The Mechanisms of Information Technology's Contribution to Economic Growth

Timothy F. Bresnahan*

For presentation at the Saint-Gobain Centre for Economic Research

Conference "What do we know about the sources of technical change?"

The mechanisms by which business use of information technology contributes to growth in living standards are badly understood, leading to a confused debate. In this paper, I start from the microeconomic and management literatures that study use of IT at the individual firm level. IT-using firms, and the IT vendors who support their efforts, are the agents whose incentives, innovations, and opportunities matter in the first instance for understanding the mechanisms. Aggregating those agents' incentives toward a nationwide or worldwide level leads me toward growth. My goal is to explain what existing knowledge about the economics of the IT-using firm and IT markets has to say about the concerns of growth theory, rather than to provide evidence for or against a particular view.

Modern aggregate growth theory analyzes the macroeconomic and microeconomic conditions for sustained economic growth. From the macro perspective, the key idea has to do with accumulating assets of sufficient economy-wide importance to raise aggregate well-being yet not subject to diminishing returns over an extended period of time. The relevant assets can be a stock of capital, human or physical, or a stock of scientific or engineering knowledge, or a combination of all of these. What is critical for growth to continue is the existence a force raising the return to investment in the accumulating asset over time.[1]. Without such a force, diminishing returns set in.

The microeconomic underpinnings of this theory have drawn particular attention to stocks of knowledge. Since knowledge can be re-used, with sufficient leverage it can have economy wide implications. Cumulative knowledge, together with a large technical opportunity, can raise the social return to learning for a long time. The microeconomic part of the theory thus draws our attention to increasing returns to scale and to the strength of complementarity between early knowledge accumulation the same activity later on, i.e., to complementarity over time. The same reasoning leads to gaps between the social and private return to invention of new knowledge. Understanding both macroeconomic and microeconomic perspectives is key in this analysis, for only then can

* Stanford University. Gordon and Betty Moore Senior Fellow, SIEPR, and Professor of Economics. Contact at tbres@stanford.edu

[1] A number of theories with very similar formal structures variously draw our attention to aggregate increasing returns (in several forms) to formalized R&D, to learning by doing, product invention in "quality ladders," to education and human capital formation, and to many other things.

Modern aggregate growth theory has not been written in a way which leads toward empirical analysis. Indeed, most empirical work proceeds either by comparing growth across countries or by examining changes in growth across times in the same countries , attempting to predict them by the assets identified in the theories, surely two of the least promising empirical strategies ever used in economics.

the mapping from the incentive effects of economic policy to their growth impact be understood.

Information technology has key features related to the key structures of the macroeconomic and microeconomic approaches to growth. The first is substantial technological opportunity in the narrowest engineering sense: it is physically possible, with enough resources, to make computers and telecommunications equipment perform better and better. The second is a potential for broad applicability and therefore economy wide impact. IT is a general purpose technology.[2] Its most widely useful attributes, such as ever-cheaper computation, data processing, and bandwidth, underlie a wide array of more specific technologies that are in turn used in a number of very different applications. Like all GPTs, IT advances enable but do not compel improvements in applications. Applications themselves are the product of complementary co-invention by IT-using firms. Advances in IT shift the innovation possibility frontier of the economy rather than directly shifting the production frontier. The third reason is also related to this complementarity between IT invention and co-invention of applications. Inventions in applications increase the size of the market for IT, improving the economic return to IT invention, which is subject to considerable increasing returns. Typically, inventions in applications lag a few years behind IT inventions. This injects a dynamic feedback loop in which IT advance leads to unpredictable applications invention, which in turn raises the return to improvements in IT whose rate and direction can only be understood after the applications invention is complete.

The related microeconomics follows from the same set of forces. Any GPT is associated with substantial social increasing returns to scale. The externalities among co-inventors are considerable. The externalities between inventors of IT and co-inventors of applications are considerable. The dynamic externalities between inventors of different types of IT are considerable. Attempts to internalize these externalities have been wildly successful on some occasions and a thorny problem to overcome on others. Contractual, organizational, technological or public policy (e.g., IP) schemes to internalize these externalities have been partially successful, and remain as difficult as they are important. As a result it is difficult to completely appropriate the returns to some kinds of invention and to some kinds of co-invention. These contracting difficulties lead to a higher social than private return to much of the invention and co-invention and the possibility that the dynamic positive feedback loop with have weak causation at some junctures. As a result, there is scope for public policy and business policy innovations to affect technical progress (and thus growth) through a wide variety of incentive-altering mechanisms.

Understanding the mechanisms by which these forces rise to economy-wide importance and the means by which they maintain that importance over a long period of time calls for a broad approach. I will take up 5 main topics. Section 1) examines co-invention in narrow IT segments, taking the behavior of the single IT-using firm vis a vis a single technical advance as the unit of analysis. Section 2) examines markets for specific IT capital goods. Those complete the most microeconomic partial equilibrium portions of the analysis. Section 3) looks at the widening span of application of IT over time, and Section 4) examines the mechanisms by which early IT invention – and co-invention – have come to be complements to, rather than substitutes for, later technical

[2] See Bresnahan and Trajtenberg (1995) and Helpman and Trajtenberg (1998).

advance. Section 5) considers the forces holding social return to invention above the private return.

1) IT in Use

Let me begin with three analytical distinctions. We have already distinguished between invention by the sellers of IT and co-invention by the buyers. We need also to distinguish a stock of knowledge, called information technology (IT) from a stock of capital which embodies that knowledge – computers, software, and telecommunications systems -- called information technology capital, ITK. Finally, we need to distinguish between different kinds of applications of IT. I focus on two very different kinds of applications, large-scale applications, typically undertaken by an entire organization, which I call Business Information Systems, and applications where one person is the user, personal applications. While there are other demand segments, these two capture the extremes of the process of co-invention.

Improvements in IT as a *technology* expand the scope of what can be invented by IT users. Some IT inventions are well understood as incremental improvements in existing capabilities, e.g., a microprocessor which can undertake the same functions more cheaply and faster than last year. Other IT inventions, by their nature, enable new or considerably improved applications. I focus here on the process which begins when either accumulation of incremental progress or substantial breakthroughs makes new or improved applications economic. In most years, that happens several times in different parts of IT, marking the beginning not the end of a process that shifts out the production function.[3]

The rest of this Section briefly sketches the next step of that process, co-invention by a single applier.

a) IT in Use—Co-Invention of Business Information Systems

The most valuable IT applications, and those which directly transform the production process of the firm, Business Information Systems (BIS) are associated with substantial co-invention costs. These are costs that a using firm needs to bear up front, before getting any value out of the application. The co-invention costs include the costs of programming to make a technology useful in a business information system, the costs of training workers to use it, and, importantly, the costs of inventing a business purpose for the system and organizational change to use it effectively. Co-invention costs are very substantial. Only a small fraction of the costs of inventing and implementing IT-based applications in large firms are costs of ITK.

The reasons for the high costs of business co-invention are well documented in the literature.[4] Co-invention of new BIS often involves linked improvements in the white-collar production process and in difficult-to-value "soft" output attributes. This co-invention involves the translation of an underlying technical opportunity into business systems, which will produce an output using fewer resources (sometimes) or a better

---

[3] IT vendors announce dozens of such "paradigm shift" advances per month. The actual rate, while far lower, is well above zero.

[4] A substantial management research literature, and a smaller econometric one, have characterized the invention process in organizational applications.

output that customers might choose (often.) The translation is itself not trivial, nor is the conceptual leap of understanding how customers will value "soft" output attributes like timeliness, convenience, flexibility, and availability. Improvements in the white-collar production process often call for changed organizations. Job definitions, division- or even firm-boundaries, incentive schemes, and other elements of the organization's structure and functioning will be different in the more-computerized business information system than in the pre-existing one. This co-invention takes time and resources, including flashes of inventive brilliance. Inventive brilliance which changes everyone's job is not usually the most common or the cheapest resource in organizations. Nor is deep understanding of customer valuation of radical changes in "soft" output attributes.

The underlying features of IT-based production that leads to these substantial adjustment costs lie primarily in the tensions of routinizing and regularizing white-collar work. As databases and business logic have improved and accumulated, bureaucracies have become better over time at two very different kinds of functions. Bureaucracies know more about their workers, suppliers, and customers, and analytical bureaucracies can take advantage of this knowledge to incrementally improve service and incentives. Second, the comparative disadvantage of routinized and regularized processes, inadequate personalization to the circumstances of particular workers, suppliers, or customers, has been steadily lessened, often by combining IT-based systems with a human front end.

Once a BIS is installed and working, it also creates two other assets that may be of sustained importance. Almost all BIS continue to interact with the rest of a human organization, and the system will begin to generate knowledge at the interface. Gripes about the BIS and "wish lists" – guides to future technical progress –- are half of this. The other half depends on the analytical use of the BIS. If there are analytical uses, however, they generate knowledge about how the organization itself (outside the BIS) functions, knowledge which itself is, after it is understood and put to use, valuable in future technical progress. Of course, not all BIS produce such knowledge, and not all managers are capable of taking advantage of it when they do. The creation of these stocks of knowledge in the aggregate, however, has been sufficient to drive the compensation of the kinds of managers who can take advantage of it quite high.

As a result, co-invention takes time and scarce resources, and is the source of adjustment costs at the individual using firm. IT-using firms designing new business systems 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. Since co-invention is risky, co-invention costs include costs of experimentation, costs of failed systems, and so on. There is some learning from experience, and from the experience of other users, so that co-invention costs will sometimes be declining over time (when imitation is cheaper than original co-invention.) Nonetheless, much of the costs of inventing a new BIS are local, as many of the problems that are being solved are local and specific.

After the period of incremental experimentation with a new BIS, any of several outcomes can occur. A frequent one is that the system isn't useful, so that it is either abandoned or goes unused. (Some surveys suggest that this is the most frequent outcome.) Another frequent outcome is that the system is discovered to be useful, and that the discovery also leads to incremental improvement and refinement. In this case, the demand for the services of the system is likely to put pressure on computer or networking capacity, leading to an increased demand for capacity. Much of the demand for ITK is conditional on an assessment of its usefulness as a result of this process. The incremental improvement (and capacity pressure) process may well continue for a period of some years, with resulting "upgrades" to the underlying ITK capacity.

This structure of co-invention activities leads to a very particular structure of costs. Look first at narrowly technical activities such as designing computer hardware or writing computer software. When we count up all the software that is written in connection with a single use, we find that more – far more – of this technical invention on a cost basis is done by buyers rather than sellers of IT[5]. In the aggregate, buyers of IT do somewhere between half an order of magnitude and an order of magnitude more technical invention on a cost basis than do sellers, even in the US[6].

If we then look beyond the technical costs to the business costs, i.e., the costs of inventing new uses for IT and new IT-based organizational structures that can achieve those new uses, we find even more investment by buyers. Indeed, the costs of co-inventing the nontechnical parts of applications appear to be, again by perhaps half an order of magnitude, larger than the total costs of technical invention.[7] The total cost of invention in IT in use can be divided into three categories: most of it is the cost of business co-invention, the second biggest category is technical co-invention, and the smallest category is "technical" invention by the sellers of IT.

The econometric literature has found several results that are consistent with that finding.[8] First, it is the usual interpretation, and likely the correct one, of the very high measured rate of return to IT investment and use at the firm level.[9] In a number of studies of large firms, Erik Brynjolfsson and coauthors have measured that high return. If the lower-IT firms who are not using new IT-based business co-inventions, but instead must overcome them, that is a rationalization of the results. Perhaps more directly related

---

[5] These remarks draw on Bresnahan [2000] which has definitions and sources. The underlying data behind this particular remark are surveys of expenditures on purchased technology vs. technology built in-house in the corporate sector. Thus, this part is literally cost accounting.

[6] I emphasize the cost basis because this does not establish that buyers do more "technical progress" than sellers do. Sellers achieve enormous economies of scale by shipping the same hardware or software design to millions of buyers (this is an important part of the equilibrium that defines the boundary, in invention space, between buyers and sellers) and thus may be doing "more" technical progress in whatever the appropriate physical units might be.

[7] This result appears both in studies of the demand for IT, such as Bresnahan and Greenstein [1997] and Ito [1996] and in studies of the firm-level return to ITK investments, such as Brynjolfsson et al [1996, 1997, 2000]. This is less of a literal cost accounting exercise and more an econometric one. Both literatures have needed to assign large business co-investment costs to ITK to rationalize firm behavior.

[8] Systematic statistical work on shifts in computing architectures has found substantial adjustment costs [Ito, 1996; Bresnahan and Greenstein, 1997], and the case literature on IT implementation highlights difficulties in implementing concurrent organizational changes (e.g. Kemerer and Sosa, 1991 and Zuboff, 1988]. Moreover, there is additional evidence that monetary and non-monetary costs of these adjustments is larger than the capital investments in many cases [Brynjolfsson and Hitt, 1996; Brynjolfsson and Yang, 1997; Bresnahan, 2000].

[9] 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 costs of co-invention. Cf. Bresnahan and Greenstein [1997] or Brynjolfsson et al [2000]. For modeling purposes in this paper I will be more interested in this heterogeneity ex ante than in the uncertainty of co-invention, which I shall treat as an adjustment cost to the user.

to the point in this paper is the literature on diffusion of new IT technologies across large firms. Studies such as Saloner and Steinmueller [1995] and Bresnahan and Saloner [1996] find that business co-invention is a major barrier to early adoption. Of course, barriers to adoption at the individual firm level and high measured returns to adoption are consistent in equilibrium if one believes that the adopters have made a large investment in co-invention.

In one econometric study, Shane Greenstein and I looked at the transition to new architectures for BIS in the late 1980s and early 1990s. This transition was quite marked, and strong enough to undercut the rents of pre-existing dominant vendor IBM. The earliest users to switch, however, were not those who had the most to gain from using the new architectures. Instead, the earliest users to switch were those who had the smallest adjustment costs at the individual using-firm level, often lower value applications later on. While this "reverse diffusion" finding may not be general to all cycles of new IT in BIS, the general point is that obviousness or ease of organizational implementation (low adjustment costs) as well as high value are important determinants of early adoption.

This was a long narrative of what happens at a typical single using firm that began with a piece of IT invention which was the stimulus to the applications co-invention. That single using firm process creates a number of changes. The most immediate of these are that, if the BIS works, the firm demands more ITK and uses it in its own production. That is the simplest part; IT having raised the invention possibility frontier of the economy, complementary applications co-invention raises the production function.

Co-invention in the single using firm also leaves behind a number of assets. More of the production process is routinized and regularized, often with databases and the logic to work on them embedded in software. The using firm will often have learned something about its organization and its customers' demand, knowledge that may not only useful locally but also partially portable to other firms. The using firm will have learned, especially in its technical co-invention, what features it desires in purchased ITK, and will have communicated these to IT sellers.

Many different using firms build these assets in response to any substantial new IT invention. There is enough failure in co-invention, and in efforts to learn the general, broadly applicable lessons of any particular new co-invention, that society is lucky that there are a number of distinct replications. This wide variety of explorations in co-invention increases the likelihood that one or more of them will yield broadly useful new knowledge. While I have been writing about the response to a particular new invention, exploration and co-invention recur regularly in response to new IT inventions.

b) IT in Use – Personal Applications

When the user of an application is not a firm or other large organization but rather an individual, adoption and co-invention have different features. The most important of these is that up front adjustment costs at the individual level are not nearly so important.

Many individual adopters buy systems that are ready to run the applications they will use. An individual's PC can run a word processing program or spreadsheet, or an engineer's workstation a design program, without large investments in technical invention at the using firm, and often without large up-front investments in changing the firms' organization or products. The simple solution of "buy computer and run application" is available.[10] To be sure, not all the value of that application may be available without co-invention. Most of the valuable applications of spreadsheets had to be invented, and all the organizational changes that come from having presentation graphics programs and email had to be invented as well. Yet all these could be simply undertaken as learning-by using based on PC systems that could be installed and used without up front business or technical adjustment costs.

The scope of this characterization of demand is limited to personal applications, not uses of personal computers; it is a point about demand not technology. When PCs or PC software are used as one technical element of a BIS, the BIS will be subject to the usual adjustment costs.

This is not to say that there are no adjustment costs in personal applications areas, merely that they are not primarily borne by the individual adopter. Instead, someone – an application software developer -- has to write the applications software that lies behind the ready to go systems. Applications software developers learn from their customers' learning by using, so that, over time, invention and co-invention together lead to better systems.

Once again, much of the demand for ITK comes from repeat purchasers conditioning their decision to acquire further capital equipment on earlier experiences with it, or, in these segments, from users who have seen other users like themselves adopting the PC. For some kinds of applications, the value of adopting is directly raised when many other users have themselves adopted. Email, for example, is more valuable when one's friends or colleagues use it. Word processing is more valuable when others can read the file formats, and so on. For almost all applications, however, there is an "indirect" network effect. Since the costs of writing an application can be spread over all the users, the more users, the more the economic return to making new applications.

As a result, the structure of IT demand in personal applications is quite different from the BIS we saw in Section 1)a). Here, strong network effects contribute to social increasing returns to scale. The external effects, however, do not arise through individual users' failure to co-invent, but failure to coordinate in creating large markets.

Widespread adoption of particular technologies in personal applications creates human and physical capital. The human capital can be understood simple as "complex computer use." Workers know how to use personal computers to analyze and communicate. The physical capital is partly the installed base of computers, partly the data and documents that are stored on them.

We have now looked at the very partial equilibrium of an individual adopter. Let us now move up to the level of the markets in particular IT segments.

2) Two Dynamic Externality Structures

In the last section, we looked at the individual user's decision to adopt a new IT technology in two kinds of markets, differentiated by adoption adjustment costs of co-invention at the individual user level. Let us now turn to the implications for all users and for sellers of ITK. In this section we continue to examine the implications of a single

[10] This is not to say using firms do not need to undertake technical activities to use PCs. Indeed, the costs of purchased PC hardware and software are just a down payment on the lifetime cost of ownership, which includes maintenance and support. But those are well understood as a cost of having the ITK in place, not as adjustment costs or invention costs.

important new IT invention, but now at partial equilibrium single market level. I maintain the distinction between BIS-supporting and personal application-supporting markets.

Many of the common implications of the two kinds of markets arise from their general-purpose technology structure. The economic return to inventors of IT and/or vendors of ITK arises from large markets with many different users. External effects determine usage. The external effects may be direct network externalities, as when users communicate with one another or share files or ideas for applications. Or the external effects may be indirect, as when the number, inventiveness, and willingness to pay by all users determines the quality of the ITK on offer to each user. In either case, the implications are that a new IT invention is associated with social increasing returns to scale because of the GPT structure.

The distinction between high individual user adoption cost markets and high network effects markets does, however, imply a distinct dynamic for the realization of the social increasing returns.

a) Social Increasing Returns under Adjustment Costs

In the BIS-supporting markets, substantial co-invention costs at the individual use level are a bottleneck after a piece of new IT invention. If early adopters of a new technology are successful in co-invention costs, they typically earn substantial returns. Imitation by laggards follows, with the laggards taking advantage of some of what has been learned from the early adopters' struggles. ITK sales rise as a result of this collective success. IT inventors learn from their customers' experience and improve their offerings in support of what their customers actually want.

This story of co-invention cycles offers the prospect of social increasing returns to scale, arising through two distinct external effects. The first external effects come from falling co-investment costs. Early users' experience lowers later users' costs of co-invention. This is a one-way externality, spilling out from early to late adopters, and there are few available institutions for internalizing it. The one-way nature of the externality creates an incentive to wait.

The second externality flows through the supply of improved technology. With foresighted technology vendors, the rate of technological improvement in IT will depend on the number of total adopters and their likely adoption dates. The more, and the earlier, the adoption, the higher is the incentive for technologists to invest, race, or compete for the business. This incentive for IT inventors to anticipate co-invention may not always be effective. The anticipation is conceptually quite demanding for important IT inventions, as it would involve guessing what will be co-invented. Follow-on inventions and incremental improvements are conceptually easier in anticipation, and for these there is considerable opportunity for IT vendors to advance the entire market equilibrium by anticipatory technical investment.

The timing of social returns to important IT inventions is thus often determined by the timing of the co-invention and spill outs, including IT that follows on co-invention. This tends to lag a while behind the period of early adoption of the technology or the period of the greatest public attention to the technology.[11] There are two lags, the

[11] Friedman and Cornford (1989) call this "the problem of the tenses."

first between invention and initial serious co-invention, and the second, between early co-invention and the widespread spill out of co-invented ideas. Historically, for important IT inventions these lags have been in the half-decade to decade lengths. As a result, there is little reason to look for a contemporary relationship between the introduction of valuable technologies for organizational computing and the resulting value creation, and little expectation that there will be a great deal of coordination across co-inventors. The productivity impact will come later and be spread out over time.[12]

The cycle between co-invention and invention in business information systems often leads to standing relationships between using firms and particular vendors or technologies. Users find that their very valuable co-invention and selected vendor's invention are together designed in to systems. Incremental improvement of those systems using compatible technologies lowers competition but economizes on adjustment costs. Once users begin to chafe under the lock-in, diminishing returns of a different form have set in. Existing technologies may be a constraint on the rate of social increasing returns, but escape from the lock in may require very difficult to coordinate common movement to new technologies, a difficult task if there is considerable co-invention. Such shifts are correspondingly rare and require special circumstances.[13]

b) Social Increasing Returns under Network Effects

While the personal applications segments are not an area of up-front co-invention, they share the same structure of social payoffs in the long run. Many technologies that are used by a single person have this structure. Above, I suggested the PC, but other technologies such as word processing, the browser, email or instant messaging have this flavor as well. Each of these has substantial network effects. If there is a substantial body of applications, many users will want PCs, and if there is plenty of online content, many users will want browsers; and those applications and that content will be forthcoming if the technology is widely adopted. As long as there are plenty of other users to communicate with, users will want email or instant messaging. In that case, the network externalities are direct (users want to communicate with one another) while in the case of PCs or browsers, they are indirect (a large installed base of users will draw rich applications or rich content.)

The implications of this particular structure of payoffs to users are familiar from the literature on network effects, and it emphasizes change in the nature of equilibrium over time. The first phase is before the invention of the enabling technology. The second phase is one of waiting for network effects to take effect – perhaps I don't adopt the technology because you haven't yet? The third phase is one of positive feedback once the network effects are taking effect. If an important motivation for users of a technology is network effects, one should expect some coordination in the adoption time, as the figure shows. That coordination can either start quickly or start more slowly, as the length of the second phase is typically not well determined.

---

[12] Links between particular IT technical progress in the late 1990s – the commercialization of the Internet, the increase in the rate at which integrated circuits grow smaller – and the rapid rate of productivity growth in selected heavy IT-using countries at the same time, such as the US, are very unlikely to be causal.

[13] See Bresnahan and Greenstein (1999) for analysis of many of the major examples.

The transition from a regime of not realizing network effects to one of realizing them can add considerable momentum to these markets. Consider the PC market once it was realized that the spreadsheet and word processor would be the dominant "personal productivity" applications. A series of technical and marketing activities, culminating in the introduction of the IBM PC, put the industry on the high equilibrium in which ordinary business people, not computer hobbyists, knew of the PC and its main uses. Once that was accomplished, there was ongoing technical progress that made PCs more powerful, easier to use, cheaper, and so on. This technical advance, plus the larger and larger network effects led to a steady increase in the number of users, a positive spiral.

The key to the speed of positive spirals for some personal applications software or support ITK is their low co-invention costs. These, in turn, tend to follow from the comparatively simple and non-transformative nature of these applications at the time of adoption. Their power comes later, with the invention of useful systems in which they are embedded.

While the creation of positive network effects in personal applications supporting IT can be a very powerful force, it can also create and sustain dominant positions for established vendors, especially when the network effects are linked to a proprietary product or standard. This will lead the incumbent dominant firm to seek to maintain control, perhaps by attempting to prevent widespread distribution of alternative new technologies. Users would like competitive experimentation outside, or partially outside, the existing network effects cluster. As long as they get to choose whether to switch, or partially switch, to a new technology, their interests are preserved. Vendors whose technology embodies a standard will discourage such experimentation and switching.

c) Implications of single market equilibrium

Knowledge creation creates economic opportunities as well. The beginning of the cycle in the last section was an advance in IT. This new knowledge creates an opportunity for applications. The center of that section was the invention of those applications and the creation of the required knowledge. In this section, we see that that knowledge, in turn, creates an economic opportunity for sellers of ITK. In the short run, it creates a market for them. In the longer run, new knowledge about applications of IT creates an opportunity for IT to change to accommodate those applications.

Many of the benefits of large markets and positive feedback between buyers' and sellers' inventions have been achieved in IT. The large scale of modern IT markets has led to the opportunity to put more and more resources into fundamental hardware improvements, and the emergence of more and more high quality software of both infrastructural and applications form. Not all of the advantages of large scale have been achieved (experimentation and competition are lacking in some contexts) a topic to which we return later.

At this stage, what we have seen is the complementarity between invention of new IT and co-invention in the IT-using sectors, and the structure of the social increasing returns to scale and resulting externalities those results. The limitation so far is that the discussion has been confined to a single round of IT invention and related co-invention in a single segment of IT application. We now turn to the question of sustained technical progress over time, and then to the degree to which existing market institutions have internalized all the externalities.

3) Widening Span of Application over Time

As IT has advanced, and as ITK has grown cheaper, the span of BIS IT application has broadened in several directions. I examine three in this section. This takes us out of the partial equilibrium framework of the last two sections toward general equilibrium and sustained economy wide mechanisms.

The first of the three broadenings stems from the re-use of fundamental IT inventions in a number of segments. A PC of today is as powerful as a mainframe of a generation ago, and far cheaper, smaller business units may now take advantage of BIS similar to those deployed only in large enterprises a generation ago. I take this up in Subsection a). The second broadening of the application of IT is moving the boundary between the specific and the general. What used to be co-invention, localized and thus expensive, becomes invention with the resulting exploitation of scale economies. I take this up in Subsection b). The third broadening, like the first, arises from cheaper and cheaper IT. In this case, however, both cheaper computing and cheaper communications move toward larger and more complex applications, notably those that cut across the boundaries of a firm.

a) Falling Fixed Costs of Computing

Early IT markets saw investments in fundamental technologies, notably in hardware and systems design. This meant that costs of IT costs of ITK have been falling, particularly that the costs of the smallest capable computing device have been falling. This led to the emergence of new markets. One way to understand these would be as a movement down a fixed demand curve –which suggests that diminishing returns set in quickly. Since the smaller size categories of computers have, historically, had vastly larger unit sales than the larger size categories, the diminishing returns would, on this story, have been temporarily offset by the shape of the demand curve. Demand must have been elastic over the relevant range to explain the relative size of the PC and Mainframe market segments, for example.

Examples of mere diminishing returns do exist. On the data processing side, we have seen the extension from core applications that undertake the most valuable dp activities in the firm to far more trivial ones, keeping track of the football pool in the mailroom. In terms of computers literally "computing", diffusion of cheaper and cheaper floating point arithmetic from very critical problem in physics (with urgent military applications) to drawing bad guys on the screen (a highly floating point-intensive task) for computer games. One consequence of far cheaper ITK has been diffusion to lower value uses. It would be quite odd if it did not have this consequence.

If that were all the lower costs achieved, one would conclude that diminishing returns were setting in. But the smaller size categories of computers have had their own feedback loops, with new classes of applications being invented and widely used. Most of the elements of IT have increasing returns to scale, from the first-copy-costs of a piece of applications software to the invention costs of new kinds of storage devices. As a result, the large market size of the newer categories of computing, such as the PC, is both fed by and feeds back to considerable supply of new IT of a number of kinds. A fundamental nonlinearity has been at work in the system. IT cheap enough to be used in a mass market draws mass-market complements and improvements. This drives the size of the mass market in IT ever larger, creating an economic return to invention of new IT specific to that market

Diminishing returns do not quickly set in, despite the widening span of application, because the positive feedback loop is renewed when the span of application widens. Of course, sooner or later this particular dynamic is going to run out of economic activity. Perhaps handheld computers, or smart cell phones, won't generate nearly as much of a positive feedback loop as did PCs. Through now, however, the re-use of fundamental IT developments in new, cheaper, categories has renewed the growth cycle.

b) Converting Activities from Specific to General

When many inventors are working on very similar problems, the possibility of exploiting scale economies in invention arises. Often, the solution will arise in the form of vertically disintegrating the common part of the invention and creating a new industry or discipline that focuses on only that, i.e., a new "general specialty." The result increases the degree to which scale economies are exploited at some cost of lost specialization. This mechanism can powerfully change the pace of technical change.[14] Since much of co-invention of IT applications appears to be the same inventive activity undertaken over and over again, there is a long history of efforts to split parts of it out into new general specialties.

Steady progress has been made in this direction, but there appear to be limits. Within the "technical" parts of co-invention, the history shows tremendous successes. As Friedman and Cornford (1989) document, the market IT sector has made tremendous steps in helping IT-using firms with the technical parts of co-invention. IT products, of the type generally called "infrastructure" and "tools" by using firms, have taken over many of the tasks of managing hardware and managing the writing of applications software. Perhaps the most important of these developments were the invention and improvement of (a) applications-development environments and related programming tools and (b) database management systems, but there are many other examples. The overall effect of this technical progress has been to replace duplicative effort by a large number of technical people in IT-using companies with tools that undertake the effort for them. MIS departments have become, in the narrowest technical sense, considerably more productive as a result of this progress.

This extension of the general into taking over from the specific has, however, met important limitations. The business part of co-invention has resisted movement into a more scale-economies exploiting mode, remaining stubbornly local. A long series of efforts to move all of business co-invention into general specialist firms has largely failed. The most recent two failures are both quite spectacular, the "business process reengineering" movement and the horizontal .com movement. The horizontal .com movement hoped that all IT-using firms and industries would move into new applications areas in electronic commerce using universal ("horizontal") applications software and services. The slightly earlier BPR movement hoped that the tasks of deciding on the

[14] See Bresnahan and Gambardella (1999) for theory in the spirit of Stigler, and Arora, Landau, and Rosenberg (1998), for how the creation of the chemical engineering discipline was centered around making the analysis of processes routine and regular, so that each new plant-design or process-design task could build on prior knowledge.

business logic of new applications could be turned into an "engineering" activity that would have a new specialty, like, say, C++ programmers, that one could simply hire to get the job done. While some good came from both efforts, neither of them – nor the many other similar efforts – led to widespread exploitation of scale economies in the business part of co-invention, in contrast to the successes in the technical part.

c) Making more and more of the production process ordered and routinized

As IT has grown more capable in both hardware and software, and as ITK has grown cheaper, the possibility of larger and more complex applications has grown. The boundary between IT-based production in white-collar bureaucracies and other systems of production has moved steadily as a result.

The familiar example of the airline reservation system can serve to illustrate this point. It was originally invented to count something important, unsold seats on airplanes. That is a fairly common start for regularization and routinization, there is something that is obviously countable and valuable to count. Once something is counted, it can support a simple operational behavior, such as do not sell any seat that has already been sold 1.48 times or more. Of course, one would like that operational sales decision to be made at a wide number of distinct sales points, including one's own sales people, distribution agent (travel agents) and large customers. That calls for some invention, i.e., IT that can deliver the number of unsold seats to all those distinct people, to some technical co-invention, i.e., building a system that reliably delivers that information to all those different people at all those different sales locations and decides which of them sold the seat firs, and some co-invention, i.e., deciding what exactly is a large customer permitted to learn from the database about capacity and inventory. In this simple operational narrative, I have swept together decades of invention (database management systems, networks, etc.) and co-invention under the long run goal of an integrated operational system. Operations thus become more tightly controlled.

That database of unsold seats becomes a database of ticket sales and capacity utilization, which permits building analytical models management can use to make valuable decisions. What kinds of routes have systematically empty seats? That can guide allocation of aircraft to routes, and allocation of sales efforts to fill planes. What kinds of sales lead to customers actually showing up to fly, and what kind to no-shows? That can guide the joint determination of the overbooking rate (pegged at 1.48 above) and capacity. More and more complex analytical models can be built as more and more of the production process is measured and monitored systematically. Management functions thereby become more analytical and systematic. Managers feed back to the tightly controlled operational system what they would like it to accomplish.

The database of ticket sales also becomes a marketing tool for building complex customer relationships. Volume discounts and other mechanisms for treating different classes of customers differently with regard to specific transactions (e.g., frequent flier programs) come into use. Learning from customer behavior becomes both easier (big databases) and more IT-intensive, and feeding back to marketing the information learned about customer behavior in order to build more complex relationships itself calls for more IT systems. Customers, especially large ones, want to manage their side of that complex relationship, leading to more opportunities for cross-firm communication. Since much of the production process in service industries is shared between buyer and seller, the marketing function may include customer-management functions as well, which here would be connected at first to measuring no-shows, later, in attempting to minimize the impact of no-shows on costs by incentives or planning. The entire buying selling relationship can end up being routinized and regularized, with complex contracts between buyer and seller emerging in explicit (software based) form.

At each stage, routinization leaves behind not only databases full of facts but also the structures that systematically capture economic data and put it into databases and the other structures that increasingly rely on precise calculation from those data to support or perform decisions. The overall production process of the firm – and to some degree, beginning later, perhaps twenty years ago, of its customers, suppliers, and distributors, grows more routinized and regularized.

Of course, a different version of this history, adjusted for the different information economics of other kinds of firms and industries and for the different kinds of commercial relations between buyers, sellers, and intermediaries, has gone forward in other contexts. In many industries, however, the broad general themes of this section are playing out. The important general themes are those of operational support, analytical decisionmaking, internal reorganization, discrimination across customer or vendor classes, and the emergence of complex contracts based on multiple observables. At every stage, these involve not only invention but also co-invention.

d) Summary on the Extensive Margin

What each of the sections above shows is how the historical pattern of IT use and IT investment does not have the character of falling costs taking us down a fixed demand curve. Instead of taking the economy toward diminishing returns in the first instance, IT as a general capability has found new extensive margins at which to work. Each extensive margin has set off its own cycle(s) of positive feedback. As a result, the demand curve for ITK shifts out and diminishing returns do not set in.

The GPT logic of IT applies not only within each kind of application, as we saw in Section2), but across kinds of applications, building new kinds of IT partly on the foundation of the old kinds. This re-use or nonrivalrous nature of IT knowledge arises from the GPT nature of IT in the large (across very different kinds of applications) just as in the small (across different applications of the same specific technologies.)

Development along each new margin also has inherent limitations. Ultimately, one would expect this kind of expansion to hit decreasing returns as well. For now, however, the expansion has continued with great success.

4) Complementarity over Time

For continued costly invention and co-invention to be profitable, some force must keep the return to the use of ITK from falling over time. So far, we have looked at the general purpose technology nature of IT, one powerful force keeping the return high. A second dynamic force arises because, just as invention of IT spurs co-invention, ideas from co-invention spur new invention. This leads to ongoing improvements.[15]

[15] The outcome is much like that of the “quality ladder” models in the growth literature, e.g., Grossman and Helpman (1994). In those models, as here, early steps in the invention process form part of

IT improvements over time build on the basis of earlier inventions in IT. In that sense, there is a fundamental complementarity between inventions in IT at one time and inventions at another time. That narrow remark, however, applies to all cumulative bodies of knowledge, i.e., to all science-based technologies. The particular structure of invention and co-invention leads to a sustained pattern of inventive opportunity for reasons that go beyond that universal force. Invention and co-invention are not only static complements, but their duality creates dynamic complementarities, as inventions in each domain raise the economic return to inventions in the other and provide intellectual stimulus to one another.

a) Inventions in Two Domains: Complementarity over Time

The centerpiece of the argument about knowledge creation in Section 1) and 3)c), above, was that a cycle of invention and co-invention create two distinct bodies of knowledge, one about IT, another about the applications of IT.

This feedback system is inherently dynamic because of the way that knowledge about those two areas works. It is very difficult to anticipate, even with a great deal of communication, how the co-development of IT and its applications should go. Typically, sellers of IT announce new technical capabilities, and these either do or do not find a market. Sellers learn from the co-invention that actually occurs what inventors of applications really want. At each new stage of learning what can be accomplished with IT-based production in the using industries, the state of broad IT knowledge is higher. Thus, vendor firms attack problems at a later stage with both better "technology" in the sense of being able to make smaller and faster ITK devices and better knowledge of the market for new IT inventions.

Another force, growing more important with time, is relationships of complementarity between existing installed ITK and existing BIS, on the one hand, and new advances in IT and the possibility of new BIS, on the other. Some of this is simple cumulativeness. After the installation of BIS that permit operational control, analytical opportunities arise later. We saw examples of this in the airline discussion, above. Another example of this would be that, in one generation, accounts receivable software permits simple operational control of billing, getting letters out to dun late payers. In another generation, the existence of the operational database in the AR software permits analysis of what kinds of customers pay on time, and what kinds of incentives lead customers to pay on time, valuable invention that builds upon the earlier valuable invention. In yet another generation, those AR systems and several other operational databases and their analytical overlays are combined into a unified customer relationship management system. At each stage, the cumulativeness arises from the character of the investments made at the earlier stages: more information held explicitly in BIS lowers the costs of new BIS. This is a force that plays out at the individual IT-using-firm level, and which creates market opportunities for invention of new IT and for co-invention of new BIS at later stages. Because of the combination of cumulativeness and new

---

the foundation for later steps. In the models, however, economic agents know from the beginning which way is up the ladder, but the ladder is not instantly scaled because the incentive scheme rewards single steps as much as multiple steps. My analysis here is driven by agents discovering which of a number of apparently attractive rungs leads upward.

opportunity, there is ongoing positive feedback between past investments and new investment opportunity for inventors of IT and sellers of ITK.

This feedback between invention and co-invention leads to a repeated cycle of improvements in two stocks of knowledge. This process can carry buyers and sellers a long way with improvements that are entirely complementary over time. If the standards permitting different hardware and software components, including those which are written by using firms, change slowly, then the process of repeated invention and co-invention leads to cumulative change. This path will be encouraged by existing dominant firms selling IT, often with considerable success. Of course, change of that type sometimes stays on the incremental improvement path for too long a period, and then radical breaks with the past are called for. These tend to be sponsored by entrant sellers of IT, not from completely outside the IT market, but from a different part of the market for IT.

b) Dynamic complementarity between different IT

Not all of the relationships of complementarity happen within such narrow confines. New forms of IT are sometimes complementary to old forms. Consider the commercialization of the Internet. The Internet had been use primarily in universities and governments. A series of developments, most notably the invention of the browser, brought the Internet into widespread commercial use, where it appeared as a genuinely new technology. The complementarities between this new technology and technologies already in use are striking. "The Internet has sold a lot of PCs," said Microsoft's CFO. Absolutely, and it has sold a lot of large computers, and a lot of software, and telecommunications services and equipment as well. So here we see new technical progress raising the MVP of capital. The complementarity also flows the other way. The prior success of the PC meant that many people could access the Internet cheaply. The prior success of BIS meant that many business datasets could be re-used or cheaply created to show, for example, the price of the item you would like to buy on a web page. New inventions of valuable applications such as email, instant messaging, or the many forms of electronic commerce are complementary to existing IT (both as technology and as capital in place.)

The complementarity between old inventions and new inventions, and the related complementarity between old investments in ITK and new ones, is heightened by the secularly growing scope of BIS. Where BIS originally applied to narrow, specific tasks, such as collecting bills or scheduling a machine, they have grown to encompass or at least affect much of the production process (in service industries) or much of the management and control of the production process (in all industries.) The production process includes buying inputs and selling outputs – indeed, in modern economies, far, far, more workers are engaged in those activities than are engaged in, for example, manufacturing – so BIS have grown to include many activities in which a firm interacts with its customers and suppliers, as we saw in the airline example. The growth of the BIS to itself cross the boundary of the firm was the next step, a path down which the economy began to go some 20 years ago. That is a demand-pull agenda drawn by the value of the applications that might be generated.

Demand pull from co-inventors creates a series of new economic opportunities for inventors. Co-inventors wish to design new applications that take advantage of the assets in place (in existing, "legacy" applications) and which can be accessed over the now-near-universal Internet from the now-near-ubiquitous PC. In the first instance, this shifts out the demand curve for existing forms of ITK, as the new systems will need more capacity, newer equipment, and so on. More powerfully, this creates a demand for new forms of IT invention and ITK which would make it easier to design and implement those new applications systems running on many different computers over a network. The trend toward networking in computer applications brings advances in telecommunications into the positive feedback loop. The economic return to making bandwidth cheaper (or more portable) rises along with the economic return to making computing cheaper (and more capable.)

c) Convergence of Formerly Separate kinds of IT

Convergence is another powerful force preventing the onset of diminishing returns, especially for complex applications which have large potential impact. The existence of distinct clusters of IT technology gives new, complex applications a menu of supporting technologies, and creates the opportunity for new technical progress linking those technologies together. Further, past investments in ITK in the distinct clusters let new applications that cut across the boundaries of the firm assume that there will be ITK in the other firms. This creates a very powerful complementarity between the inventions of the past and the inventions of the future. Rather than crowding out the inventions that will support interorganizational computing in the future, the inventions that supported organizational and personal computing in the past [labels] raise their value.

The long run dynamic which gives IT the power to make contributions over a long time span has two very distinct elements, one of specialization, the other of convergence. Very general scientific and technological knowledge underlies the basic engineering of computer hardware and software and telecommunications and networking equipment and systems. That general knowledge is specialized into distinct clusters or nodes of technical progress. Within each cluster there are powerful external effects and positive feedback between the (specialized version of the) general purpose technology and applications. Clusters such as telephone systems, mainframe computers, and PCs have historically grown following distinct positive feedback paths. Since there is sharing of basic technologies across these clusters, the founding of new clusters is a powerful force preventing the onset of diminishing returns

d) Complementarity over all

The point of these particular examples is that diminishing returns in the process of accumulating knowledge in specific kinds of applications and IT need not become diminishing returns to the aggregate use of IT. Widespread success in one applications area accumulates assets which will be the foundation for new applications areas. The most important examples in the present have to do with extension of applications across the boundaries of the firm. Will that economic and technical opportunity be quickly exhausted? So much of employment in the rich countries is in white collar bureaucracies that buy or sell, it seems highly likely that automation followed by process improvement will have high payoff. New ways to use old IT and combine it with new IT, as yet far from fully exploited, suggested the ongoing process of positive feedback and accumulation will continue.

5) Incentives and Goal Conflicts

The general purpose technology structure of IT means that there are a number of important externalities among co-inventors and between co-inventors and inventors. If all of these were internalized by contract, the economy would undertake the socially optimum rate of invention and co-invention. The modern economy is far from that standard, and for important forms of co-invention and invention the social return to innovation lies above the private one. I discussed some of the information gaps which make literal contracting a difficult task above. In this section, I take up the question of why and how existing structures for the organization of invention and co-invention do not always achieve the optimum. The existing structures that I discuss are **(1)** the seller of general purpose ITK, from computers, to telecommunications switches, to operating systems to databases and tools, **(2)** the seller of packaged applications software, **(3)** the seller of custom or semi-custom applications software, and **(4)** the using firm. Obviously, considerable scale economies are achieved by having firms sell ITK to many using firms. What incentive problems remain?

I have already mentioned the re-use externality for co-invention. When the tasks of an application are reasonably well defined and common across many users, this externality can be partly internalized by moving technical progress in applications from users **(4)** to an independent applications vendor **(2) or (3)**. There are some limits, as we have seen above, because many applications require very extensive localization to work in particular using firms. Thus the widespread movement of applications development from the comparatively expensive **(4)** or **(3)** to **(2)** has been limited. The most direct accommodations are applications software – simply taking the ideas from applications first invented in a clumsy or one-off way and writing a piece of software. This permits widespread use of the idea in an economical form. When firms or industries cannot easily use the same application software (because they are different in their organizational structures, relationships to vendors and customers, or products) other forms arise. Semi-custom software is used for some problems, or software written by consultants over and over for somewhat different using firms, based on the same knowledge. For other problems, market IT can be altered to make writing certain kinds of applications in using firms easier, lowering the costs of applications variety.

There is also an important incentive problem among co-inventors. Suppose firm A undertakes an important new application of IT that it thinks will give it lower costs or better products than its competitors. Firm A will not be eager to have the key ideas in its new application embedded in a piece of application software and sold to competitor firm B. Firm A will not be eager to have its consultants or other semi-custom software vendors finish the application and then go over to firm B for a new engagement. Rightly, firm A will think that much of the co-invented part of the application is its own invention and that to have it spill out to competitors is a problem. Accordingly, efforts to fully achieve the broad re-use of co-invented ideas are likely to face difficulties.[16] Only those things firms do not find strategic are quickly and widely re-used by the comparatively

[16] Ross Perot, inventor of applications software resold to many users, sensibly chose as an early body of customers the local monopoly health insurance firms in the US in the 1960s. The semi-custom software industry he founded has had its greatest successes selling to PTTs, governments, and the like ever since.

efficient spill out mechanisms (3), instead of clumsy leakage among (4) by, for example, movements of employees to competitors (labor market sales of trade secrets.)

The other important externality associated with a general purpose technology lies between the sellers of general products and co-inventors. In the case of IT, this externality falls between inventors of IT products and co-inventors. Each group would like to see the other experiment with a wide variety of approaches, as each group would like to have successful (co)invention by the other group pushing out the demand curve for its products. Inventors can typically expect this behavior from co-inventors, as there are many of the latter.

There is a problem achieving adequate experimentation on the invention side in some circumstances. Certainly when an established proprietary standard is at the center of a great deal of applications development, there is a goal conflict about competitive experimentation between sellers and buyers. This applies to either the case of adjustment cost co-invention or the case of network effects.

The famous example from the past is experimentation with replacements for the IBM mainframe platform. Further inventions of compatible extensions to the platform by IBM were one of the highest private rate of return investments in the economy, and IBM prosecuted them with vigor – and with care to make sure that customers always stayed with IBM. Competitive alternatives were one of the lowest rate of return investments in the economy, and for a generation, as customers grew more and more to wish that IBM were a bit more experimental, competitive investment and experimentation was largely lacking. Indeed, the technologies which ultimately became competitors to the IBM mainframe standard were not invented as competitors, but underwent a long "indirect entry" path, at first paid for by other uses. Only long after customers would have wished it were there serious alternatives.

Much the same story is being replayed in the Microsoft era. That firm has every incentive to have such important technical developments as the commercialization of the Internet introduced to its customers in a way which permits it to manage customers' technology choices, so that customers remain loyal to Microsoft platforms and technologies. Customers would prefer, given that they would have the choice to stay with Microsoft products, considerably more experimentation in response to the convergence of PC and Internet. Those experiments have lately been blocked by Microsoft's interference with the widespread distribution of products and technologies it finds competitively threatening.

As one might expect when the social return to innovation could be appropriated by any of a number of different firms, including existing sellers of IT, new sellers of IT, sellers of different types of IT, and users of IT, contracting to internalize all the externalities is incomplete. Accordingly, social returns to invention in the aggregate remain above private returns.

6) Conclusions

Much of the mechanism of IT's contribution to economic welfare follows from the social increasing returns to scale rooted in complementarity between very different kinds of innovative activity. Within the many particular markets for IT, the complementarities take the form of positive feedback between invention of new IT and co-invention of applications. Invention and co-invention are both areas with substantial innovative possibilities, so the cycle of positive feedback can be long-lived, as important inventions in one domain lead to new inventive opportunities in the other. Existing structures fail to internalize all the externalities among inventors and co-inventors, first because of limitations what they can foresee about one another, second because existing IT dominant firms seek to control and manage the arrival of new technology, sometimes succeeding in reducing the amount of experimentation and user choice far below the feasible competitive level. These mechanisms have delayed the onset of diminishing returns for decades. Socially beneficial routinization and regularization of white collar bureaucracies is now moving into its second half century without signs of technological exhaustion.

The range of applications of IT has been expanding as well, another force that tends to offset diminishing returns. Cheaper ITK has meant (of course!) that some of the new applications are of lower value than the older ones. Yet it has also set off new positive feedback cycles which have pushed out the demand for these new types of IT and created new sources of value. The personal computer, for example, has not appeared merely as a cheaper form of mainframe, but instead has served a number of new functions, most recently involving communication and information retrieval, which increase its aggregate contribution. Here, the social return once again exceeds the private because the founding of new categories occurs by spill out from old categories.

A final complementarity arises at an even wider scale. Over time, a number of specific IT technologies have come into widespread use. New applications today, for example, can assume a wide installed base of PCs and of large organizational BIS. Thanks to more recent developments, new applications today can assume widespread connection to the Internet. Networked applications that take advantage of those assumptions can be developed more cheaply and quickly. The technical inventions of the past become complements to the technical inventions of the future through the expanding scope of applications. By this mechanism, the Internet has accelerated the 20-year-old trend to applications crossing the boundary of the firm.

All of these forces will eventually be exhausted, as the contribution of any particular technology, even one as powerful and levered as IT, must ultimately hit diminishing returns. They show no near term prospect of being exhausted, as white collar bureaucracies continue to employ many while not yet doing quite what their customers or managers would like.

7) Bibliography

Ashish Arora, Ralph Landau, Nathan Rosenberg, eds. (1998) Chemicals and long-term economic growth New York: Wiley.

Barras, R., "Interactive Innovation in Financial and Business Services: The Vanguard of the Service Revolution," Research Policy, (1990), No.19.

Bresnahan, Timothy, "The Changing Structure of Innovation in the Computer Industry," in D. Mowery et al., editors, U.S. Industry in 2000: Studies in Competitive Performance, (National Academy Press, Washington, 2000).

Bresnahan and Gambardella

Bresnahan, Timothy, and Shane Greenstein, "Technical progress and co-invention in computing and in the uses of computers," Brookings Papers on Economic Activity, Micro, (1997) pp.1_83.

Bresnahan, Timothy and Shane Greenstein "Technological Competition and the Structure of the Computer Industry" in Journal of Industrial Economics, March 1999, v. 47 n. 1, pp. 1-40.

Bresnahan, Timothy, and Garth Saloner 'Large Firms' Demand for Computer Products and Services: Competing Market Models, Inertia, and Enabling Strategic Change," in Colliding Worlds: The Merging of Computers, Telecommunications, and Consumer Electronics, David B. Yoffie (ed.), Harvard University Press, 1996.

Bresnahan, Timothy and Manuel Trajtenberg, "General Purpose Technologies: 'Engines of Growth'?," Journal of Econometrics, Special Issue, January 1995, v 65, n 1, pp. 83-108.

Brynjolfsson, Erik and Lorin Hitt, "Paradox Lost? Firm Level Evidence on the Returns to Information Systems Spending," Management Science, (1996), Vol. 42, No.4, pp 541_558.

Brynjolfsson, Erik and Shinkyu Yang, "The Intangible Costs and Benefits of Computer Investments: Evidence from Financial Markets," Proceedings of the International Conference on Information Systems, (1997).

Brynjolfsson, Erik, Timothy Bresnahan, Lorin Hitt (2000) "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-level Evidence" forthcoming in Quarterly Journal of Economics.

Davenport, Thomas, and James Short, "The New Industrial Engineering: Information Technology and Business Process Redesign," Sloan Management Review (1990) 31.4: 11_27.

David, Paul. "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox," American Economic Review, (1990), Vol. 80, No.2, pp. 355_361.

Friedman, A. & D. Cornford, (1989). Computer Systems Development: History, Organization and Implementation. Wiley, Chichester.

Garicano, Luis and Steven Kaplan, "The Effects of Business-to-Business E-Commerce on Transaction Costs." NBER Working Paper W8017, November 2000.

Gordon, Robert J. "Does the 'new economy' measure up to the great inventions of the past?" NBER Working Paper, 7833, August 2000.

Greif, Avner, Paul Milgrom and Barry Weingast, “Coordination, Commitment, and Enforcement: The Case of the Merchant Guild.” Journal of Political Economy, (1994), Vol. 102, No.4 pp. 745-76.

Grossman, Gene M., and Elhanan Helpman (1994): "Endogenous innovation in the theory of growth." Journal of Economic Perspectives, 8:23-44.

Helpman, Elhanan and Manuel Trajtenberg, “A Time to Sow and a Time to Reap: Growth Based on General Purpose Technologies,” in Elhanan Helpman, ed., General Purpose Technologies and Economic Growth, MIT Press, 1998.

Hubbard, Thomas N. “Why are Process Monitoring Technologies Valuable? The Use of On-Board Information Technology in the Trucking Industry.” NBER Working Paper W6482, (1998).

Ito, Harumi, “Essays on Investment Adjustment Costs.”(Ph.D. Dissertation, Stanford University, 1996).

Katz, Michael and Carl Shapiro, (1994), “Systems Competition and Network Effects,” Journal of Economic Perspectives, vol. 8, no.2, pp. 93_115.

Litan, Robert and Alice M. Rivlin, Beyond the Dot.Coms The Economic Promise of the Internet. Brookings Institution Press, 2001.

Malone, Thomas and Jack Rockart, “Computers, Networks and the Corporation,” Scientific American, (1991), Vol. 265(3) pp. 128_136.

Milgrom, Paul and John Roberts, Economics, Organization and Management. Englewood Cliffs, N.J.: Prentice Hall (1992).

Paul M. Romer “Endogenous Technological Change,” Journal of Political Economy, Vol. 98, No. 5,. (Oct., 1990),

Saloner, Garth, and W. Edward Steinmueller "Demand for Computer Products and Services by Large European Organizations,” June 1995, Stanford GSB Research Paper Series, No. 1370.

Scott Morton, Michael, The Corporation of the 1990s: Information Technology and Organizational Transformation, (Oxford University Press, 1991).

Zuboff, Shoshana, In the Age of the Smart Machine: The Future of Work and Power. New York: Basic Books, (1988).