Conclusion: Computer, Productivity, and Wages; Reflections on the Economics of the Information Age

Timothy F. Bresnahan

The arrival of a new general-purpose technology, with its prospects for widespread changes in product and factor markets, is an occasion to reflect on what we know about the economic impact of technology. The analysis in this volume comes at the time when the commercial internet is at the earliest stages of its role as a GPT, but when computing more generally is at the fifty-year mark. This is a particularly opportune time to confront the empirical puzzles associated with the economic impact of computers so far, and to turn our understanding to the future. In this closing chapter I seek primarily to make that bridge in time.[1] This tasks appears daunting because of the puzzles about the past brought to the foreground in this volume. Are computers (and information technology more broadly) a highly valuable general-purpose technology spreading spillout benefits throughout modern economies? Does computerization of the workplace make up the skill-biased technical change we need to explain recent labor market events? I argue that detailed and specific answers to these questions form the only sound basis for using our understanding of the recent past to understand the near future. How might the commercial Internet affect the growth rate of economic well-being, and how might it affect the demand for labor? I will attempt to undertake the leap from past to future with regard to both the productivity and the labor market issues.

I focus here on the micro–macro puzzles in the puzzling relationship between computers and the economy. There is an important distinction, revealed in the studies in this volume and elsewhere, between the output and productivity results, which show an important micro–macro puzzle, and the labor market results, which do not. For both productivity and labor, we see a strong relationship in the micro data. Studies of productivity see sharp difference between firms that use more information technology capital and firms that use less. The more computer-using firms are likely to be more successful on a variety of measures, including productivity. Yet this micro result does not appear as a time series macro results, for measured productivity has not grown rapidly in the era of computerization, with the exception of the late 1990s.[2] In contrast, the macro and micro evidence on the labor market appear more consistent. Computer-using firms use relatively more skilled workers, and the era of computerization of work has seen a surprising shift upward in the demand for highly skilled workers relative to those with fewer skills.

For the micro productivity results to be right, there need to be some benefits of information technology that can be measured at the firm level but not in the aggregate. Obvious candidates include difficult to measure product quality improvements from computerization, which tend to raise a firm's demand at the expense of competitors (thereby becoming measurable in micro data) but which elude the quality-adjustment efforts of productivity and output measurement frameworks in the macro data. For the macro productivity results to be right, there must be a measurement problem at the micro level. Obvious candidates include reverse causation based on micro-level heterogeneity. Some firms may be more successful than others for reasons that have nothing to do with computerization. The more productive firms may choose to invest in information technology for whatever reason (perhaps computers are fun for managers, and successful firms have managerial slack?).

The labor market results do not immediately bring forward a macro–micro puzzle to reconcile, for there is an effect both at the firm level and the economywide level. They do, however, force us to confront the contradiction between the macro-level results on productivity and those on labor. How did computers manage to shift economywide labor markets without shifting output markets? Further, there is heterogeneity at the firm level in both the output and the labor market results. Why does it wash out at the aggregate level for the productivity measures and not for the labor results? What is the nature of the heterogeneity across firms that affects their labor demand and their productivity; is it a single cause or a coincidence?

Reconciling these puzzles of the past involves understanding the forces pulling computing into highly productive uses as well as the limitations on its productivity. Especially important is understanding the sources of heterogeneity at the firm level. I argue that they are consequences of heterogeneous adoption costs for a new and valuable technology. First, we should understand the first fifty years of computing as the reasonably rapid diffusion of a series of increasingly better and cheaper technologies. Second, whether we think of these technologies as skill-biased or as productivity enhancing (or both), we should examine the forces that have limited the extent of their diffusion. I argue that the limitation has come at the boundary of organizational computing (as distinct from technical computation or personal computing). As a result I argue that the real promise that comes from the commercialization of the Internet arises from new opportunities to extend organizational computing into new spheres: interorganizational computing. But getting to that result about the future calls for a review of our knowledge of the past.

7.1 Invention in Organizational Computing

I base my explanation of the productivity results and of the labor demand results in the same specific theory of the effects of IT. I focus on organizational computing (OC) and specifically co-invention by firms using IT in business information systems. Business information systems are productive and important applications of organizational computing, which is a new (fifty years old) technology for white-collar bureaucracies. In sectors like finance, banking, brokerage, and insurance, OC is the technology of production. In much of the rest of the service sector, OC is a crucial technology for control of production processes or is the production process for key subtasks. In the rest of the economy, OC serves as a control technology for many production processes and is the production technology for the buying of inputs and the selling of intermediate outputs.

Transacting, remembering, allocating, budgeting, planning, rewarding performance, and managing are all heavily influenced, or accomplished, by organizational computing. Although OC is about fifty years old, the underlying powerful computer, data communications, and software technologies have undergone constant and rapid technical change.[3] Adoption of OC began in those firms, and those functions within firms, where there are large-scale white-collar processes. These were in such industries as financial services, and in such functions as accounting in other industries. OC has since diffused to many more kinds of firms and functions.

A firm that desires a new business information system and that decides that IT has advanced enough for that system will nonetheless bear large co-invention costs to design and implement that system. By co-invention I mean invention closely complementary to the use of IT that takes place in the IT-using firm itself. For business information systems, co-invention often takes the form of linked improvements in the production process and in difficult to measure elements of output, such as response time, availability, or reliability in provision of a service.

Co-invention involves two things that are not conceptually trivial. First, it involves the translation of an underlying technical opportunity—an advance in computer or networking hardware, or in fundamental software such as database management systems—into business systems that produce the same output using fewer resources (sometimes) or which produce a better quality output (often). The translation is itself not trivial, nor is the conceptual leap of understanding how a firms' customers will value the "soft" output attributes that OC is particularly suited to, such as timeliness, convenience, flexibility, and availability of the output. Second, changing the organization using the business information system often involves re-organization in several senses. Job definitions, division or even firm boundaries, incentive schemes, and other elements of the organization's structure and functioning will be different in the computerized business information system than in the preexisting one.

Effective use of IT in organizational computing calls for substantial co-invention by the using firm, and this co-invention takes time and resources, including flashes of inventive brilliance that become real in changing work systems. Inventive brilliance that changes everyone's job is not usually the cheapest resource in organizations. Deep understanding of customer valuation of radical changes in "soft" output attributes is often scarce as well.

There are at least two important implications of expensive and uncertain (calling for brilliance) co-invention at the firm level. The first implication is that the rate of technical progress in IT-based production processes like OC is slower than the rate of technical progress in IT itself. Economywide advances in computer, software, or telecommunications technology are just one step in a multistep process of invention. To become valuable in use, they need co-invention. Further IT-using firms are very aware of co-invention costs, which are one of the central points of discussion in the professional and business literature of IT users. Because co-invention costs, and the success of attempts at co-invention, vary across firms, there is substantial heterogeneity in the use of IT in the cross section of firms at any given moment. The heterogeneity is driven, in no small degree, by differential success in co-invention and differential ability to co-invent.

Co-invention is linked both to productivity (via unmeasured output) and to skill-biased technical change (via organizational change).

7.2 Output Measurement and Productivity

Let me illustrate what I mean about the difficult to measure elements of business information systems' output contribution with a simple and familiar example, the automatic teller machine. This banking technology is about half as old as computers, just over a quarter of a century, and uses both computer and data communications technology intensively. Its value to the consumer arises largely through flexibility and availability or convenience.

Withdrawing money at night, or in a grocery store, and with a quick check on account balances, offer considerable product quality advantages over traditional banker's hours and locations.[4] Banks do not find these service quality improvements easy to value, though they obviously think they are large (partially based on internal studies of higher customer attraction and retention rates). By their technology investments, banks reveal their assessment that the service quality improvements will be beneficial to customers. They hope that these customer benefits will attract customers and contribute, when the returns are not competed away, to profitability. A bank that successfully innovates in a new customer service system—as ATMs were in the 1970s, and improvements to them were later on—would expect to see its market share rise and, typically, its profitability rise as a result. None of this return is captured in the productivity measurement statistics. It would, however, show up in micro-data studies that compare across firms.

Some elements of the banking anecdote are systematic. Many business information systems have customer convenience, flexibility, and availability as their primary value. Measuring the private return to those values is difficult—the relevant engineering professional literature talks of the "problem of metrics." Whereas cost savings would be measurable, at least *ex post*, the value of service quality improvements is difficult to measure. The anecdotes are not at all limited to the difficult-to-measure sectors of the economies like financial services. Manufacturing firms, for example, compete on delivery terms, availability, and so on.[5] Just as is the case with the improvements resulting from the ATM in banking, these returns are not captured by productivity statistics. Further the nature of the returns suggests a substantial spillout to society at large. Customers get much of the return to OC improvements in product quality, particularly after lagging IT-using firms catch up to the leaders in their industry and undercut the leaders' ability to obtain a price premium for their superior quality.

These sources of value from OC in use and the costs interact with the difficulty of co-invention. Because of the latter, there is considerable variety across firms in the degree to which they adopt the latest OC technologies. This variation shows up particularly between firms that are leaders and followers in economic success. After a period of time, however, the adoption process allows the followers to take advantage of the technical advances that leaders had adopted earlier. These could be ordinary mechanisms of technology diffusion, as decreases in the price–performance ratios of IT can make certain technology choices attractive to ever more firms over time. Co-invention that is idiosyncratic to a firm may even spill out to other firms over time. Some mechanisms are immediate and direct, such as in the consulting and custom programming sector which enables OC knowledge to be spread from one client firm to another. Further technical progress may lower the costs of co-invention over time, as in the invention of a database management system that dramatically simplified the creation of many OC applications. Other classes of OC application become embedded in software products, as in SAP's enterprice resource management products. A firm that acquires such software can avoid co-invention costs to the extent earlier co-invention is embedded in the software. All of these forces mean that there are substantial increasing returns to OC at national and international levels, as the costs of invention and co-invention are both, with a lag, spread out over many firms.

There are important limitations on the valuation of IT in use. The first I have already discussed at some length. This is the costs and delays associated with co-invention of business information systems. Those costs and delays result in the benefits of improvements in the underlying hardware, software, and communications technologies lagging several years behind the initial invention, even as much as a half a decade.[6] OC has improved in a series of small steps, each setting off a wave of co-invention that diffuses throughout the economy over a matter of years (not decades). The scope of application of OC has also been limited by the fixed costs of designing and implementing business information systems. Within the firm, this means that OC's benefits are far higher for processes which are repetitive and used by many workers, and lower for processes which are infrequent or limited to individuals. That limitation comes from the large fixed costs associated with OC technologies, such as mainframe computers, themselves, and the fixed costs of co-invention, such as training individual workers to use OC applications. A related limitation arises at the boundaries of the firm. Many valuable applications of OC cross firm boundaries. Traditional organizational computing technologies are better at mediating the transactions between trading partners who deal in standard transactions at a very large volume, such as the automated parts of bourses and the reservations systems that are usable by travel agents (regular trading partners of airlines). They have been less effective for transitory commercial relationships because of the fixed costs of design and training.

The personal computer was greeted as the end of these limitations. This turned out to be partially correct. PCs are strong were OC technologies are weak. PCs are smaller, cheaper and easier to use, than mainframes. This permitted use of computing power in tasks such as document production, simple accounting using spreadsheets, and desktop publishing, where the preexisting OC technologies were effective. These "personal productivity applications" are largely separate from the organizational productivity applications built around OC. From a macro perspective, the stand-alone PC is a cheaper form of computing that, while valuable, has diffused to lower-value uses than the earlier OC technologies.

From the early days of the PC, those designing business information systems wished for the opportunity to draw on the strongest features of both OC and PC technologies. This "best of both worlds" vision was announced for a number of technical advances, for it was responsive to a very high value demand opportunity, the limitations on OC. Business information systems within the firm or crossing firm boundaries that connected to the central OC applications but that could be used intermittently, causally, or with little training would be highly valuable. For many years, however, the new technologies that were invented to support these kinds of applications failed to deliver on "best of both worlds." It is the commercial Internet that finally will deliver on that process, a topic to which we will return shortly.

In arguing here that there has been a substantial social gain to IT in use, I do not mean to make any argument about the productivity slowdown whatsoever. There is the rest of aggregate technical progress to deal with, and the rest of the macroeconomy's growth to understand. I argue merely that the social gain to computerization has been very large, that it has covered most sectors of the economy,[7] and that it is badly captured by the aggregate productivity statistics. These remarks are only distantly related to the behavior of productivity aggregates over time.

In sum, organizational computing has been an important technology in producing gains for consumers. It has also, however, had substantial adjustment costs at the firm level and as a result diffused slowly and unevenly. Important technical limitations have given OC a range of applicability that, while substantial enough to generate a social return to invention at the economywide level, leaves obvious boundaries beyond which applications have not yet penetrated.

7.3 Labor Markets

Over the last several decades, demand for highly skilled labor has grown substantially more than demand for less skilled labor. When not offset by countervailing movements in labor supply, this labor demand shift has been the proximate cause of a widening spread in income distribution. The literature on economywide labor markets identifies skill-biased technical change as the source of large changes in labor demand. OC is a technology that effects enough of labor demand to form an important part of the macro labor market explanation. Much of modern employment is in the white-collar bureacracies where OC is an economywide and substantial piece of technical change. The key to understanding OC as skill-biased technical change lies in the co-invention of business information systems. OC does not merely involve the use of computer and telecommunications technology in the firm. Instead, the firm makes substantial changes in the organization of its workforce. The specifics of these changes vary from firm to firm and industry to industry, of course, but there are strong systematic tendencies toward complementarity and substitutibility relationships with different types of labor. This section examines those relationships with an eye to the question of whether OC is old enough, widespread enough, and sufficiently linked to labor demand to be an important part of the skill-biased technical change seen in the macro labor market.

Limited Substitution

The first mechanism by which organizational computing has been skill-biased technical change is that of limited substitution. Over the first fifty years of computer use, routine bureaucratic work has gradually become systematized and partially automated through IT-based organizational computing. As a result many of the clerical tasks once performed by low- and medium-skilled workers have been automated.

The relevant tasks here are not so much handling of paper, which later becomes handling of data, but the replacement of a great deal of remembering, recording, and rule-based decision making. Is it time for a letter to a customer, reminding him his bill is overdue? That task was once the province of accounts receivable (AR) clerks, now of AR software. There are many such routine tasks in white-collar bureaucracies. They can be at least partially automated because the essential elements of the tasks can be carried out by people or by business information systems.[8]

The notion of "partial automation" is important here. It is rare that the computer completely replaces human workers in performing such medium-skill tasks. Instead, OC changes the whole business information system so that work tends to be shared between people and computers in a variety of complex ways. The computer system is quite good at remembering a large number of small things, like how many times customer X has been late with his bill. Continuing to use the AR example to stand for much of OC, I note also that the computer is good at making the same routine decision over and over again, like "payment is late 30 days, send bill." At a task level then, computers are good substitutes for human workers at repetitive, routine, but care-intensive tasks—the center of work for medium-skilled and low-skilled people in white-collar bureaucracies.

As technical progress in OC has advanced over the last half century, the number of such automated tasks has slowly but steadily increased. One driver is that OC has gotten cheaper, so firms find it even more worthwhile to automate tasks. This is partly due to exogenous technical progress in electronics, both computers and telecommunications, as I suggested above, and partly to network effects among firms using OC and scale economies in the invention of OC technologies. Another cumulative piece of technical progress subject to some network effects arises in connection with the automation of work. As experience with OC has grown, computers have gotten better at making more complex decisions because they have been programmed to make more complex decisions as the number and complexity of databases has grown. Relatedly firms have learned over time how to design business information systems to provide better quality service to their customers (the accounts receivable program sends a far sweeter dunning letter to late payers the firm deems valuable customers, the inventory control system directs production toward those customers' urgent needs, etc.). All these forces have expanded the tasks that can be undertaken by routinized, rule-based systems, permitting more substitution of computers for human cognitive skills.

That substitution has historically been limited by perfectly ordinary diminishing marginal returns considerations, as parts of a large complex task are shifted one by one from human to machine function. One source of diminishing returns is at the margin of substitution with human cognitive capabilities. OC has been a more effective substitute for modest bundles of human cognitive skills than for large bundles. The routine and the repetitive have been more automatable than the carefully thought through. As a result there has been far more substitutability for the less highly educated than the more highly educated workers.

A second source of diminishing returns is the rest of human skills, the noncognitive ones. Much clerical work has been changed into being the eyes, ears, voice, and fingers of the computer. While the computer can learn that the payment is late, the computer cannot call the customer on the phone to learn why, nor are computers quite as good as (some) people in dealing with an irate customer who is late paying her bill. Thus noncognitive or "people skills" have not been the subject of much substitution.

The "people skills" limitation has mattered for the impact on labor demand. As I said above, one of the main purposes of OC-based business information systems has been to improve service quality (broadly understood) and to make bureaucracies more rapidly responsive. This leads to a need for people in the OC-based business information system with good communication skills and other good "people skills."

Further, as complex organizations make decisions faster, the internal organization skills associated with directing and being directed become more valuable. Thus the impact of OC-based production has been to shift labor demand not just from low skill to high skill but to revalue skills of different types. Almost all of this revaluation comes from the co-invention of business information systems, not from the invention of information technology.

There is a lively debate about the exact form of the *organizational* change in firms that is at the heart of all this co-invention from the narrow perspective of hierarchy. The routinization of some aspects of work has led both to increases in the degree to which low- and medium-skilled workers are controlled by the business information system and to increases, in slightly different contexts, in the extent to which the logic of the revised organization calls for increases in autonomy and independence. We do not need to reach a conclusion on this issue to understand that the organizational change at the firm level, understood more broadly than hierarchy, is the primary mechanism by which labor demand is shifted.

Complementarities with Skill (in Large Bundles)

There is another important mechanism that reinforces the view that labor demand has been shifted from the less skilled to the more skilled. Routinization and regularization and partial automation of the low-paid part of white-collar work is the direct impact of use of OC. This leads to bureaucratic production processes that are more thoroughly understood by the firms that contain them—or at least more systematically and more analytically understood. An OC-based bureaucratic production process is also more directable by the firm that employs it. There is a mechanism for changing the entire system. Further, OC-based bureaucratic production produces a great deal of information—about workers, about departments, about customers and suppliers and decisions. The information is an input into more analytical decision processes about production and allocation in the firm.

The widespread change in bureaucracies toward more analytically understood production processes, toward more directable production processes, and toward a flood of information have all led to an increase in the demand for analytical skills in the firm. The control of OC-based technologies calls for large bundles of cognitive skills, for people who can conceptualize and solve problems, who can understand the goals of the organization as well as its rules, and so on. This is the technological root of a *complementarity between computerization and highly skilled workers*. It is a source of a rising demand curve for highly educated labor.

Counterintuitively, as managerial work grows more analytical, the demand for managers with high levels of human interaction skills grows as well. Interactions with subordinates that are more analytical—perhaps mediated through explicit monitoring and rule-based incentives—make it more difficult to undertake the human part of motivating, leading, counseling, and mentoring. Even more, the change of the organization to a systematized one calls for managing and being managed work that makes the junior people feel valued and motivated.

Note that the complementarity arises because the white-collar production process is OC-based, not because the highly skilled worker herself uses a personal computer. In my story, the shift from low-skilled to high-skilled labor demand does not arise because of an effect that arises at the individual worker level. Instead, it arises at the firm or organizational level. The widespread diffusion of the personal computer, and the associated personal productivity applications (e.g., spreadsheets and word processing) does not play much of a role in my story of changed labor demand.[9] As with the discussion of economywide productivity and personal computing above, the main point is that the stand-alone PC and the individual productivity application represent marginal use of a technology after it had grown cheap, not the extension of high-value uses to a new domain.

The end result of the labor demand shifts associated with OC has two parts. First, there is a shift in the demand for human cognitive skills from the less skilled (less educated, less trained, or simply less smart) to the more skilled. Second, there is an across-the-board shift in the demand for "soft," "interactive" or "human" skills. To the degree that high levels of cognitive skills and of interactive skills present in the same person are particularly rare, this leads to an even larger increase in the price paid for that skill bundle. A potential important result for my analysis of labor markets is this story of failure to unbundle: it is plenty of intelligence combined, in the same person, with plenty of people skill that is scarce and highly rewarded.

7.4 The Future

Despite the early stage of commercialization of the Internet at present, it seems to me that enough can be seen of technical opportunity at this stage, and enough has been learned from the past economic impacts of computing, to make some preliminary forecasts about the size and nature of its economic impacts. Let me begin with a forecast of the direction of technical change. The commercialization of the Internet permits a combination of previously separate elements in information technology.

The Internet brings together much of the power, system, and connectedness of traditional organizational computing with the small scale, ubiquity, and ease of use of traditional personal computing. As a result some of the boundaries limiting traditional organizational computing, emphasized throughout this chapter, will now break down. While it is never possible to forecast all of the implications of such an advance,[10] enough can now be foreseen to say much about the future economic impacts.

One should, in general, be wary of that kind of "best of both worlds" forecast, for it has been heard before. Because of the well-known limits of organizational computing technology, the users of information technology have been calling out for a combination of OC and PC attributes since the early 1980s. Several not very important technologies have been announced, for marketing purposes or out of general optimism, as bridging the gap between OC and PC. At this juncture, however, we may safely forecast, for there is enough co-invention activity among the users themselves, not merely hopes on the part of technologists, to see some of the utility of the recent technological advances.

We are at the beginning of a new wave of organizational computing. This time, the relevant organization is larger than the firm, including suppliers, customers, contractors, and workers. We might call this new wave *interorganizational computing*. It has a solid foundation in recent changes in networked computing. The technical progress we have already seen in connection with the commercialization of the Internet forecasts the removal of the most important purely technical limitations on the diffusion of organizational computing in the past, difficulty of use and difficulty of low-scale (casual, intermittent, or merely associated with small trading relationships) use. Together with ongoing improvements in computing and communications power, these new technical directions will permit the construction of new IT applications that cut across organizational boundaries.

The widespread distribution of the technologies that would permit this kind of valuable advance was cut short by Microsoft's anticompetitive campaign against commercial Internet technologies. The widespread entrepreneurial invention of technologies to support "best of both world" applications that characterized 1994 to 1998 has been cut off. Nonetheless, Microsoft promises us that it will deliver such technologies in a form that preserves its monopoly position only about ten years later than the entrepreneurs would have. The demand opportunity remains large, so the delay, while costly, does not mean the end of the long-run opportunity for a new positive feedback loop between invention and co-invention in applications that cut across organizational boundaries.

The value of automation and systemization in this area is clear. Much of distribution and communication about it has the kinds of obvious inefficiencies that call out for rationalization. Many middleman industries have market power that has accrued along with (one imagines) socially valuable functions such as learning and acting on reputations for reliability in carrying out contract terms. The new technologies permit reorganization and thus removal of the market power in an improved information environment in which achievement of the other functions is undertaken by market participants directly. Many supply relationships are longstanding because of the costs of identifying trading partners who are reliable, honorable, and so on. The resulting inflexibility could be avoided if information relevant to that identification were to arise in a lower-cost way. One could build examples like this at some length. The general point is this: much of what goes on in markets and supply chains is information processing now. Also many of the costs associated with running markets and supply chains are costs familiar to information economics: the costs of agency, of monitoring, of reputation-building, and so on. Information technology, widely deployed in these new domains, has the prospect of being highly valuable.

On the other hand, the very nature of that value suggests the need for a great deal of difficult, and sometimes brilliant, co-invention. Market institutions for reputation, assurance, resolution of disputes, and so on, have been built up over long years of trial and error. Replacing a complex web of human relationships and institutions with new business information systems will not be trivial. It is a very safe forecast that the benefits of automating markets, while huge, will be realized by a process that is uneven and sometimes surprising. Co-invention, whether by buyers, sellers, or new intermediaries, will take time, effort, and invention.

As benefits are realized in leading examples—leading markets and supply chains in this instance—they will feed back to technological developments in IT. Those in turn will, later, enable more co-invention and adaption in other parts of the economy. There is no reason to suspect that a critical aspect of information technology use in the recent past, heterogeneity between leading and following adapters, will decline.

We will continue to see a pattern in which firms vary considerably in their productivity, positively correlated with IT use. The positive correlation will continue to be a blend of rents to successful co-invention and of complementarities between IT and co-invention. The social gains to all this new invention will be substantial in the long run.

Much of the dot-com boom was a bet that my last few paragraphs would be wrong.[11] The events of 2000 and 2001 show us how those paragraphs are right. Many dot-com firms planned to replace much of the distribution system—intermediaries, or the buying and selling bureacracies in firms—in a radical shift to a new, IT-based form. This was a bet that co-invention was not necessary for interorganizational computing. New entrepreneurial firms would cast aside the knowledge that lived in the existing bureacracies and quickly achieve widespread success with new kinds of organizations. That bet, whether it was ex ante wise or not, has been revealed to be wrong. Instead, we will see that the widespread diffusion of interorganizational computing, just like its predecessor OC, will take time and co-invention. The short-run prospects therefore are of a period in which established firms undertake improvements to the organization of buying and selling, engaging in substantial risky and difficult co-invention. That short-run forecast, however, takes nothing away from the long-run forecast of where that co-invention will ultimately lead, toward highly valuable improvements in the interorganizational computing technology supporting buying and selling.

Some of the implications for labor demand seem forecastable. If we look to the production processes that will be affected by interorganizational computing, they are the ones involving buying, selling, and co-ordinating across firm boundaries. Some economists may think that this is a narrow and specific kind of enquiry I am suggesting, to look at specific production processes like those. Yet these production process involve, depending on how one does the definitions, approximately a third of the employment in the rich economies. There is plenty of scope for economywide impact. Further there remains a great deal of unautomated or only partially automated work in the white-collar bureaucracies that border the firm, buying or selling. While many of the production processes at the center of firms and industries have been very successfully automated, the ones at the boundaries of firms and beyond have, so far, changed less. Much of the work in these parts of bureaucracies is very suitable for automation using IT. Much of the white-collar work itself remains simple, suitable for partial automation. Many of the tasks these bureaucracies undertake involve remembering, making sure, coordinating, and controlling, which are exactly the tasks for which IT is quite suitable. The new technical progress opens up the geographic and organizational span of those tasks considerably. Further the prospect of flexibility with control that IT offers is even more important with the new technologies of IT that the commercial internet brings us. Transitory relationships with many suppliers, for example, are more possible after the technology of connecting has lower relationship-specific fixed costs. It seems likely that the forces of limited substitution will still be at play in the future as they were in the past, on that wider geographical and organizational span.

These new applications will exhibit a familiar *complementarity with analytical management*. The reasons will be much the same for the interorganizational computing era as for the OC era, and will if anything grow more important. Interorganizational commercial relationships are more complex than are organizations, as they extend across more domains of knowledge. A typical interorganizational application will involve buyers as well as sellers, and perhaps middlemen as well. This is inherently more complex and cognitively difficult.

Extending the reach of organizational computing technologies across the organizational boundary, to smaller and more fleeting economic relationships than those we typically see within the firm, and to a great deal more of economic activity, seems very unlikely to change its cognitive or business-social character. Let me turn first to the business-social elements. It seems to me likely that the extension of people management across boundaries, and the resulting extended span of these new applications, will exhibit a familiar bias toward soft skills or people skills. Negotiating and managing across firm boundaries has at least as many traps and pitfalls as that within. In the present we see firms deploying employees with above-average people skills toward relationships with other firms or with workers or consumers. Further, to the extent that interorganizational computing fosters a shift toward more flexible and transitive economic relationships, it will likely push either toward labor market institutions that are based on flexibility or toward work organizations that have inflexible employment relationships but are highly redeployable. In either case, many of the soft skills and people relationship skills that have recently seen rising wages will likely grow ever more valuable.

I would summarize the force of the last three paragraphs as saying the past impact of computing on the demand for labor by skill type will likely continue into the future as to direction. We can expect the relative demand for highly cognitively skilled workers to rise relative to less cognitively skilled workers, and the relative demand for workers with good social or interactive skills to rise as well. How far, how fast? It would be foolish to forecast either the tendency to pay a premium for cognitive skills or for social ones as advancing at any particular pace as the exact size and timing of the relevant technical change—co-invention—is very hard to predict.[12]

I take that uncertainty as a wonderful research opportunity. The introduction of IT changes the information economics of firms, markets, and supply chains. Yet information economics is notably subtle, and adding information to a system can change its institutions in a variety of directions. As interorganizational computing becomes an important process technology, we have a wonderful opportunity to watch its economic impact, particularly its impact on labor markets. A related research opportunity comes from the opportunity to study the theory of organizations. Much of the successful empirical work in that area has focused on the boundaries of the firm, or on relationships between closely linked firms. There is now the opportunity to see those boundaries redrawn and those relationships redesigned on a wholesale basis—and by business people who would like to be studied, for they have no clear idea where they are going.

Notes

1. I am grateful to Jacques Mairesse, Yannick L'Horty, Nathalie Greenan, and several conference participants for helpful comments, and especially to the late Zvi Griliches for discussions full, as always, of gentleness and wisdom.

2. The result does reappear in a simple cross-sectional comparison across countries, but the United States is very different from European countries on so many dimensions as to undercut the value of this observation as a tiebreaker.

3. Sometimes major branches of computing are categorized by the kind of computing used so that organizational computing was for years called mainframe computing. I wish, however, to do a categorization by type of *use*, and not technology. The partial replacement of mainframes by smaller computers in OC leaves the uses in place, for example.

4. The other benefit of the ATM is as a cost-saving technology, literally an automatic teller. To take advantage of the service-quality improvements, however, ATMs were deployed far earlier in time, and to far more locations, than cost savings alone would justify.

5. Comparing easy-to-measure with hard-to-measure sectors of the economy is thus ineffective as a method for learning anything about the valuation of IT in use.

6. The idea that this is a half century, and that what has been going on so far in the first fifty years of organizational computing is a low-payoff activity that will now finally blossom into a high-payoff one, is laughable.

7. Limiting the value of cross-industry research in understanding the productivity impact of IT, unfortunately.

8. A recurrent debate about whether "computers can think" sometimes confuses onlookers. It is irrelevant to the substitution story. Checking whether a bill is overdue and calculating the overdue penalty can be done either by a person (who is thinking) or a computer (which may or may not be thinking).

9. This analysis applies primarily to such mainstream personal productivity applications as the word processing program and the presentation-slide program. While useful, their impact is limited by the span of the individual worker's influence. The most common uses, word processing and spreadsheet, do not tend to have dramatic impacts on the productivity of the individual. There are some specialized individual productivity applications that do have such an impact, such as those used by com-

puter programmers, graphic designers, and so on. These are a small fraction of workers, however.

10. Broader changes to society, changes in the nature of individuals' work, changes in the autonomy/control nexis, changes in social interaction, and so on, are very difficult to forecast.

11. I am proud to report that they are printed here just as I spoke and wrote them years ago.

12. Authorial pride once again leads me to point out that sentence pre-dates the dot-com crash by several years.

This page intentionally left blank