

ACM Copyright Notice. Communications of the ACM, August 1998, Vol. 41(8): pp. 49-55.
 Copyright © 1998 by Association for Computing Machinery, Inc. (ACM). Permission to
 make digital or hard copies of part or all of this work for personal or classroom use is
 granted without fee provided that copies are not made or distributed for profit or commercial
 advantage and that copies bear this notice and full citation on the first page.

Beyond the Productivity Paradox

Computers are the catalyst for bigger changes.

*Erik Brynjolfsson
 and Lorin M. Hitt*

Why Should We Care About Productivity?

An important question that has been debated for almost a decade is whether computers contribute to productivity growth. Productivity isn't everything. However, as noted by the economist Paul Krugman, in the long run it is almost everything. Productivity growth determines our living standards and the wealth of nations. This is because the amount a nation can consume is ultimately closely tied to what it produces. By the same token, the success of a business generally depends on its ability to deliver more real value for consumers without using more labor, capital, or other inputs.

Productivity is a simple concept. It is the amount of output produced per unit of input. While it is easy to define, it is notoriously difficult to measure, especially in the modern economy. In particular, there are two aspects of productivity that have increasingly defied precise measurement: output and input. Properly measured, output should include not just the number of widgets coming out of a factory, or the lines of code produced by a programming team, but rather the value created for consumers. Fifty years ago, tons of steel or bushels of corn were a reasonable proxy for the value of output. In today's economy, value depends increasingly on product quality, timeliness, customization, convenience, variety, and other intangibles.



Similarly, a proper measure of inputs includes not only labor hours, but also the quantity and quality of capital equipment used, materials and other resources consumed, worker training and education, and even the amount of “organizational capital” required, such as supplier relationships cultivated and investments in new business processes. The irony is that while we have more raw data today on all sorts of inputs and outputs than ever before, productivity in the information economy has proven harder to measure than it ever was in the industrial economy.

Where does productivity growth come from? By definition it doesn’t come from working harder—that may increase output, but it also increases labor input. Similarly, using more capital or other resources does not necessarily increase productivity. Productivity growth comes from working smarter. This usually means adopting new production technologies and techniques.

The greatest increase in productivity has historically been associated with a particular class of technologies: “general purpose technologies.” The steam engine was an important general purpose technology. It could be used in a variety of new applications, from driving spinning looms in a newly mechanized factory to powering locomotives in a new transportation system. Electricity was another key technology that set off a chain of innovation in the 1890s.

What general purpose technology might hold a similar promise in the 1990s? The obvious answer is information technology. Driven by Moore’s law—the doubling of the number of transistors per chip every 18–24 months—computer technology has advanced at an exponential rate for several decades (see Figure 1). Ultimately, however, these trends in basic computer power only provide greater inputs into production. The question remains: Are computers increasing output? Are computers pulling their weight?

On the one hand, amazing success stories abound: the billions of dollars already being transacted by firms like Dell and Cisco via the Internet are only the latest example. On the other hand, there is no shortage of stories about cost overruns, abandoned systems investments and other IT failures. Some authors have even described the idea that computers have substantial business benefits as “the big lie of the information age” [12]. Anecdotes can be found to bolster either side of the debate.

A better way to determine if computers are living up to their promise is by studying broader data sets that contain hundreds or even thousands of observations. The idea is that unusually “lucky” or “unlucky” experiences with computers will tend to average out and we will be left with a clearer picture of the underlying relationship. We have reviewed several such studies, many of which were originally presented at the Workshop on Information Systems and Economics (WISE), across a wide range of technologies, industries and applications. We find that a consensus is beginning to emerge: Computers *are* pulling their weight [1–3, 8, 10].

In addition, a second, even more important finding is clearly evident in the data. While the average returns to IT investment are solidly positive, there is huge vari-




Figure 1. Moore’s Law.

The number of transistors that can be placed on a semiconductor die doubles approximately every 18 months.

ation across organizations; some have spent vast sums on IT with little benefit, while others have spent similar amounts with tremendous success. Today, the critical question facing IT managers is not “Does IT pay off?” but “How can we best use computers?”

Fortunately, the same methodologies used in investigating the first question can be directed toward the second question, and a number of provocative results are emerging. Most importantly, the greatest benefits of computers appear to be realized when computer investment is coupled with other complementary investments; new strategies,

new business processes and new organizations all appear to be important in realizing the maximum benefit of IT. This change is rarely easy since many organizations will require a painful and time consuming period of reengineering, restructuring and organizational redesign in order to best utilize their IT investments. However, once these investments in change are made, these companies will be positioned to reap the benefits of continued technological progress in the computer industry, while others may be left further and further behind.

The Productivity Paradox

Attention was first drawn to the “productivity paradox” by a simple but provocative study, “America’s Technology Dilemma: A Profile of the Information Economy” by Morgan Stanley’s chief economist Steven Roach published in Morgan Stanley’s April 22, 1987 economics newsletter series. He attempted to explain why the measured productivity growth rate in the U.S. economy has slowed substantially since 1973. Roach observed that the amount of computing power per white-collar worker in the service industry was growing dramatically over the 1970s and 1980s, yet the measured productivity of this sector was flat. His conclusion was that the tremendous increase in computerization has had little effect on economic performance, particularly for those sectors of the economy with large numbers of “information workers.”

Other studies also showed little evidence of a link between computer investment and productivity using data on computer investment in manufacturing industries or in a sample of business units of large firms. A few studies found positive effects on intermediate factors such as cost efficiency or market share, but it was still difficult to tie these benefits to the bottom line. Furthermore, despite the tremendous advances in computer power, the aggregate statistics suggest that productivity has grown more slowly since 1973 than it did between 1950 and 1973. By the late 1980s, the conventional wisdom was that computers were not contributing significantly to productivity. As succinctly stated by Robert Solow in the *New York Times Book Review* (July 12, 1987) “we see the computer age everywhere except in the productivity statistics.”

However, while these results generally found little evidence of a relationship between IT and productivity, there was also little evidence that computers were unproductive. In particular, many people pointed to the inadequacies of productivity measurement. One problem is that until recently overall computer investment was relatively small compared to overall capital investment and labor expenditure.

Economist Zvi Griliches in his presidential address to the 1996 annual meeting of the American Economic Association likened the search for IT value to looking for a needle in a haystack. However, even as the magnitude of IT investment grows larger, he notes that there are still systematic biases in conventional productivity measurement that prevent an accurate assessment.

Most productivity metrics are oriented around counting things: number of employees, pounds of nails, or number of checks processed. As long as computers allow firms to produce more of the same product at lower costs, these metrics work reasonably well. But there is strong evidence that managers are not simply making IT investments to cut costs. When managers are asked, “Why do you invest in IT?” surveys suggest that customer service and quality consistently rank above cost savings as the prime motivation for making investments [3].

The quirks of productivity measurement are easily seen in banking. ATMs reduce the number of checks banks process so, by some measures, banking output and productivity decrease. The increases in convenience ATMs have created go uncounted in conventional productivity metrics, while their costs are counted. At an aggregate level, banking labor productivity is measured, like all sectors, as the ratio of an output metric to number of employees. But since the aggregate level of the true “output” of banks is difficult to measure, most conventional analyses have shown that labor productivity has essentially been flat. Not surprisingly, when you can easily count the costs of computer investment but have a difficulty assessing the benefits, particularly those that take time to be realized, IT can look like a bad investment.

An “Information Payoff?” In the early 1990s data became available that allowed a reexamination of some of the previous results on IT productivity. This data, for the first time, enabled researchers to look at the IT investment behavior and productivity of large numbers of firms rather than focusing on higher level aggregates such as manufacturing industries or the whole economy. This microlevel approach had a number of advantages. While there is only one U.S. economy and only a few dozen manufacturing industries, this data allowed analyses to be conducted on hundreds of firms over several years. The increase in sample size enabled much more precise estimates of IT contributions, improving the chance of identifying the needle in the haystack.

Firm-level data also enables the measurement of at least some of the intangible value that was being created by computers even if this value could not be directly observed. If consumers are willing to pay

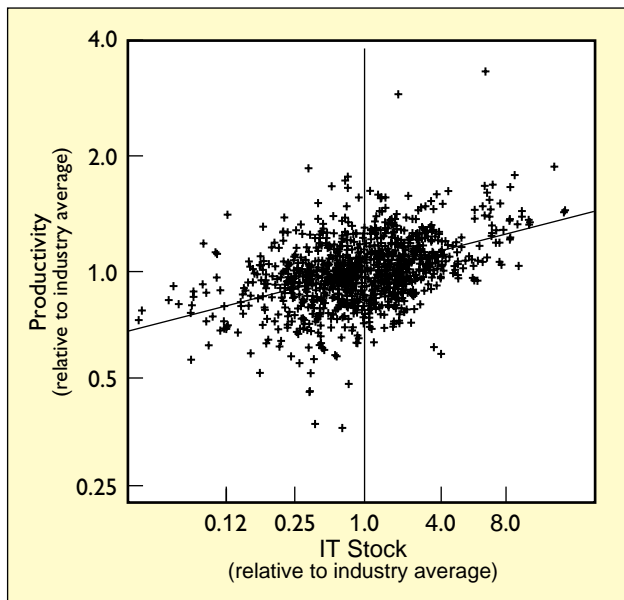


Figure 2. Variation in productivity and IT investment across firms.

The vertical axis (labeled “Productivity”) is multifactor productivity, defined as output divided by a weighted sum of inputs (in constant 1990 dollars). The horizontal axis (labeled “IT Stock”) represents the total IT inputs in a firm. Both productivity and IT input are centered at the industry average. The points are approximately 1300 individual observations representing an individual firm in a particular year.

more for increases in quality or convenience, then a firm’s revenue will reflect some of this increase in intangible value. However, these differences will not appear at the industry level; high-quality firms force low-quality firms to lower their prices to remain competitive. Therefore, overall industry revenues will not necessarily increase as firms computerize. While some of the value from IT investments made by firms and passed on to consumers through competition will not be observed—at least some of this intangible value can be captured in productivity measurements.

Initial firm-level studies of IT and productivity found that a dollar of IT capital is associated with a substantial increase in revenue each year [1, 2, 11]. Other analyses have replicated these basic findings using different sets of econometric assumptions, different characterizations of IT (mainframes, PCs, IS staff or some combination), and different subsets of the economy (manufacturing vs. services) [3, 8]. Across all these studies there is a consistent finding that IT has a positive and significant impact on firm output, contradicting claims of a “productivity paradox.”

In fact, in these studies the returns to IT appear to be quite high. This raises the possibility that com-

puters are not only pulling their weight but contributing substantially more. However, at least part of this high rate of return is required to compensate for rapidly falling prices of computer equipment. In addition, these returns may represent more than just the returns to the technology. Technology is only one component of an IT investment; there are usually large expenditures on training, process redesign and other organizational changes accompanying a systems investment. This doesn’t change the conclusion that computers contribute to increased output, although it does make exact rate of return calculations more difficult.

Beyond the Averages

While computers on average appear to be productive, this fact alone is not enough for an IS manager to make good investments. In fact, the difficulty of establishing the overall value of IT may indicate that the value that IT brings to a firm varies enormously from company to company. When we plot the relationship between IT and productivity in Figure 2, two features stand out. First, when a line is fitted through these points, it slopes upward, which suggests that firms with more computers are compensated by increased output. However, more strikingly, there is an enormous amount of variation around this line; some firms have high IT investments and are highly productive, others have similar investments but poor performance. What explains the difference?

One way to start thinking about the sources of variation is to divide the benefits of IT into two parts: those that are unique to a particular firm, and those that appear due to variation in spending across firms. These two dimensions can be distinguished by a statistical technique known as a firm effects model. Applying this technique, we found that the measured benefits of IT were reduced by almost half when firm effects were included [1]. One interpretation of this result is that about half of IT value is due to unique characteristics of firms, while the remaining part is shared generally by all firms. What goes on inside the “black box” of the firm has a substantial influence on the productivity of IT investments.

To obtain a better characterization of the organizational factors that affect IT value, we can look at the relationship of IT investment to productivity over different time periods. If organizational changes can be made instantaneously with IT investments, then it should not matter whether we look at one-year changes in the firm or five-year changes. However, if there is some lag or adjustment time required to match organizational factors and IT investments,

we would expect to see more benefits over longer time periods. The statistical results were striking (Figure 3). While short-term benefits were about what would be expected if they had “normal” returns, long-term benefits were substantially larger: from 2 to 8 times as much as short-term benefits [3].

Our interpretation is that the organizational factors that unlock the value of IT are costly and time consuming. This could explain why the effects rise (these changes take substantial time and are put in place incrementally) and why IT appears unusually productive in the longer term: the long-term benefits are not just the returns from IT but returns from a system of technology and organizational changes; for every dollar of IT there are several dollars of organizational investments that, when combined, generate the large rise in measured firm productivity and value.

The Arrival of the “New Organization.” Tom Malone [11], Peter Drucker, and others recognized that general changes in the economy as well as the increased diffusion of IT into the workplace would facilitate and necessitate a dramatic restructuring of organizations. For example, Drucker’s article, “The Coming of the New Organization” [9], predicted that technology-rich firms will increasingly shift toward flatter, less hierarchical organizations in which highly skilled workers take on increasing levels of decision-making responsibility. Similar ideas underlie other management trends such as business process redesign, the emergence of “high-performance work systems” and the shift from “mass production”-style manufacturing to flexible “modern manufacturing.” In essence, these all represent organizational changes that exploit low-cost communications and information processing capabilities created by IT.

Do these types of practices actually make a difference? The initial answer appears to be yes. Recently completed studies suggest that organizations that utilized decentralized decision making and have employees with greater levels of skill and education appear to invest more in information technology [10]. Irrespective of how IT is measured, there is a consistent positive relationship between the use of these technologies and a set of work practices that include the use of self-directed work teams, greater levels of individual decision authority, particularly over method and pace of work, increased investments in training and screening for education, and incentive systems that reward and encourage high team performance. Part of this relationship is due to the fact that organizations that employ large numbers of educated workers, particularly professionals, or




Figure 3. Productivity of IT investments over time.

The vertical axis represents estimates of the productivity growth contribution of IT capital. The numbers are estimated output elasticities of IT capital, which represent the percentage change in output for a small percentage change in the quantity of IT. The value would be approximately 0.01 if IT has a “normal” rate of return. These estimates were computed by linear regression and the different lines represent different statistical techniques. “OLS” refers to ordinary least squares. “SRF” (semi-reduced form) is similar to OLS except that labor expense was not included in the list of inputs to reduce biases on the IT estimates from reverse causality between output and labor expense. “IV” represents instrumental variables regression which is an alternative way of addressing reverse causality. Further discussion of this analysis appears in [3].

employ technology- and skill-intensive production processes are likely to use more IT and adopt decentralized structures. However, the relationship between IT and the new organization of work goes beyond that which would be predicted by the composition of the work force and is present both within and between industries.

A cynical explanation of these results is that firms adopting the new work practices are wasteful users of IT; they spend too much or are too quick to adopt various management fads including IT investment and the new work practices. In fact, the opposite appears to be true. In addition to spending more on IT, these firms also appear to receive slightly higher returns on their IT investments. When we look at combinations of IT and work practices in the 2x2 matrix shown in Figure 4 we see that firms that couple IT investments with decentralized work practices are about 5% more productive than firms that do neither. However, firms can actually be *worse off* if they invest in computers without the new work systems.

In addition to getting more total benefits from IT, these organizations also appear to be adopting IT at a faster rate. In 1994 this amounted to only about a 50% difference in overall IT investment intensity,

however, this gap is growing by 10% per year. Over the next decade, these decentralized and empowered organizations may begin to pull away from their industrial age counterparts in performance as they are better able to exploit increasingly inexpensive information technology. These results suggest that it is becoming increasingly important to organize in ways that leverage the value of IT. While these types of results may not hold across all possible settings, the general trend is clear. So why do so many organizations still retain the old structure?

Decentralization \ IT	Low	High
	High	0.161 (.0191) N=47
Low	0 (n/a) N=69	-.0366 (.0197) N=47

Figure 4. Productivity effects of IT and decentralization.

Each quadrant contains the average productivity, the standard error of the productivity estimate (in parenthesis) and the number of firms from our sample in each group. Productivity is defined as multifactor productivity: output divided by input costs (in constant 1990 dollars). The productivity numbers are relative to the “low-low” quadrant (which is set to zero as a reference point). More firms line up on the diagonal (high-high, low-low) than on the off-diagonals indicating that IT and decentralization are correlated. Reproduced from [4].

A plausible reason is that these types of organizational changes are time consuming, risky, and costly. Redesigning management infrastructure, replacing staff, changing fundamental firm practices such as incentive pay and promotion systems and undertaking a redesign of core business processes are not easy. In many cases this may involve abandoning business practices that may have been successful for decades in favor of work systems with which the organization has little experience, or adopting an abrupt, radical, and discontinuous change in organizational structure. The large number of documented difficulties and failures of change of this magnitude suggests that making these types of changes is indeed costly.

The experience of one firm we visited is instructive. It spent millions of dollars to implement a new computerized manufacturing process. Top manage-

ment was wise enough to understand that fundamental changes in work practices would also be required. For instance, to exploit the new, more flexible equipment, they proposed a sharp reduction in work-in-process inventories and more frequent product change-overs in production lines. Despite the best of intentions, there were initially no significant gains in either productivity or flexibility [5].

The reason was that workers maintained the old ways of doing things, not in a conscious effort to sabotage the new manufacturing system, but simply because they had too many ingrained habits. For example, one worker explained, “The key to productivity is to avoid change-overs and keep the machines running at all times.” This was a very useful rule of thumb with the old, inflexible equipment, but it nullified the benefits of the new machines. Effectively, the huge investment in flexibility machinery was being used mainly to make the new machines work just like the old machines! Interestingly, management’s attempts to empower line workers with more decision-making responsibility were not much more successful. Several workers confessed that they had no interest in having more decision-making responsibilities—when chatting with their colleagues at work, they preferred to discuss sports rather than statistical process control.

Eventually, the firm was successful in managing this change, and they ultimately surpassed their production goals. The unmistakable lesson was that purchasing computerized equipment was the smallest part of the overall cost of creating a new manufacturing system. The biggest costs were in changing the organization.

One way to think about these changes is to treat the organizational costs as an investment in a new asset. Firms make investments over time in developing a new process, rebuilding their staff or designing a new organizational structure, and the benefits from these investments are realized over a long time period. Our earlier results suggest that these types of organizational assets need to be matched to information technology assets in order to be maximally valuable.

To get an idea of the potential magnitude of these assets, one can look at how the stock market values different types of assets owned by the firm (see Figure 5). For large corporations, a dollar of most types of capital is valued by the stock market at about a dollar. However, for these same corporations a dollar of computer hardware appears to be associated with about \$10 of market value [6]. While it could be true that IT is extraordinarily productive, it seems more reasonable to argue that this extra \$9 represents the value of hidden complementary organiza-

tional assets. For instance, with each dollar a firm spends on enterprise resource planning software like SAP's R/3 system, it typically spends \$3–4 on consultants who implement the new system. Even bigger costs are incurred in employee retraining and time spent redesigning business processes. However,




Figure 5. Relative market value of computer capital.

Estimates of the market value of computers relative to the market value of other assets from [5]. This analysis suggests that a one-dollar change in IT capital is associated with a change of about \$10 in market value for the average firm in our sample. For this to be an equilibrium, there must be about \$9 of unmeasured intangible associated with each dollar of measured IT capital.

in the end, the firm has a new system with lasting value—they own a new asset. These types of assets don't show up on a firm's balance sheet but accompany and complement IT investments.

What We Now Know About Computers and Productivity

Research on computers and productivity is entering a new phase. While the first wave of studies sought to document the relationship between investments in computers and increases in productivity, new research is focusing on how to make more computerization effective. Computerization does not automatically increase productivity, but it is an essential component of a broader system of organizational changes which does increase productivity. As the impact of computers becomes greater and more pervasive, it is increasingly important to consider these organizational changes as an integral part of the computerization process.

This is not the first time that a major general purpose technology like computers required an expensive and time-consuming period of restructuring.

Significant productivity improvement from electric motors did not emerge until almost 40 years after their introduction into factories [7]. The first use involved swapping gargantuan motors for large steam engines with no redesign of work processes. The big productivity gains came when engineers realized that the factory layout no longer had to be dictated by the placement of power transmitting shafts and rods. They reengineered the factory so that machines were distributed throughout the factory, each driven by a separate, small electric motor. This made it possible to arrange the machines in accordance with the logic of work flow instead of in proximity to the central power unit.

It has also taken some time for businesses to realize the transformative potential of information technology to revolutionize work. However, the statistical evidence suggests that revolution is occurring much more quickly this time. **C**

REFERENCES

1. Brynjolfsson, E. and Hitt, L. Computers as a factor of production: The role of differences among firms. *Economics of Innovation and New Technology* 3, 3–4 (May 1995), 183–199.
2. Brynjolfsson, E. and Hitt, L. Paradox lost? Firm-level evidence on the returns to information systems. *Manag. Sci.* 42, 4 (Apr. 1996).
3. Brynjolfsson, E. and Hitt, L. Computing productivity: Are computers pulling their weight? MIT and Wharton working paper, November, 1997; ccs.mit.edu/erik/.
4. Brynjolfsson, E. and Hitt, L. Information technology and organizational design: Firm-level evidence. MIT, Stanford, and Wharton Working Paper, January 1998; ccs.mit.edu/erik/.
5. Brynjolfsson, E., Van Alstyne, M., and Renshaw, A. The matrix of change. *Sloan Management Review* 38, 2 (Winter 1997), 37–54; ccs.mit.edu/mo
6. Brynjolfsson, E. and Yang, S. Intangible benefits and costs of computer investments: Evidence from the financial market. In *Proceedings of the International Conference on Information Systems* (Atlanta, GA., Dec. 1997).
7. David, P.A. The dynamo and the computer: A historical perspective on the modern productivity paradox. *American Economic Review Papers and Proceedings* 1, 2 (1990), 355–361.
8. Dewan, S. and Min, C. The substitution of IT for other factors of production: A firm-level analysis. *Manag. Sci.* 43, 12 (Dec. 1997), 1660–1675.
9. Drucker, P.F. The coming of the new organization. *Harvard Business Rev.* 66, 1 (Jan.–Feb. 1988), 45–53.
10. Lichtenberg, F.R. The output contributions of computer equipment and personnel: A firm-level analysis. *Economics of Innovation and New Technology* 3, 3–4 (May 1995), 201–217.
11. Malone, T. IS empowerment just a fad? *Sloan Management Review* 38, 2 (Winter 1997).
12. Schrage, M. The real problem with computers. *Harvard Business Rev.* 75, 5 (Nov.–Dec. 1997), 178–183.

ERIK BRYNJOLFSSON (erikb@mit.edu) is an associate professor at the MIT Sloan School of Management and a visiting professor at Stanford Business School.

LORIN HITT (lhitt@wharton.upenn.edu) is an assistant professor of Operations and Information Management at the Wharton School, University of Pennsylvania.