:

$$C_n = \frac{P_n}{T}. \qquad (4)$$

Hence, the revenue, net of this cost, will be:

$$\overline{R}(t)-C_n = R(t)-B\frac{T}{2}-\frac{P_n}{T}. \qquad (5)$$

We wish to choose the replacement interval, T, so as to maximize revenue, for given $R(t)$ and P_n. Setting the first derivative of (5) equal to zero, we get:

$$\frac{d(\overline{R}(t)-C_n)}{dT} = -\frac{B}{2}+\frac{P_n}{T^2} = 0. \qquad (6)$$

Whence,

$$T^* = \sqrt{\frac{2P_n}{B}}. \qquad (7)$$

That is to say, the optimal replacement interval for programs, T^*, will vary directly with the square root of copying costs and inversely with the square root of the technological change coefficient. Substituting T^* in the revenue function (5), we get:

$$\overline{R}(t) - C_n = R(t) - 8\sqrt{P_n B}. \qquad (8)$$

The second term on the righthand side of (8) is the penalty for failing to use the best available technology. This penalty is larger the greater the cost of copying programs (because the lag will then be greater) and the more rapid the improvement in technology.

Suppose there are also rental costs associated with navigators, human (C_H), or automatic (C_A). Then if human and automated navigators are just competitive:

$$\overline{R}_H(t) - C_{NH} - C_H = \overline{R}_A(t) - C_{NA} - C_A. \qquad (9)$$

That is:

$$-\sqrt{P_{NH}B} - C_H = -\sqrt{P_{NA}B} - C_A. \qquad (10)$$

Presumably, $P_{NH} > P_{NA}$, and $C_A > C_H$, so that:

$$C_A - C_H = \sqrt{P_{NH}B} - \sqrt{P_{NA}B}$$
$$= (\sqrt{P_{NH}} - \sqrt{P_{NA}})\sqrt{B}. \qquad (11)$$

Now if the rate of invention increases (B larger), the cost balance will be tipped in favor of automated navigators. If human learning can be made more efficient (P_{NH} reduced), the balance will be tipped in favor of human navigators.

Shipbuilding technology. The argument is easily extended to deal with optimal replacement rates for ships and for shipwrights' programs. The average age of the technology embedded in ships will be $\frac{1}{2}(T_W + T_S)$, where T_W is the replacement period for shipwrights' programs and T_S the replacement period for ships. Ignoring navigation, we have, analogously to (3):

$$\overline{R}(t) = R\left(t - \frac{T_W}{2} - \frac{T_S}{2}\right). \qquad (12)$$

Assume that the number of shipwrights required is proportional to the number of ships built. If the improvement in technology is linear, as before, then, for each ship:

$$\overline{R}(t) - \frac{P_S}{T_S} - \frac{P_W}{T_W}$$
$$= R(t) - B\frac{T_S + T_W}{2} - \frac{P_S}{T_S} - \frac{P_W}{T_W}. \qquad (13)$$

Maximizing, we obtain:

$$T_S^* = \sqrt{\frac{2P_S}{B}}, \; T_W^* = \sqrt{\frac{2P_W}{B}}. \qquad (14)$$

REFERENCES

This is an address given in December 1966 to the Industrial Relations Research Association meeting in San Francisco, California.

1. (New York: Harper & Row, 1965). I have been reinforced in the conclusions reached in those essays by noting their agreement with the subsequent report of the National Commission on Technology, Automation, and Economic Progress, *Technology and the American Economy* (Washington, D.C.: United States Government Printing Office, 1966), and the recent book by Charles Silberman, *The Myths of Automation* (New York: Harper & Row, 1966).

2. The technological developments that underlie the possibilities of using human and computer programs interchangeably in a wide range of tasks are discussed in my paper, "Decision Making as an Economic Resource," in Lawrence H. Seltzer, ed., *New Horizons of Economic Progress* (Detroit: Wayne State University Press, 1964), pp. 71–95.

3. *Ibid.*, pp. 91–93.

4. "The Economic Implications of Learning by Doing" are examined in Kenneth Arrow's well-known essay with that title, *Review of Economic Studies*, XXIX (June 1962), 155–173, and in the reference cited there.

5. "Decision Making as an Economic Resource," *op. cit.*, pp. 80–82.

6. My intention here is simply to state a fact, not to offer social criticism. In particular, I do not mean to argue that transmission of production programs should be the sole, or even a major, goal of formal education. I simply observe that it is easy to jump to the conclusion that this is what education is all about. The jumping has been encouraged by studies of American education that have made much of the largely spurious correlation between the amount of education and earnings. The correlations are spurious because they have been uncorrected, or inadequately corrected, for differences in ability, in ambition, and in family status.