2

What Computers Do

So far, we've treated computers and information technology as merely a general kind of investment, like gold stocks or mutual funds, and asked how well it performs financially. But computers are supposed to do particular jobs and to have more specific, concrete effects.

We are interested in how computers affect productivity and how they affect workers. These related issues have two aspects: how computers influence employment and what they do to wages, morale, and the quality of work life. These issues are critically important in their own right, and they form essential parts of the bridge between technology and its economic consequences.

Employment

Total white-collar employment rose strongly through the 1970s and 1980s; between 1972 and 1982, the number of clerks increased by almost 30 percent and the number of managers by over 40 percent. Blue-collar employment grew over the same period by less than 10 percent (Hunt and Hunt 1986).[1] These data could mean that computers fueled a great expansion of activity in the service sector, that they failed to reduce the labor needed, or both. Osterman (1986) tried to partial out the effects by econometric modeling of changes in deployment of computers and employment of labor. The analysis compared earlier and later years that were at comparable points in the business cycle between 1970 and 1978. Overall, he concluded that greater computerization was associated with less growth in clerical and management employment for the average of a variety of manufacturing and nonmanufacturing industries. On the other hand, he found evidence that computers tended to reduce employment in these categories for the first two years after their introduction and then to increase it for the next two years. He says, "The displacement effect seems to be concentrated in the period immediately following the expansion of computing power. Within a reasonably short time, employment expands, presumably in response to the coordinating or bureaucracy effect. The evidence of this pattern is stronger for managers than for clerks, as should be expected since managers are more likely to benefit from the added bureaucracy effect."[2]

When it comes to employment quality, the situation is even less clear. Real wages have not improved as much in recent decades as in earlier eras, a reflection of labor productivity trends as well as forces that have led to upward redistribution of income. Effects of computers on the wages of their users are mixed. Some computerization, such as bank bookkeeping and utility company billing systems, is intended to allow the same work to be done by workers with less skill and lower pay. However, the overall impact of this strategy is hard to assess. In the insurance business, fewer employees did low-paid routine filing and data entry, and more carried out previously difficult functions like deciding who was a good risk (Baran 1987). In many businesses, the typist of yore has been recast as a word processing clerk with greater training requirements and a higher salary.

There is a host of articles and books on the effects of technology, and of computers in particular, on the quality of work life. Two leading authorities sum it up thus: "Research literature on the impact of new information technologies on job content and job satisfaction provide a mass of contradictory findings" (Attewell and Rule 1984). Apparently phase two computerization can improve work, or degrade it; enrich it, or demean it; make it more challenging, interesting, and comfortable, or duller, more repetitive, and stressful. It figures. Computers are infinitely flexible, powerful devices; they can be used in many ways to make work richer or poorer, better or worse, to bring workers either sickness or health.

One careful study of the effects of a narrowly aimed computerization found both good and bad effects. Kraut, Dumais, and Koch (1989) surveyed and observed several hundred service representatives of a large telephone company before and after a new computer-based billing system replaced the old paper and microfiche records. They found that self-reports of work challenge and overall job satisfaction went down, while indicators of mental health and stress—like headaches—improved. The difficulty of the job decreased for the service representatives but increased for supervisors. In addition, the size of all these effects differed dramatically by office. (We will recount more such facts from this methodologically elegant study in chapter 14.)

Individual Firm Experience

What happens in individual firms as they computerize? Unfortunately, private companies are generally reticent with detailed financial data of this kind, especially if they might reveal mistaken business decisions. There are many anecdotal reports and strong opinions in the popular press and the academic management literature as well, many lauding computers strongly, some critical, but almost none with trustworthy quantitative data. Nevertheless, there are a few useful reports.

Franke's study (1987) of the finance industry included a detailed analysis for one large northeastern bank that shared historical data. The econometric equations for this firm showed the same overall pattern as the industry as a whole, with capital returns somewhat more negative than for the average. An interesting detail from this analysis was the finding that capital was even less effective during the years after the bank began installing automatic teller machines (ATMs).

On the positive side of the bank ledger is a report in *The Office* for February 1976, entitled, “We Increased Typing Productivity 340%” (O’Neal 1976). It recounts the experience of Illinois National Bank of Springfield in moving correspondence typing to a word processing center using telephone-accessed dictating machines and electric typewriters with magnetic card storage. With just six people, according to the author, the new center was able to do work previously done by twelve full-time secretaries and twenty half-time clerk-typists, while cutting turnaround time for letters by 70 percent and increasing typing accuracy and letter image quality tremendously.[3] Among other positive reports is a 1979 survey of insurance companies whose respondents reported average typing output gains per operator of 70 to 85 percent (Baran 1987). Examination of the survey questionnaire suggests, however, that the output gains were estimates based on how much typing was done before word processing was introduced and what proportion was boilerplate and retyped sections of documents. The authors apparently assumed that word processing would save all that time and took no account of any of its negative effects.

A report of another extensive and possibly more pertinent questionnaire study came out of Quebec in the mid-1980s (Benoit, Cossette, and Cardillo 1984). One hundred twelve organizations, most of them private service businesses, answered detailed questions about technology and employment changes and gave estimates of how many manual typists would be needed to replace the ones that had been equipped with word processors. The subjective estimates agreed closely with the insurance company answers—an average output increase of 80 percent. Actual figures on the number of secretarial employees before and after introducing the machines were not as impressive: a total *increase* of 19 percent. At the same time, the businesses reported that work flows had gone up an average of 58 percent. Dividing one by the other gives a net efficiency gain of 33 percent. Unfortunately, we're not told the source of the greater work flow—for example, whether it represented more business product, new uses of documents, or merely additional drafts.[4]

These reports seem more in line with the kind of effect we've hoped and supposed IT to have. The stampede of American business to word processing centers would suggest that similar experiences were widespread (although many firms subsequently abandoned them because of undesirable organizational side effects). Considerable caution is needed in interpreting the early reports. As we will see shortly, word processing tends to multiply the number of drafts of each document. If the efficiency gains that people claimed were based on pages of output per operator, a common measurement in typing work, the effects could be an illusion. I suspect that the absence of any more recent, more objective studies reflects dawning industry awareness that overall output gains are much less than promoters had claimed. Still, suppose we take these reports at face value. How much help for white-collar productivity would they imply?

A very small proportion of labor time is actually devoted to typing, and typing jobs are relatively low paid, so efficiency gains of this magnitude, especially when offset by expenses, might not even show up in overall white-collar productivity measures. Roughly 5 percent of nonfarm–nonmanufacturing workers are typists or secretaries (Hunt and Hunt 1986). They spend perhaps 60 percent of their time actually typing. If we add in the typing time of people like me whose hunt and pecking is not in their job descriptions, we still can't get more than 5 percent of white-collar hours at the word processor keyboard. Even if the efficiency of this work had increased for everyone by 80 percent, it would be a once-only productivity gain of only 3.2 percent. Amortized over the twenty years since its introduction, its contribution would have slowed the decline by less than 0.1 percent per year.

If there were real gains of 80 percent across all white-collar tasks, it would be a different matter. Then we would be looking at a compounded productivity growth of 3 percent per year. Unfortunately, it appears that effective phase two applications so far exist for very few tasks, and even the best—realistically including word processing—have rather small impact.

To determine what size of efficiency gain *would be* impressive for an individual application, we can look at the effects of technology in industries that have experienced major productivity gains. One of many late-nineteenth-century advances in cotton thread production resulted in three times as much—a 200 percent increase—in the number of yards produced per worker per hour. Cumulative inventions over a fifty-year period reduced required labor hours per yard by a factor of over 150, for a 37,000 percent improvement (Mokyr 1990). Even during the productivity doldrums of the 1970s and 1980s—the very age of the word processor—inventions for thread spinning and weaving brought 400 percent efficiency gains [Baily and Chakrabarti 1988].) Such technological efficiency enhancers are the main driving force of productivity growth.[5] How do such huge improvements in particular processes translate into overall industry productivity? Not directly, of course, and not nearly on a one-to-one basis. The local component of production that is speeded up has to be imbedded and managed in an overall process that may hedge in and limit its impact. Thus, the textile industry exploited its amazing new technologies for a healthy but “mere” 4.2 percent annual multifactor productivity growth from 1973 to 1985. If IT is to fuel major productivity growth, it will probably need to find hundreds of applications with at least tenfold effects, and the improvements will have to be repeated over and over in succeeding decades or less.

Individual Worker Efficiency

Finally we can ask what happens to individual worker efficiency when phase two computer aids are introduced. The most informative data are results of controlled experiments comparing the efficiency and desirability of computer-assisted work to the same work done without a computer. We would like such evidence for all major applications of phase two computer work aids: text editors, spreadsheets, information storage and retrieval systems, order entry and billing systems, inventory management systems, meeting support software, message systems, desktop publishing programs, graphics drawing programs, computer aids for design, automatic teller machines, point-of-sale devices and so on. Unfortunately, very few such programs have been tested against preexisting work methods.

Before loosing drugs on the public, pharmaceutical companies do extensive controlled tests of efficacy in both the laboratory and carefully monitored clinical trials. But software is not usually perceived as a health threat, so most “testing” of its efficacy is left to the uncertain vagaries of the marketplace, where snake oil sometimes triumphs. Nevertheless, there has been a sufficient number of well-controlled studies of several types of systems to provide an instructive, if not conclusive, sample. The results of such studies are of special interest.

Let’s start with one of the most encouraging—a case study of telephone service representatives and their new billing records system (Kraut, Dumais, and Koch 1989). The researchers did not have access to company productivity records but did get reports of the number of tasks completed per day from the reps themselves, who keep track of their performance as part of their job. With the new system, they reported processing approximately 50 percent more customer interactions each day. The time saving came mostly from being able to pull up records instantly on the screen rather than engaging in the clerk's three Fs, fishing, ficheing, and fetching. We don't know whether they were more effective in making collections or satisfying customers (and there's some reason to believe these aspects could have suffered), but at least we have here a likely case of significant labor efficiency improvement. An important aspect of this finding is that the system studied is only one of a large number of similar systems introduced by telephone companies over the last twenty years, and we will shortly see circumstantial evidence that its achievements were not unusual.

Next to put in evidence are some controlled experiments on the effectiveness of text editing (table 2.1). Gould (1981) enlisted ten office professionals from his organization, IBM research, as subjects. They were all accustomed to writing letters in the traditional way—penciling a draft and passing it to a typist—but were also familiar with a computer-based text editing program. Each wrote four business letters the old way and four with the text editor. The four letters were of different types, ranging from dull sales information transmittals to unpleasant late-payment negotiations. Letters produced with word processing were modified an average of forty-one times, handwritten ones only eight. Other office professionals found no discernible differences in quality. Handwritten letters took an average of twenty-one minutes of the professionals' time for composition and proofreading, plus fourteen minutes of secretaries' time, for a total of thirty-five minutes labor. Text editor letters took an average of only thirty minutes, but all by the professional. If, for the sake of illustration, we assume the professionals' time to be worth $20 per hour and the typists' $10 per hour, the text editor–produced letters cost $9.83 apiece in direct labor costs, handwritten ones $9.44.

Stuart Card of Xerox's Palo Alto Research Center, a laboratory that nurtured many widely acclaimed advances in user-friendly computing, found it difficult to accept Gould's results. Gould, he said, had used an early model of word processor, a so-called line editor, one in which the user has to tell the machine where to make a correction by typing arbitrary "commands" (for example, "7s/hte/the" to correct a misspelled word in the seventh line). Card knew that newer full screen or display editors, in which the user simply positions a mark on the screen to indicate the text that needs changing, were easier to use (Egan and Gomez

Table 2.1
Two experiments on text editors versus handwriting

	Number	Time (min.)	Pay rate	Cost
Gould (1981)				
By hand and typist:				
Composing		19.2	$20/hr.	$6.40
Typing		13.9	$10/hr.	$2.31
Modifications	8.5			
Proofreading		2.2	$20/hr.	$0.73
Total		35.3		**$9.44**
With a computer text editor:				
Composing		29.5	$20/hr.	$9.83
Typing				
Modifications	41.3			
Proofreading				
Total		29.5		**$9.83**
Card, Robert, and Keenan (1984)				
By hand and typist:				
Composing		21.7	$20/hr.	$7.23
Typing		13.9	$10/hr.	$2.31
Modifications	5.0			
Proofreading		2.2	$20/hr.	$0.73
Total		37.8		**$10.27**
With a computer text editor:				
Composing		22.8	$20/hr.	$7.60
Typing				
Modifications	23.5			
Proofreading				
Total		22.8		**$7.60**

1985; Gomez et al. 1983; Roberts and Moran 1983). However, such editors had not, like the one Gould tested, been compared to precomputer technology for the same tasks, so Card performed an experiment of much the same kind as Gould's, using a better, more up-to-date text editing program (Card, Robert, and Keenan 1984). Appropriately to the changing times, participants were habitual and proficient users of text editors. Except for the text editor, the Xerox group copied Gould's experiment in most respects: the same paired sets of four business letters, half by handwriting, half by text editor, and so forth. An expert on English composition did the blind quality ratings.[6]

The new results were barely more encouraging than the old. Again, authors made about five times as many revisions with the text editor, and again there was no appreciable difference in quality. In time spent, there was less disadvantage to the text editor method for Card's authors than for Gould's; they took twenty-three minutes to compose a letter with the computer, hardly any longer than the twenty-two minutes they spent handwriting one. Card did not report transcription times for the handwritten letters. However, since the instructions were the same and the letters were of almost identical length, it seems appropriate to use Gould's transcription times to calculate total costs. This has been done in table 2.1. As shown, if the author's time is worth twice a typist's time, the estimated total labor cost of a letter is now about 26 percent less using the text editor, an encouraging but not impressive difference, especially if the additional labor costs for learning and maintaining the technology are taken into account.[7]

The results from these text editor experiments are particularly distressing for two reasons. First, they are controlled experiments, not subjective reports of managers who have committed themselves to the technology. Second, the investigators in these experiments were exceptionally competent and worked for companies with a vested interest in computerized office technology, so it's unlikely that the results were biased against computers.[8]

There are probably better uses for text editors than having professionals use them to produce their own business letters. The survey results from the insurance industry and from Quebec businesses suggest that typing pools equipped with word processors are 33 percent to 100 percent more productive than typists with typewriters, although we need to worry whether the respondents were counting drafts or finished output. Presumably the use of text editors for legal documents such as contracts, which often need a large number of revisions and final drafts that are letter perfect, would show even bigger benefits. I know of no good data for such specialized applications, although they may have formed part of the work behind the insurance company reports. Of course, there may also be applications of word processing that are less productive than the ones studied in the experiments. For example, many professional scholars, university professors, and industry and government scientists have been migrated away from the use of secretarial help toward the typing of their own manuscripts on computers. Here, since the salary of the person to whom the job has been moved is usually very much higher than that of the people whose job it used to be—perhaps three to five times as high—and the scholar or professional's skill at typing relatively lower, it seems likely that the true economics are positively terrible. One survey (Sassone 1992) confirms this expectation.

The other day I was printing out this section for local distribution. Table 2.1 got automatically separated onto two different pages, with a long footnote in between. At one point there were three expert users in my office for fifteen minutes trying (fruitlessly) to figure out how to fix it.

The area of phase two computer application for which the most extensive comparison data are available is computer-aided instruction (CAI). Instruction was the domain of some of the earliest attempts to use computers to help people better use their time and brains. Emerging in the mid-1960s, CAI even predated the popularization of text editing programs. Probably because the early versions of CAI were largely the work of academic educational researchers, for whom the testing of novel teaching methods against standard approaches is a professional obligation, the relative effectiveness of these systems was evaluated in literally hundreds of experiments.[9] Not all of the results were positive, but most were. Applications to the teaching of well-defined technical topics were most successful, and their use in industrial and military training was the most extensive. Researchers reported that such applications, when done well, reduced learning time by around 30 percent (Clark 1985; Eberts and Brock 1988; Kulik, Kulik, and Shwalb 1986; Kulik, Kulik, and Cohen 1980; Orlansky and String 1979).[10] Because industrial training is a big expense, and getting bigger with the information revolution, improvements of such magnitude, if repeated in following generations, could make a significant contribution.

Unfortunately, some caveats are needed. First, most of the gain from "computerizing" instruction apparently comes from careful analysis and pruning of the content, followed by detailed planning of the order in which core facts and skills should be introduced, and especially from individualized instruction. Individualized instruction allows students to progress at their own rates, rather than all moving at the rate that most members of a class can handle. Thus, if ten students each learn at their own speed, the average completion time is the average of ten individual learning times. If the same ten students learn in a standard classroom, the teacher will probably go at a speed that keeps the eight best students on board—likely to be about two-thirds as fast.[11]

Although good analysis and individualization is essential, or at least obviously desirable, for computer-based instruction, it is often employed without computers. Indeed, it has been institutionalized under the title of "instructional technology" (Clark 1985; Eberts and Brock 1988; Gagne 1974, 1987; Kulik, Kulik, and Schwalb 1986; Kulik, Kulik, and Cohen 1980; Orlansky and String 1979), which customarily uses only lectures, standard audiovisual aids, and printed materials. Even when administered with only traditional media, individualized instruction produces large improvements in learning efficiency (Block and Bums 1978; Bloom 1984). Thus, the role of the computer as such in CAI gains is not obvious.[12] Of course, if putting a lesson on a computer encourages better teaching, or if it makes the process cheaper, replacing all or part of a teacher's time without adding compensating expenses, there would still be good reason to adopt CAI.

After almost thirty years, however, CAI has made only small inroads into education and training. The main reason is probably not lack of effectiveness, although CAI is easier to apply effectively to some kinds of material than others. Desirability is probably an important factor; many people want personal interaction between students and teacher. Cost may be the biggest hurdle. The cost of writing good materials and programs for computer-based lessons can run two to many hundred times that for preparing traditional classroom lessons (Avner 1979). Add outlays for hardware, maintenance, operation, space, and administration, and CAI may not compete well with its paper-based instructional counterparts. Thus, the promise of 30 percent efficiency gains, which also, apparently, did not multiply with succeeding generations, did not succeed in greatly reducing human labor in instruction.[13]

Another phase two computer application that has been compared with noncomputerized benchmarks is the retrieval and presentation of textual information. These systems have seen widespread use in bibliographical searching—that is, finding not the actual article or book that the searcher wants to read but a reference to where to look for it in paper or microfilm. Usually the electronic reference is identified by title, author, publisher, date, a few subject category nouns, and, sometimes, a short abstract of its contents. There is a large literature of theory and tests concerning how well such systems succeed in returning all and only those references that a user wants (we will discuss some of this research later.) Unfortunately, few compare how well the highly evolved prior technologies of card catalogs and paper indexes serve the same purposes. Generally the proponents of electronic information retrieval have been pleased if people using the systems do as well with them as with so-called manual methods using printed sources (Cleverdon 1979). The two techniques usually produce rather different qualitative results; online methods search more potential titles, but manual methods detect relevance more intelligently. One doleful conclusion is that "an on-line search cannot substitute for a manual search and vice-versa: the two methods complement each other" (Murphy 1985, 178). If libraries do *both* kinds of searches for a request, there will be no cost savings.

A handful of rather limited studies have sought to compare search times and overall costs (East 1980; Murphy 1985; Roose 1985), but the cost accounting and evaluation methods in these studies leave much to be desired; for example, they do not always involve the same search topics, or there are no exhaustive search data available with which to compare, or the cost of equipment or training, on the one hand, or of printed source material, on the other, are omitted. Nevertheless, the results are encouraging. Librarian work times range from around two-thirds as long for online as for manual searches down to as low as one-seventh. That is, labor productivity for reference librarians in this aspect of their job improved by factors of 70 percent to 600 percent. Labor costs for provision of the online service were not broken out but are included implicitly in the service provider's charges. Overall the total cost of manual and online searches appears to be roughly the same. However, the most recent of these studies is now eight years old (a noteworthy and discouraging fact itself). It is likely that computer and communications costs have come down considerably, although other components, such as average royalty payments, probably have increased.

An interesting aspect of this case is the way IT has been used. For manual searches, each library maintains its own extensive and expensive set of reference books and indexes. Online computerized systems use a single centralized set of databases to serve a large number of libraries. If the systems allowed libraries to dispense with their own reference-work collections, the efficiency advantage could be significant. Individual libraries should reap large savings in book purchases and storage space, as well as library staff time. But experience so far has been duplication rather than replacement of resources. I have been unable to find any documentation of such "cooperative operations" savings, or of any remarkable decrease in relevant library staff employment.[14] Indeed, a whole new employment category of library personnel has evolved: information specialists, who know how to use the various computer systems and the myriad databases, each with its own method of operation, to which the systems are connected.

Often such systems allow the searcher—or the searcher's trained intermediary—to locate materials in many far-off places: libraries throughout the California university system, for example, or legal abstracts stored on a single immense, centrally located, commercial database. Such a facility can be a powerful tool for scholars, lawyers, physicians, business analysts, and other intensive knowledge workers. The commercial success of the remote database industry, represented by such firms as Dialog, Mead Data, and Dow-Jones, implies that these new computer-based services add significant real value to the economy. Yet their net productivity, as compared to other methods for doing the same work, is hard to assess. The same intensive culling, indexing, and abstracting by subject matter experts is required, and insofar as similar compendia have been provided in paper books (for example, *Index Medicus* or *Chemical Abstracts*), the principal advantages of the computerized versions are speed and economy of delivery and updating—that is, telephone line access instead of printing and physical shipping—plus the savings in shelf and office space. Each lawyer, chemist, or stock analyst needs only a terminal instead of a multivolume paper document. Obviously, if these economies result in more useful knowledge for more people, there is another significant gain. But how the dollar value of all these advantages stacks up against their added computer, communication, and specialist costs is not known.

A recent addition to the area of computer applications is hypertext. Hypertext uses computers to deliver and display the full content of written information rather than just references to it; the information is supplemented by links between segments of text, and between text and figures, data, references, and notes, all of which can be brought to the screen with a click of the mouse. I have found nine studies comparing hypertext with old-fashioned paper book technology as a means for people to find and absorb information. In only two of these studies was the computer-based method superior, and this particular system is quite different from almost all the many hypertext systems on the market. In most cases, hypertext actually *decreased* user efficiency relative to paper and ink. In one especially telling case, both objective performance and user satisfaction were significantly worse with hypertext, despite the fact that the text was explicitly designed for a popular hypertext program. We will go into the matter of hypertext, and the design of SuperBook, an exceptional case, in detail later.

There are several other phase two applications for which strong claims have been accompanied by reports of study "data." Although some of these results are tantalizing, none is convincing. For example, it has been claimed that "groupware"—computer systems to help run meetings—reduces meeting time needed to plan and manage projects by half. Such an effect could have significant productivity implications because a large portion of the time of highly paid managers and professionals is spent in meetings. But the "controls" in these studies are merely estimates by either the authors or the participants of how much time would have been needed without the computer system. Such estimates are much too vulnerable to bias to be taken seriously. By contrast, a survey of all published controlled behavioral experiments (McLeod 1992) found that although they improve decision quality slightly, such systems usually significantly *increased* the time groups spent making decisions.

There are also phase two systems for which no objective comparison data are available but which are so widely used with such apparent success and strong testimonial acclaim that it is hard to doubt their effectiveness. The prime example is computer-aided design (CAD). These systems help designers of integrated circuit chips, airplane wings, or buildings get all the parts in all the right places. Circuit design is probably the most important use of this technology. Very large-scale integrated (VLSI) circuits can have millions of components that have to be connected by millions of wirelike conducting paths. The components have to be arranged so that they can dissipate heat without destroying each other. All the paths have to go to and from the right places. To be efficient, and sometimes even to work at all, all or most of the paths have to be as short as possible. And the whole design has to be verifiable by inspection of some sort. All this poses a daunting, probably impossible, task for an unaided human. But the task is not yet fully automated; a highly skilled human who can see and think the whole layout is still needed. The computer system helps by making it easy to copy, multiply, and move components, by suggesting component placements and path arrangements based on computationally intensive algorithms, and by doing much of the checking. It is said that the use of such systems reduces the time to design complex circuits by large factors.

The productivity effects of these tools has been felt most directly in the computer industry itself, which has shown extraordinary growth. The effect in other industries has apparently been less, probably because design is not so difficult without computers and because design work is not as large a contributor to costs or effectiveness. The fraction of labor time spent in design is limited. CAD systems take months to learn, sometimes years to master. No one can make a computer chip without CAD, but for other products, their value is often marginal. An architect may prefer to sharpen her pencil. Much of the popularity of CAD is not attributable to actual salary savings for designers but to belief that CAD will reduce the cycle time to get a high-tech product to market. This may be a critical consideration for a competitive company, but the productivity effect is negligible.[15] Since the real payoff for CAD is in the new or improved products they make possible, we would have to turn to indirect economic and business success data for evidence, discouraging evidence that we have already reviewed.

Let us take one last well-known, and disappointing, example. ATMs, automatic teller machines for banking, are almost always thrown up as a counterexample when I give talks on this material. Members of my usual audiences—primarily computer scientists, human factors professionals, R&D managers, high-tech business types—are near-unanimous fans of ATMs. Yet banks have experienced neither overall productivity gains nor reduction in clerical staffs as a result of their introduction. Nor have they generally offered higher interest, lower transaction fees, or other incentives for the use of ATMs. Indeed, many banks make additional charges for ATM use. What's up? A fascinating study of a large western bank chain tells the tale (Haynes 1990). Of all customers with debit cards for the company's wide network of machines, only a third ever used an ATM, and among those who did, the average use was once per month. Overall, only about 5 percent of all customer banking transactions were completed with ATMs. Most of the planned economies from the machines were based on expectations of much higher use, of many more labor-saving transactions to offset the costs of hardware, servicing, and real estate for their location. Haynes quotes estimates that ATMs as used reduced a bank's cost for a transaction by roughly half as compared to a teller transaction. Over the whole country this would represent a multimillion dollar cost saving. However, the small proportion of business costs involved, especially in relation to the rather large capital and management expenditures, does not add up to a significant influence on bank productivity. We will have more to say later about why those who love ATMs love them, why the transaction rate is so low, and why banks have chosen to deploy them so aggressively.

What about all the myriad other uses of computers? Sadly, for both the analysis and for progress, I have about exhausted the available direct data on productivity and efficiency effects (table 2.2). For some applications there is indirect evidence from which we can draw tentative infer-

Table 2.2
Summary of work efficiency effects of phase two computer applications

Application	Evidence	Improvement with computer
Word processing	Experimental studies	Worse to +25%
	Survey reports*	+33 to 100%
Computer-aided instruction	Meta-analyses*	**+30%**
Information retrieval	Studies of users	**None** to +500%
	Estimates for librarians*	+70 to 600%
Electronic document delivery	Studies of users	**Worse** to +100%
Meeting support systems	Experimental studies	**None**
ATMs	Usage analysis, $ benefit	Almost nil

Note: Estimates with asterisks require cautious interpretation for reasons given in the text. Estimates in boldface type are based on several to many studies, those in regular type on only one.

ences. For example, consider all the applications of PCs for use in the home. The productivity effects of home computers, and services and software to go with them, are quite difficult to assess. The government does not try to measure the productivity of domestic households, and none of the economists interested in the productivity crises has either. We have already reviewed evidence on several kinds of programs that such machines often support: text editors, information retrieval systems, and instructional software. But home computers are also usable for many other purposes for which we lack similar data. It appears from various surveys that a high proportion are not used at all, and that among those that are, the primary use is for playing games. The extent and time-efficiency effects of home computers for tasks like budgeting, bill paying, and tax return preparation have not been properly evaluated, to my knowledge. Because home computers are used only occasionally by a small number of people, they must have a minor effect on total private time utilization. (Compare their time-saving effects, price, and learning time with washing machines, vacuum cleaners, and clothes dryers.)

There have been several attempts to apply home computers to facilitate other services, for example, information services supplied over telephone lines. In the United States, Canada, Japan, and most of Europe these have gone through several cycles of failure. Although fairly large numbers of network memberships have been sold, such services have not been financially successful. The nearest to successful is the French Minitel system, which, begun in 1980, will break even on the enormous investment of the French government in 1993, 1995, or 1998, depending on which of three audits you believe (Pinsky 1991). The secret of French success, in addition to the large social investment (France Telecom initially distributed half a million free terminals and will have spent $7 billion by break-even), appears to have been a comparatively simple, easy-to-use interface and the market evolution of many thousands of different services that can be accessed through the network. Recent evidence suggests that the French data network has reduced transaction costs for the businesses that use it but that their overall return on investment has yet to show an effect.[16]

Because there are so many uncontrolled factors, inferences about productivity from facts about market success and usage are not very compelling and will not be pursued further in this chapter. Similarly, there have been many reports of good and bad user experiences with various computer applications, but it is impossible to add them together for an estimate of net productivity. These kinds of evidence will be important in later chapters, where we consider how to make things better.

Silver Linings

In this long stream of dismal productivity trends, comparisons, and individual efficiency studies, a few fish swim against the tide. These successes are particularly important to examine because they may offer clues about what is wrong (in the data or in the effort) in the other cases, and they may point a better way. We will consider the source of exceptions at the individual efficiency level—where we have just seen a number of modest successes—in later chapters. Here we concentrate on the only major service industry that has maintained strong productivity growth into the information age: the telephone business.

Despite all the negative results to date, the promise of computers for productivity is tremendous. But given the length of time they have been around, one would think that if they really have great potential, some-

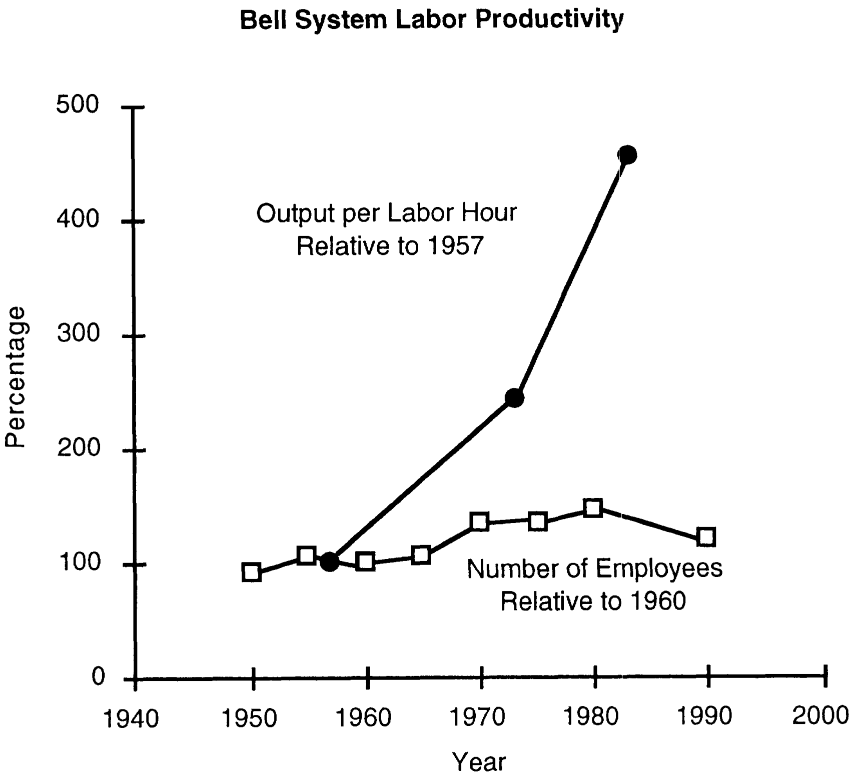

Figure 2.1.
Bell Company (plus AT&T after 1984) employment and labor productivity. Data from FCC and Baily 1986.

where someone should have pulled off a major, lasting productivity success with them. There is such an example: the telephone industry.[17] We saw evidence of this before in the relatively healthy overall productivity gains of communications as a whole in the critical phase two computerization years following 1973. We also saw that the telecommunications sector is the only major service industry with top-ranking productivity results. According to the Bureau of Labor Statistics, dollar output per labor hour grew by 6.1 percent per year from 1957 to 1973 and by 6.7 percent from 1973 to 1983 (figure 2.1). There was no post-1973 slowdown but, rather, a modest acceleration in labor productivity growth.

There is even more direct and impressive evidence. Between 1950 and 1990 the number of access lines (essentially lines with separate telephone numbers) increased fourfold, while the number of employees needed to install, operate, and maintain them grew by less than 40 percent. Measuring productivity as customer lines per full-time employee, the companies averaged 2.7 percent gains each year from 1950 to 1970 and 3.6 percent from 1970 to 1990. Employees per line is a rather interesting pure labor productivity measure but somewhat conservative.

Later lines were much more useful; they connected to many more destinations and did so faster and with better sound quality.[18] At the same time real (after inflation) wages of telephone employees increased, work hours per employee decreased, the cost of telephone service went down, revenues grew steadily, and shareholders earned respectable returns and appreciation. Some of the data are shown in figures 2.2 and 2.3.

Thus in telephony there was obvious and substantial improvement over the time period in which phase two computing was integrated into the business. (Note that we can't measure telephone business productivity meaningfully in terms of return on investment because the public utility commissions set a fixed rate.) The telephone business was able to continue America's traditional double-time march of progress: the industry cut costs, improved its product, lowered prices, and generated ever increasing demand whose supply created ever more jobs at ever higher wages.[19]

There were other credible contributions to these gains besides the effects of computerization: lower unit costs of transmission facilities, lower per line cost of switching equipment, and increased labor efficiency due to more reliable and easy-to-handle hardware. Probably the rapid expansion of the total telephone service market as more people were connected more easily to more people, and consequently made more calls, played an important role. However, a large part of the productivity improvement also has to be attributed to what telephone people call mechanization. The biggest and earliest, and continuing, mechanization was of the job of connecting a calling line with a receiving line. The switches that do that work in response to rotary dial pulses or keypad tones are, and always have been, computers. The first ones, introduced in the 1920s (initially by independents, then by Bell System companies), were fairly crude electromechanical devices (although some operating on the same principles are still in service). Nevertheless they perform the same function as their modern digital, VLSI-based replacements. They "compute" and establish a path between two points. They were, of course, special-purpose computers but computers nonetheless.

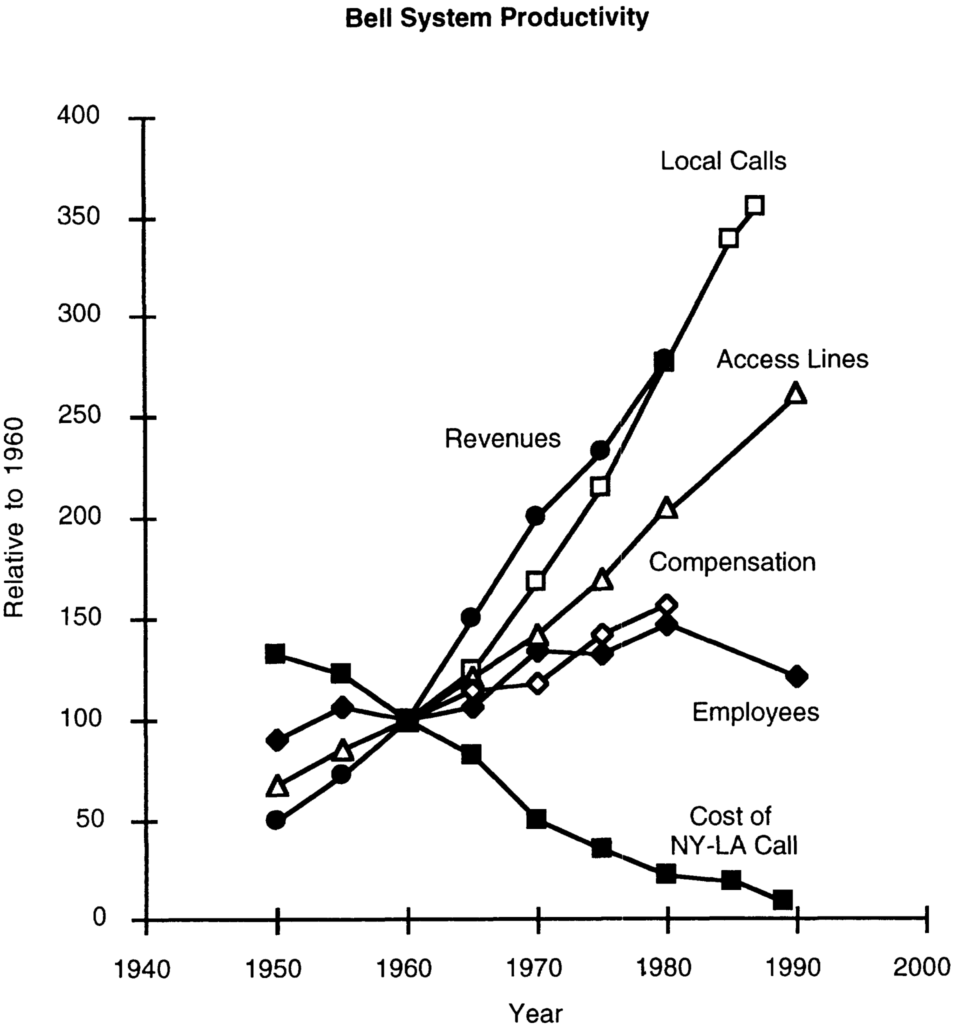

Figure 2.2.
Productivity in the Bell System before divestiture and the regional telephone companies plus AT&T thereafter. Data from FCC and BLS in constant dollars relative to 1960.

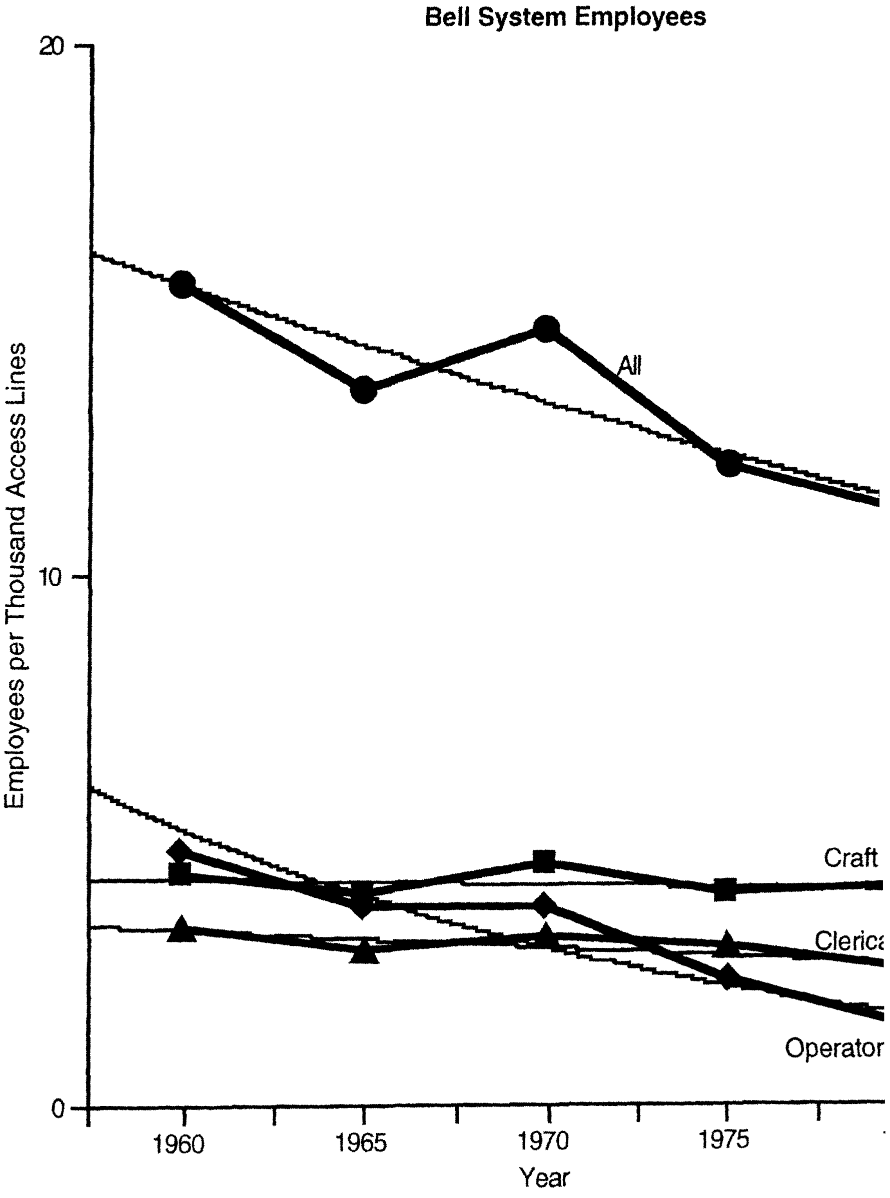

Figure 2.3.
Change in ratio of employees to customer lines in Bell System companie
1980. Data from FCC and Bellcore.

Switching machines are the purest of pure automation computers. Each succeeding generation of switching computer has enabled telephone companies to connect more lines, faster, and at lower capital and operation labor cost per line. (If we were still doing it entirely by hand we would need over 3 million telephone operators to handle today's calls, if we could.) However, their effect in reducing the total number of employees needed was slowing rapidly long before the last manual switchboard was melted down. As the number of lines and the volume of traffic grew, there was a continuing potential increase in labor requirements because each line entailed a telephone number, an order to process, a monthly bill, and an opportunity for repair, and each call presented the possibility of a request for assistance in finding a number or collecting a toll. After about 1960, we might have expected at least a linear increase in labor costs as the network expanded. Each new customer or call added increased demand for employee services. Actually, the number of employees needed to provide the services would grow even faster than the number of new customers because the complexity of some of the individual jobs depends on how many total customers and calls there are. When a caller asks for a telephone number, the directory assistance operator takes longer to find it if there are a hundred thousand listings than if there are only twenty thousand. It also takes a billing clerk longer to find a record, a cable splicer longer to find a particular wire pair, and so forth the more of them there are. Some such tasks may tend to require labor in proportion to the square of the number of lines rather than just in proportion to the number.

To keep costs down as the customer and traffic base grew, starting in the mid-1960s, the telephone companies began to develop phase two computer applications to ease and speed the work of employees; deployment on a large scale got underway in the 1970s. These operations support systems have been primarily databases used to keep track of all the wires, cables, repeaters, plug-in circuits, trucks, addresses, telephone numbers, bills, and payments and to help in the planning, designing, and administering of the network. Most of these functions still require human intelligence; only certain record-keeping, calculation, and printing or display aspects of the jobs are performed by the computers, and usually in close collaboration with workers. Although the development and deployment paths of these systems have often been rough and their net effects on labor productivity not always immediately favorable, in the aggregate long term they have made a major contribution to keeping costs down. The number of employees per line continued to fall even after full automation of switching and even as the size and complexity of the network expanded rapidly; from 1965 to 1990, the number of lines more than doubled—and the traffic and complexity grew much more—but employment rose by only 14 percent.

Prior to the break-up of AT&T, the Federal Communications Commission collected detailed information on the number of telephone employees and their jobs (figure 2.3). In terms of number of workers for each thousand customer lines, the biggest gains from 1960 to 1980 were in operators, representing a continuation of automation of connection, especially long distance, plus augmentation of directory assistance and specially billed calls. There were also significant gains in efficiency for clerical jobs. The craft jobs of installation, maintenance, and repair continued to demand almost the same effort on a per line basis despite improvements in equipment but also despite the growing size and complexity of the system.

How did the telephone companies succeed where, it seems, most others failed? Part of the reason probably lies in the kinds of tasks supported by computing. The telephone companies had huge numbers of employees doing a number of well-subdivided, highly routinized tasks. Many of these jobs were well understood and had already been reduced to simple "if this, do that" operations that were relatively easy to automate or aid. For example, many jobs involved entering, correcting, or retrieving information from paper records, or creating tabulations, reports, or work orders on the basis of information received from a customer or employee or displayed by a switching machine. Often the computer could be an effective aid merely by putting all needed information in one convenient place so that employees did not have to search physically separated file drawers and make copies and transcriptions by hand. But the most important factor in the success of telephone operation mechanizations was the way in which they were designed and evolved. In particular, AT&T and the Bell System companies had a long tradition—to some extent mandated by their regulated status and in some ways directly assisted by statisticians at the Federal Communications Commissions and other public utility commissions—of careful, objective, quantitative monitoring of the efficiency of office and business operations. Moreover, meaningful quantitative measurements of labor efficiency were relatively easy to obtain: time to handle collect and directory assistance calls, number of equipment troubles and the time to fix them, and so forth. In addition, because of the crucial importance of reliability in the telephone system, telephone companies always test new computer systems extensively with real operators handling real work loads before they replace old procedures or systems. (It is surprising how rare such testing is in other business applications.) All this provided the kind of feedback that I will argue is the most critical need for progress in phase two computing. We will return to more detailed description of this aspect of the case later.[20]

"If you look back a few years, our personal lines operation, for example, had 80 to 100 field offices. And all of the policies . . . were in tub files and desks in those offices. So if you wanted to see all the policies you would have to go to 90 offices and read through 2 million policies in tubs and desks. Then all of a sudden a person could just issue a few commands to the computer and we'd zip through it for him. That's a huge change." Larry Bacon, senior vice president of information systems, Travelers Companies (National Research Council 1994a)

"During the three year period from 1989 through 1991, we reduced our annual output of computer-generated print from 1.3 billion to about 400 million pages . . . downloading print files electronically to client desktops and providing the ability to view information directly from computer screens." Howard Sorgen, senior vice president and chief technology officer, Merrill Lynch & Company (National Research Council 1994a)

No other segment of the service economy nearly as important as telephony has a good record with computers, but there are individual firms and activities within firms that have posted notable successes. A well-known company success story is the case of Federal Express. Through a combination of bar codes, hand-held devices, and a widespread computer network that enables it to track and manage each package, it was able to implement guaranteed next-day delivery service, a new and profitable form of business.

Within companies, the process of handling complex orders has given rise to most of the dramatic reports of efficiency gains from computers. Most of these gains are in turnaround time, but often labor is also greatly reduced by making records more accessible and reducing hand-offs, checking, and error correction. Retailers such as Wal-Mart and Kmart have plugged their stores into computer networks that collect statistics on what sizes of clothing are needed and what CDs are selling, thus reducing waste, stocking labor, and accounting.

We will see more of these tantalizing examples later when we analyze why there have been so few and how we can get more.

3

The Productivity Paradox

The subject of this book is computers—their problems and how to make them better—but the analysis has taken us deep into the broader puzzle of stunted productivity growth. Before going on about computers in particular, let us see if what we've learned suggests an explanation of the overall productivity problem.

The Productivity Slowdown: A Hypothesis

The United States, with the rest of the industrial nations lined up a few steps behind, has learned to automate its factories and agriculture so that little human labor is needed. About a third of the population produces almost everything people used to want—food, clothes, cars, trinkets, houses. It turns out, though, that the rest of us aren't ready to retire. So as a society, we've had to invent new work to do for each other and pay each other for doing—a new set of exchanges. We cook meals for each other, take in each other's dry cleaning, entertain each other, govern each other, write books and produce movies for each other, take each other on long trips, care for each other when sick, and organize gambling games—finance and insurance, lawsuits, lately banking. These activities are not entirely new, but they are being greatly expanded as a proportion of production and consumption. They are called, loosely, "services." Most of the work is "information work."

Trouble is, this kind of work, doesn't have as good a ratio of labor hours to dollars paid for the output as do the highly automated means of production for food and other tangibles. Service labor productivity is