Industrial and Corporate Change, 2019, Vol. 28, No. 2, 331–349
doi: 10.1093/icc/dty076
Advance Access Publication Date: 8 March 2019
Original article

Technological change in ICT in light of ideas first learned about the machine tool industry

Timothy F. Bresnahan*

Department of Economics, Stanford University, Stanford, CA 94305, USA. e-mail: tbres@stanford.edu

*Main author for correspondence.

Abstract

Rosenberg (1963, *J. Econ. Hist.*, 23(4), 414–443) analyzed innovation in the 19th and early 20th century machine tool (MT) industry and its industrializing customers. This paper re-applies his analysis to the invention and application of information and communications technologies (ICTs) in the 20th and 21st centuries. Rosenberg's ideas of technological convergence and of collective learning that leads to a pool of knowledge in a competitive upstream industry are emphasized. In some ICT market segments, the primary customer is a technologist; for these segments, Rosenberg's analysis works quite well. The core reason appears to be that there, as Rosenberg said of MTs, the analysis can go forward on a "purely technological level." In the commercially oriented ICT market segments where most value has been created, this assumption fails. Difficulties at the boundary between purely technological innovation and commercial innovations change the analysis considerably. The locus of accumulating knowledge, the path of technological change through the economy, the structure of the industry supplying the general components and the degree of sharing of technical advances across firms and industries, all change in an easily explained way. The idea of the boundary between two complementary bodies of knowledge may be important. The idea that commercial innovation can be unlike technical invention may be important.

JEL classification: D20, N82, O31

1. Introduction

"Technological Change in the Machine Tool Industry, 1840–1910," Rosenberg's (1963) "modest attempt" to turn Economics' attention to the causes of technological innovation, sought to examine and explain technological change in machine tools (MTs), a sector which played a central role in the industrialization process in the United States. Both deeply connected to the particulars of the economic history of that industry in the late 19th century and full of general economic insights, it is an example to us all. This paper looks at information and communications technologies (ICTs) and their applications. I first follow the Rosenberg framework closely, which turns out to be remarkably useful in explaining innovation the segments of ICT that are most like MTs. Other segments, however, have important differences, which illuminate both what is general about the Rosenberg framework and the specific elements of the history of MT industry that led to it.

I focus on two of the many very interesting elements in the Rosenberg analysis. He studied innovation both in MTs and in the many manufacturing industries that used them. He identified a positive feedback loop across those using industries. New or improved MTs invented for use in one industry would later be used in others. The same process, done by the same tool, would arise in many different industries. This re-use led to a situation in which many steps in production processes in using industries were the same, that is, led to "technological convergence." While he noted the resulting externality, Rosenberg focused more on using the feedback loop analysis to explain the innovation than to assess its efficiency.[1] I follow him in that goal and that limitation.

Rosenberg also analyzed the collective learning associated with all the MT and MT-using innovations. Learning occurred both about MTs, and about their application. To a considerable extent, the learned knowledge was stored within the MT industry, and then diffused across the economy. I agree very much with Strassmann (1963), who in his contemporary remarks on Rosenberg's paper, noted that the most important element of the analysis concerned learning. While there has been a great deal of analysis of learning and of the accumulation of knowledge,[2] the contrast between MT-era and ICT-era *learning* raises some new and important points.

I will focus on two ideas that arise from that contrast. (i) The existence of a potential general purpose technology (GPT) and of potential demand in a number of potential applications sectors is not the only determinant of GPT innovation. Instead, *markets* for the GPT play a role, as do *economic institutions* to coordinate invention in the GPT itself and in applications sectors. These markets and institutions will function well if they reflect the economics of the process of learning and of the accumulation of knowledge. The contrast also shows us that (ii) the process of learning and of the accumulation of knowledge were different in the ICT era than in the MT era. This has little to do with the content of the GPTs themselves. Indeed, in Section 4 of this paper, we shall see that there are important parts of the ICT-using modern economy, primarily in scientific and engineering computing, which had learning and accumulation very much like the MT industry in the era studied by Rosenberg. In contrast, some ICT segments and their applications, particularly the high-value inventions in enterprise and commerce, have had a very different kind of learning and accumulation process. This different process has been reflected in a different set of economic institutions for coordination and a different structure of ICT product markets in those segments.

Two learning processes are alike, (i) in MT and (ii) in scientific and engineering ICT, while a third is different (iii) in ICT enterprise and commercial invention. While the learning processes, and many other aspects, differ along a number of dimensions, I shall explain many of these differences using the implications of one feature. Enterprise and commercial ICT *applications* invention is not narrowly technical, but instead involves changes to the structure of the firm, to the structure of work and incentives, and possibly to the way customers are served. As a result, the most valuable ICT applications both require technical progress in ICT—one sphere of knowledge—and invention of new ways of doing business in the using firm—another sphere of knowledge. This leads to difficult problems of *ex ante* seeing the overlap between technical opportunity and commercial value. I call the problem of "visibility."

Further, not all applications of ICT in commerce are equally difficult to foresee, so that the diffusion of ICT uses has, to some degree, followed visibility rather than value. Another implication of the same distinction is that *ex ante* communication that might coordinate technical progress in ICT and technical progress in applications of ICT has proved difficult in the commercial and enterprise segments. The same problem of two spheres of knowledge returns *ex post* when it comes to understanding the accumulation of knowledge—Rosenberg's "learning." The two spheres of knowledge problem led economically important differences in the accumulation process between commercial ICT and the other two areas studied here. It has also led to differences in the economic relationship between providers of general purpose vendors and applications innovators in the commercial sphere.

We shall come back to those analytical distinctions only after doing some historical work. The plan of the paper is to build a stylized version of the Rosenberg (1963) analysis, then apply it to the parts of ICT where it succeeds, and only then come back to the segments of ICT application—the high value commercial and enterprise ones—in which the differences arise. To be clear, those differences arises because (i) the two complementary inventions, those in the GPT and those in applications, arise in very different spheres of knowledge, and (ii) one of the spheres of knowledge, commercial applications of ICT, involves inventions whose outcomes are difficult to foresee *ex ante*. Whether this is general or not I leave to future research; it is a core feature of the most important GPT of our time.

---

[1] A very successful literature has taken up the analysis of positive feedback loops generally. See Farrell and Klemperer (2007). It takes up efficiency questions of two broad forms—is there enough coordination in the feedback loop? Does coordination lead to later inertia?

[2] Foray (2004) has not only a point of view but also a useful overview of the field. A large literature, including Nelson and Winter (1982) and von Hippel (1986) has taken up the accumulation of knowledge through invention.

A large literature compares ICT to earlier eras in the history of technology.[3] Much of this work focuses on outcomes, trying to explain aggregate measured productivity growth. I instead follow Rosenberg, seeking to understand learning and the positive feedback loop as part of the explanation of innovation. Only on that basis can one attempt a forecast of the next wave of ICT usage, which I discuss briefly at the end.

2. Elements of the Rosenberg analysis

For analytical purposes, I divide the Rosenberg explanation into two parts, micro and macro. By the micro, I mean the positive feedback loop and the learning, which could occur at a broad or narrow scope. By the macro, I mean that the scope taken up by the Rosenberg analysis is part of an entire economy's industrialization transition.

2.1 Micro

The micro part of Rosenberg's analysis explains "technological convergence" and the accumulation of knowledge, a "pool" of inventions in an upstream (shared) industry. At its core, the micro story is one of a positive feedback loop between product inventions each of which calls for new process inventions and therefore leads to MT inventions, together with accumulation of knowledge about both the process step in the using industry and the tool. The loop and the learning lead to a number of new inventions and the accumulation of a body of knowledge. The positive feedback loop creates convergence, and a set of critical learning mechanisms along the invention path creates the body of knowledge. The positive feedback loop and the learning process interact in a way we shall explore at some length. There are two essential elements to the Rosenberg analysis of learning. First, learning is an accumulation process. More is known after learning than before. Second in addition, "learning" opens doors to new invention. That which has been learned directs each step of the feedback loop.

I'll use the picture of the "micro" cycle in Figure 1 to lay out one step in the positive feedback loop with the associated invention and learning.

How does this cycle work? A new industry is created—think locomotives, or mass market firearms, or sewing machines, or bicycles. The existence of the new industry makes visible some novel production problems, some of the solutions fall in a broad technology.

1. *Novel products have novel production requirements*: "The machine tool industry, then, originated out of a response to the machinery requirements of a succession of particular industries . . ."[4] This is a familiar step from industry life cycle analysis. Founding a new industry (or inventing an important new product within one) makes known and specific its "machinery requirements." Sometimes, this will require economic experiments.
2. Creation of a new product creates *novel production problems* and turning something into a specific problem to be solved can be very helpful in getting the engineers and technicians on the right track, because they know where they are going. Before the age of the emergence of the MT industry, Rosenberg points out, "For many years, the most intractable problems associated with the introduction of techniques of 'machinofacture' lay in the inability to produce machines which would perform according to the special and exacting requirements and specifications

---

[3] A smaller literature compares earlier rounds of "technological revolution" to the present, in which ICT is viewed through the lens of history. David and Wright (2005) usefully compare the period of factory electrification (somewhat later than the one used here) to the present. Their argument is much more about the process by which a general purpose technology diffuses, potentially leading to a later stage in which productivity rises. In this paper, following Rosenberg, I am much more interested in the learning and in the explanation of the technological progress. Freeman and Louçã (2001) compare a series of long waves, including the present one, with a methodological goal of bringing economic history into the discussion of aggregate technical progress. Atkeson and Kehoe (2007) take a macroeconomic model approach, in which technical change is endogenous. In this sense, the goal of Rosenberg to bring technological change to the forefront has succeeded.

[4] In this section, I am going to use quotation marks exclusively for quotations from Rosenberg (1963) since I am sketching his argument.

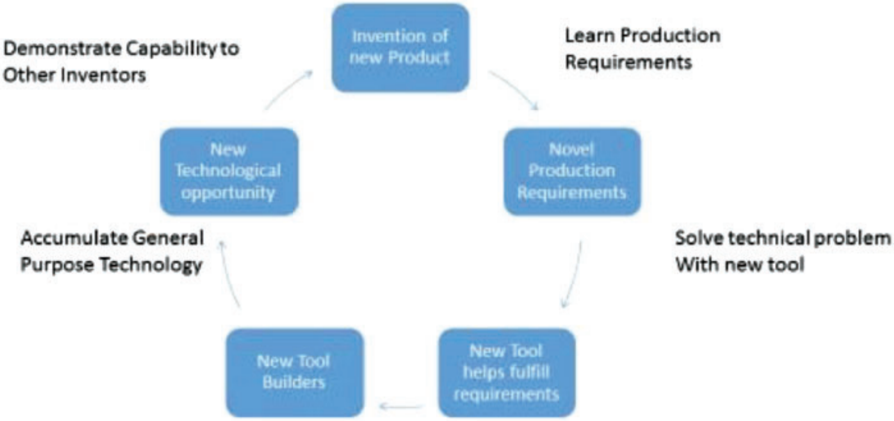

Figure 1. Schematic of Rosenberg Analysis.

of the machine users." Particular metal-using industries were "continually involved in 'setting the stage' for particular problems."[5]

3. The solution to a particular problem in a particular industry or project *elicit solutions* in the form of new technical invention.
4. "Demonstration Effect" and the founding of new markets and industries: Creation of a new tool demonstrates its capabilities, leading potentially to new uses. The tool itself is a new piece of technical knowledge. Rosenberg notes that a new tool would quickly find application outside the industry of its first invention "because the technical skills acquired in the industry of origin had direct application to production problems in other industries."
   a. *Vertical industry structure:* Once the new tool has potential users in multiple industries, it can be efficiently supplied by an upstream industry. If there are a large number of related tools, then they can be supplied by a number of separate stand-alone firms, each of which might have a narrow product range. In their totality, they are the information store.
   b. Happily, the industry that made MTs also needed MTs, using many of the same process steps as using industries. This created a further "internal" loop that reinforced the impact of the wide positive feedback loop.
5. Each trip around the cycle in Figure 1 created a technological opportunity for "other industries," typically ones that (even if not yet founded) shared a common or closely related production processes based on machinery. This had both narrow and broad implications: a very specific common problem (making parts for "reducing friction" was very important to both bicycles and autos, and to some degree to other industries), or a wide cluster of solutions: "(i) new skills and techniques were developed or perfected here in response to the demands of specific customers; and (ii) once they were acquired, the MT industry was the main transmission center for the transfer of new skills and techniques to the entire machine-using sector of the economy." The vertical industry structure just described facilitated spreading the external effects of technology across the economy.

---

[5] Note how "technical" a view of innovation this entails. Learning about requirements "sets the stage" by posing a problem. Problem solving—oh, how engineers love problem solving! Problem solving is a characteristic of engineering invention or technical invention. In the Rosenberg model of MT, the problem is also determined by a purely technical requirement. I will argue below that this feature of the relationship between MT and MT use is important for learning—the problem to be solved, and the solution, are both narrowly technical. (This will be far less true in ICT.) Choosing what problem to solve can be a technical matter, but it also can fall in another sphere of knowledge whatsoever. Rosenberg (1969) later called out the interaction between deciding what problem to solve and technical progress in the form of solving problems as extremely important for technical progress. I agree.

2.2 Macro

By macro I mean Rosenberg's observations about the nature and role of the "entire machine-using sector of the economy" over the period 1840–1910 and later. He points out that the development and advance of the "machine-using" sector formed much of the technical progress in this era. This includes an impressive list of manufacturing industries. If it were not for the macro part of the argument, we would not be paying as much attention. The macro outcomes were important, forming an important part of why mid-20th century Americans were significantly richer than mid-19th century Americans.

How did the learning and the positive feedback loop come to be of economy-wide importance? What conditions were necessary for this? Why was the process of technical change across industries not just one thing after another? All these three questions get a common answer in Rosenberg's treatment. Here, I quote his own language, but I will give the key assumptions labels for generality.

1. *Widespread technological opportunity:* "most machinery production poses a broadly similar set of problems"
2. (Related) *Widespread economic opportunity:* "All innovations-whether they include the introduction of a new product or provide a cheaper way of producing an existing product require that the capital goods sector shall in turn produce a new product (capital good) according to certain specifications."
3. *"Internal" loop:* "[M]embers of the producers' durables industry have an internal motivation to improve their own techniques in the production of the durable goods themselves."

If these macro conditions are met, then a pattern of invention like that in a series of new machinery-production industries will have economy-wide implications.

Neither the macro nor micro conditions provide a particularly complete explanation of the MT Industry *market structure*. Partly they work through an extent-of-the-market point "increasing vertical disintegration from the point of view of a single industry was accompanied by technological convergence of larger groups of industries." That part is easy to understand. However, Rosenberg also notes that the upstream industry saw "the emergence of large numbers of producers each of whom typically concentrated on a very narrow range of machines," which is harder to explain.

From a macro perspective, what is essential here is that there are elements of technical progress that are widely shared across the economy, capable of continued improvement, and ultimately highly valuable to using industries. (This is the story of GPTs). From a micro perspective, what is essential here is that there is learning based on technical and market conditions in one part of the economy that spreads out to become useful elsewhere.

3. ICT macro

ICT had a widespread technological opportunity in the late 20th and early 21st century. However, in contrast to MT, there was no "broadly similar set of problems" in "ICT production." Instead, long "stacks" of difficult-to-invent hardware and software, some of it industry specific, underlie ICT-using production. These stacks permit solution of a widely dissimilar set of problems in different industries. ICT-using production processes in different industries are dissimilar.

The invention of the long stacks of software was made economic by two forces. The first was dramatic cost falls for the basic elements of ICT—digital logic and arithmetic, data storage, and transmission—that played out over decades. The basic elements became so cheap and powerful that they enabled the solution of a very dissimilar set of problems in different domains. And the widening span of application of ICT meant that the market for the basic elements was large enough to pay for ongoing improvements in them. (Fixed) costs have indeed risen dramatically. Rock's law, a sort of dual to Moore's law, posits that the costs of a semiconductor plant double every 4 years.[6] Paying these costs has been made economic by all the invention in uses for ICs.

I study ICT and its application starting after the Second World War, approximately a century after the beginning of the MT period studied by Rosenberg.[7] The economic growth situation in the United States was far different.

---

[6] Art Rock is widely credited with having founded the Silicon Valley venture capital industry.

[7] There had long been a telephony and a telegraphy industry before the second world war; their transmissions, while analog, included valuable business information (Field, 1992). There had long been business data processing, if not electronic. The emergence of digital computers as a capital good for business data processing and as both a factor of

The society was richer, with no obvious analog to "industrialization" looming, though some observers thought that automating white collar work would be the "industrialization" of service industries. The workforce was better educated, and the transition out of employment largely in agriculture to hardly in agriculture at all was well under way. A great deal of technical change had occurred in the intervening century, and much of it had been usefully deployed into the economy.

Historians of technology and of the American Economy will doubtless think of many other points.

While the applications of MTs fall into "machinery production," as Rosenberg put it, the applications of ICT have been very diverse. MT was linked to *all of a growth pole*, industrialization. ICT has been linked to *all of the economy*. It is not just that white collar work is everywhere. ICT has been used as a scientific and engineering tool and has been the backbone production technology for many functions in many industries. Both MT and ICT are general purpose technologies, but the generality and scope of ICT is broader.

This generality and scope is, however, not a technological fact but an economic one. To understand their causes, we need to open up the *relationship* between fundamental ICT technical progress and economic value creation. As we shall see, sometimes that relationship is much like that of the MT era. Rosenberg wrote that "An explanation of many of the technological changes in the manufacturing sector of the economy may be fruitfully approached at the purely technological level," an invaluable interim conclusion of his study. The same interim conclusion is true for some of the applications in the ICT era, particularly for applications in science and engineering. However, the same interim conclusion does not hold for enterprise ICT. The invention of the applications of ICT in much of commercial and enterprise uses necessarily takes the analysis outside the "purely technological level." We will need to add some elements to the mix—while keeping the very useful insights of Rosenberg's analysis in mind.

First, commercial innovation is unlike technical innovation in important ways. Second, "learning" from earlier innovation that can guide future innovation is different when it occurs within a technical discipline than when it includes both commercial and technical innovations. As a result of these two differences, a GPT with commercial innovations in its applications will be very different than a GPT with only technical inventions in both GPT and applications. As we shall see, these differences will be key to understanding the differences between the learning process that occurred in MT (and machine-using industries) and the learning process that has been occurring in ICT and (ICT-using industries.) We shall examine these differences in learning below, looking first at ICT segments where the learning process quite resembles that in MT, in no small part because a "purely technological level" approach is effective.

An essential part of the Rosenberg analysis was of the external effects that arise through the positive feedback loop shown, alongside the learning cycle, in Figure 1. We now recognize these as social increasing returns to scale (Bresnahan and Trajtenberg, 1995; Bresnahan and Greenstein, 1999). For our present inquiry, it is important to keep several different notions of "scale" clear. The first is directly related to the external effects. The more industries are added to the loop, the more opportunity for spill outs to any particular industry—a scale economy that rises with the scope of reuse of the general ideas across industries. Second, sometimes there are scale economies at the firm level in the upstream industry itself; these are very important in ICT, where one might (e.g.) need to build an extremely expensive factory (a "fab") to make integrated circuits or (e.g.) to have an enormous server farm to store the data used to guide search, or (e.g.) maintain an expensive field sales force to sell mainframe computers or business applications or (almost everywhere) need expensive R&D, a fixed cost. This will play a role in explaining the differences in upstream industry structure between MT and ICT, in part because of the success of upstream firms capturing the returns to social scale economy.

4. Segments of ICT are like MT

Let me start in areas where the Rosenberg model of the MT can be applied, with modest changes, to provide an explanation of invention in the ICT era. That the MT model is very useful in explaining some areas of ICT rules out many broad explanations of why it is less useful in other areas, broad explanations in which the present should never be explained in terms of the historical past. This will help focus our enquiry on the narrow and specific way in which the present is, in fact, different from the historical past.

production in telephony and a key complement to it through data transmission is the start I am talking about here. "Emergence" suggests something more rapid than the multidecade process, by this I mean only that it is a good time to begin to think about the amplifying cycle in ICT and its uses.

To make the comparison close, I will first focus on the ICT segments for which (i) the customer is primarily a technical professional, such as a scientist or engineer and (ii) supply is not regulated. This rules out the largest ICT segments such as enterprise computing hardware and software and regulated telephony. Nonetheless it will serve to illustrate the generality and power of the Rosenberg analysis, and of comparisons to historical eras. The ICT segments with technical buyers, such as minicomputers and workstations (before they entered enterprise computing) parallel the Rosenberg analysis of MT in one close and interesting way. ICT segments without any "buyers" at all, such as the Internet before its commercialization, offer a different and interesting parallel.[8]

4.1 Scientific and engineering computing

As its name suggests, the original purpose of the digital computer when invented was to do .... computation. Many of the significant early inventors had military contracts or military projects.[9] More importantly for our purposes here was the content of those projects. The computer was invented as a scientific and engineering tool. Eckert and Mauchly, in their contract for the Army, were working on a military-applications engineering tool—the computer would solve differential equations numerically so that, for example, artillerymen could aim accurately even if the wind were blowing. The scientific team tasked to design the H-bomb also knew it needed to do complex calculations. As a step in that direction, they invented the stored program computer as a scientific tool.

In Economics, we are familiar with spill outs from military invention via reuse in civilian life. Many such spill outs arise through a mechanism of creating a product which is dual use. For example, the radar first located military aircraft (most valuably, hostile ones) and then located civilian aircraft (most valuably, ones not quite in the right place.) The computer spill out was, instead, much like the spill outs examined in the Rosenberg MT analysis. The invention of a scientific or engineering tool in one domain, in this case military, preceded its use on other scientific and engineering projects in many other domains, not all military. The computer as scientific and engineering tool was like a MT. The applications of the computer in a scientific or engineering task were like MT-using production steps.

If we take Figure 1, above, and make one change, it admirably illuminates the technical change process in scientific and engineering computing. We need to change "product"—new products were essential in the MT history—to "task." By "task," I mean something a scientist or engineer does, not a final product. In taking out "product," we also have to take out "production." Instead of "production requirements," think of "task size and scope" and instead of "production problems," think of "technical task challenges." There is no change in the logic, only in the nouns—so Figure 2, which refers to this part of ICT, has the same shape as did Figure 1.

This particular positive feedback loop took off in the 20th century in much the same way that the MT loop took off in the 19th. The solution of each of a series of new difficult scientific or engineering problems led to the creation of a new tool, a new kind of computer. Its existence demonstrated a technical capability of doing calculations. Other scientists and engineers who had a calculation that demanded about as much capacity could use the same tool. Other scientists and engineers with somewhat more difficult and demanding problems could see the value of a more powerful tool. An industry grew up in which suppliers—sellers of minicomputers, later workstations and embedded computers sought to assess that demand and provide the more powerful tool if there was demand for it, as there was.

In earlier papers (Bresnahan and Greenstein, 1999; Bresnahan and Malerba, 1999), I noted that these segments—the ones with technical buyers of technical products—did not need extensive investments in marketing by sellers. Buyer and seller had a common, technical, language, and sales could be based on technical criteria. In this essay, I am using that same observation to make a point about the major steps in the Rosenberg "learning" cycle. With that common, technical, language, sellers can learn what buyers' needs are; the requirements placed on sellers to provide a tool that would let buyers do a new task can be stated objectively; and, of course, the demonstration of success in

---

[8] The differences across ICT segments in the commercialization process, and the resulting implications for differences in industry structure, are ones I have written about before with Shane Greenstein (see our 2001 paper) and with Franco Malerba (see our 1999 paper).

[9] There are a number of excellent histories of these early developments—cited in my papers with Shane Greenstein and Franco Malerba in the bibliography to this paper. In this brief paper, I am not going to attempt to imitate Nate Rosenberg's command of the historical and technical literature. So those who wish to go back to important primary and secondary sources will need to jump first to more historically complete papers cited here to get pointers to them.

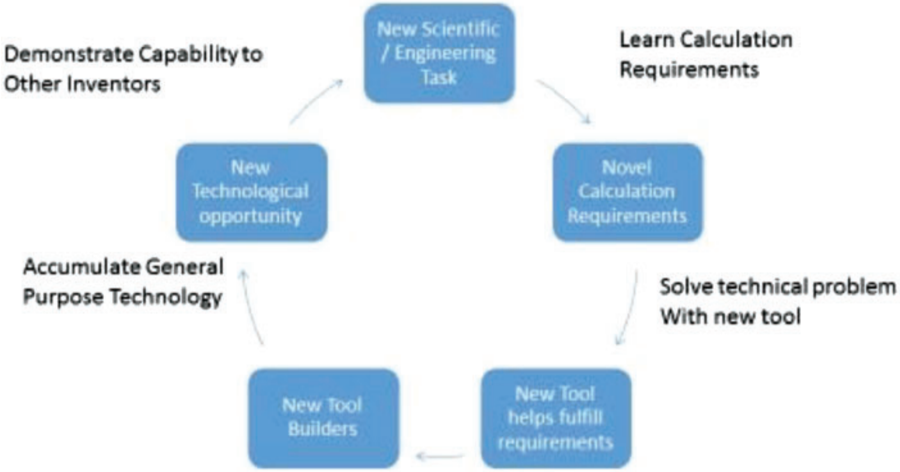

Figure 2. Schematic of Technical Computing Analysis.

one domain can be easily explained, in the common language, to potential new adopters of the tool. All of the "learning" steps in the learning cycle are made easier by their confinement to a single technological area.[10]

This argument is very close to the one that supports Rosenberg's claim that MT innovation and innovation in the use of MT could be explained "at the purely technological level." Both MT and this part of ICT have a single technological area which encompasses both user innovation and capital equipment producer innovation. This facilitates buyer–seller communication. It also facilitates learning in the Rosenberg sense. Note that it is important to this argument *both* that the knowledge is *technical* and that it is a *single area* encompassing sellers and users of capital. This appears to be the essential assumption for a positive feedback loop between suppliers and sellers—a fairly common occurrence—to have the Rosenberg learning outcome.

4.2 Growth of a specialized labor market

There is another area of ICT usage which is not a "sector" of the economy, but which has grown very large—computer programming. This has also functioned as the Rosenberg model suggested, and for much the same reason. Both the applications innovation and the GPT innovation were technical.

As minicomputers came to be used more and more in factories and laboratories, a specialized computer-programmer labor market emerged. Programmers were not exactly scientists or engineers (though technically sophisticated scientists and engineers often did their own programming), more like lab techs.[11] At first, these computer programmers dealt with the problem of maintaining the computer itself, the hardware. They also programmed it, at first in very primitive languages. Over time, they came to primarily maintain and improve the software running on the computer, and to program the scientific and engineering calculations in high-level languages (FORTRAN was an important example).

Programmer productivity improved rapidly.

Part of the reason for the rapid increase in programmer productivity was the market response to programmer needs. The growth in computing meant a growth in programmers. A tool that would make technical tasks easier for all those programmers would have a large market. This market opportunity was first fulfilled by making hardware

---

[10] The idea of a single domain of knowledge, within which learning about the inventions of others is easy, appears to be a different one than other remarks about the structure of knowledge. There is an excellent review of these in Malerba and Orsenigo (2000). [It is related to, but not the same as, recording knowledge in a way it can be widely accessed (Mokyr, Stern).]

[11] Of course, this labor market extended to cover the rest of "computing" as well. One cannot really separate ICT segments for purposes of understanding the overall rate of technical progress. There were large spill outs across segments, many of them flowing through the labor market.

easier to use and by making software tools for programmers, such as high-level languages. Programmer productivity presented a visible bottleneck, and both integrated hardware companies and (after a while) independent software companies provided programmer tools. Improvements in the quality of these tools are, of course, exactly the kind of technical progress contemplated by Rosenberg at the "solve technical problems" link of Figure 1 or Figure 2. My real point here is that the prior step, "learn requirements" was reasonably simple and direct in the industry that made computers and software as engineering and software tools, just as it had been in the MT industry. The supplier of a tool was, in each case, working within the same area of technical knowledge as his future customers. He could see their needs without crossing the boundaries of areas of knowledge. This made tool invention easier on one dimension, that is, seeing the market opportunity, and tool makers invented some true technical marvels such as the database management system. This is the point at which there is ICT invention which is most like the "internal" loop identified by Rosenberg in MT—here improved software tools are used to write new applications software, while there improved hardware tools were used to make new applications hardware.

Computers for calculation spread rapidly beyond military science and engineering to other sectors of the economy using computation, notably to engineering, first in laboratories, then in factories (fast enough computers for "control" applications were an important advance here). University and government scientific, engineering, and military calculations rapidly became cheaper. First because machines rather than people did the calculation, then because programming the computer became easier, further economizing on human time. As the cost function fell and fell, more and more complex scientific and engineering calculations could be undertaken. Eventually, more and more control functions could be undertaken—computer numerically controlled brake shoes, who would ever have thought it!

A parallel example arises in the Internet. More on the communications side of ICT, it was invented for military purposes, spread to scientific and engineering laboratory purposes, and captured a stunning amount of innovation by users and others into a common pool (see Greenstein, 2015). I shall return to its governance and to its migration out of purely technical use below.

The parallel, in process, in structure, and in scope, to the MT story is obvious. The process was multiple trips around the loop. The structure was an upstream industry—and specialized labor market—that became a supplier to a wide variety of activities. The scope was that all those activities could use a broadly similar tool. A great deal of knowledge accumulated in the upstream industry, and was broadly useful. For our present purpose, the point is an agreement with the Rosenberg model. For advancing purely technical knowledge, knowledge that is a form of an engineering tool, the forces identified in his paper have been playing out over the last 60 years in ICT.

Both the advantage and the limiting principle in this part of ICT and in MT arose from the same source: similarity and shared scope of knowledge. The purely technical nature of both demand and supply was, as we shall see by considering areas with organizational and commercial innovation, essential to the shared scope of knowledge. Going beyond this area will take us beyond the simple application of the Rosenberg analysis, though many of its principles will still be with us.

5. Other segments of ICT are very different

When we turn away from technical computing to ICT segments with nontechnical customers, the history shows us some important differences from the MT model.

As a threshold point, I should point out that these nontechnical categories, especially business data processing, have been significantly more valuable than technical computing. To be sure, there were tremendous accomplishments in scientific and engineering computing. However, it is first in enterprise-oriented computer applications in the later 20th century, then in consumer-oriented applications of the early 21st century perspective, where we find truly large societal value creation flowing from ICT.[12]

The core difference between these categories of ICT use for our purposes is the nontechnical part of innovation in applications. Many enterprise applications of ICT today are extremely sophisticated. In some areas, ICT even *is* the production process, as when a computer sells you an airline ticket over the Internet. These sophisticated applications are also very difficult to invent. They have had an indirect invention process, starting at more visible subsystems and only slowly growing out to the complex, sophisticated high value business data processing systems in use today.

---

[12] The history of ICT applications is less well studied than the history of ICT itself. However, see Cortada (2006) for a magisterial look at the history of commercial applications.

The indirection arises from problems of visibility. And, as we shall see, problems of visibility have influenced not only the invention of applications, but the coordination of ICT invention and of applications invention and also the accumulation of learned knowledge. Ultimately these will lead to a different schematic model of the GPT/Applications feedback loop than we have seen so far.

5.1 Commercial innovation in applications companies in Business Data Processing

I begin with examples where the value proposition of the business data processing systems ultimately built on top of the ICT is obvious, so one might be tempted to think that the spread of ICT toward those new applications was also visible *ex ante*. However, in these quite representative examples, rather the reverse was true, the application of ICT to the "obvious" end followed an indirect and twisting path—because of the difficulty of seeing the way to accomplish commercial innovation on top of ICT.

One impressive contribution of ICT at the company level is systems that permit complex pricing over time. At most airlines, for example, there is a system of volume discounts that benefit historically high-volume customers. These are typically implemented in the airline's reservation and sales system. Of course, discounts for particularly valuable customers existed before there were airlines, much less airline reservation systems—this was not an entirely new commercial innovation. But such discounts, before ICT, typically occurred within large enough buyer–seller relationships large enough to repay the recordkeeping, negotiating, and deciding on a customer-by-customer basis.

But the invention of that highly valuable form of application, did not follow "immediately" or "directly" from technical capability, as one might have expected from the example of MT. Instead, airlines first built reservations systems that kept track of the inventory of seats. This was a very visible payoff application, and did not require understanding the behavior of customers. It had an immediate and direct operational payoff. Only after that simple inventory-management system was in place did higher-value additions come into use. Those more complex and valuable systems were made more visible by the intermediate steps. Copeland and McKenney (1988) summarize the lessons from airline reservation systems for firms seeking to gain competitive advantage from ICT as requiring a substantial investment in ICT and:

> [E]stablished technical competence is a necessary requirement for gaining competitive advantage. … [S]ustainable advantage need not be the result of extraordinary vision, but the result of consistent exploitation of opportunities revealed during the evolution of adaptable systems.

This is a famous example. But it is broadly similar to the path to many other high-value ICT applications. The creation of new ICT capabilities typically led first to a visible application. Once user learning from the visible application had occurred, or the visible application had created intermediate inputs, such as a database, new innovation could be contemplated. That roundabout innovation process then led to the high-value innovation. Typically the earliest application was narrowly operational (like the first airline reservation system) or narrowly accounting, and later improvements were analytical or marketing or operational (like the reservation system selling tickets). Whole technological movements emerged to support those later improvements, such as the Decision Support movement.[13] Copeland and McKenney (1988) are correct when they write that in general (i) the relevant technical competence is going to need to be established *in the using company* and that (ii) learning "during the evolution of adaptable systems" by the ICT-using company is key.

The use of "evolution" here echoes an argument of Simon (1962): The "existence of stable intermediate forms exercises a powerful effect on the evolution of complex forms." In the "evolution" of a business data processing system, the "simple form" is originally visible; it involves mere automation of an existing task. But it is stable, that is, profitable, and forms the basis for later "complex forms" that may change the firm's way of doing business. While the "simple" forms can be forecast, or at least can be forecast by some of the most brilliant business people of an age, the "complex forms" for which they will later form the basis are very difficult to forecast. They are not visible.

This leads us to discover differences from the MT case and the technical uses of ICT case. If it is difficult *ex ante* to forecast the commercial applications of ICT even inside a particular company, will likely be difficult to ex ante

---

[13] A human decision could be supported by analytics running on the ICT system—a step that has led to many modern jobs for humans that are highly ICT-using, where the human sells to the customer (or negotiates with the vendor) while the ICT advises/imposes rules/sets priorities and tradeoffs/records/etc.

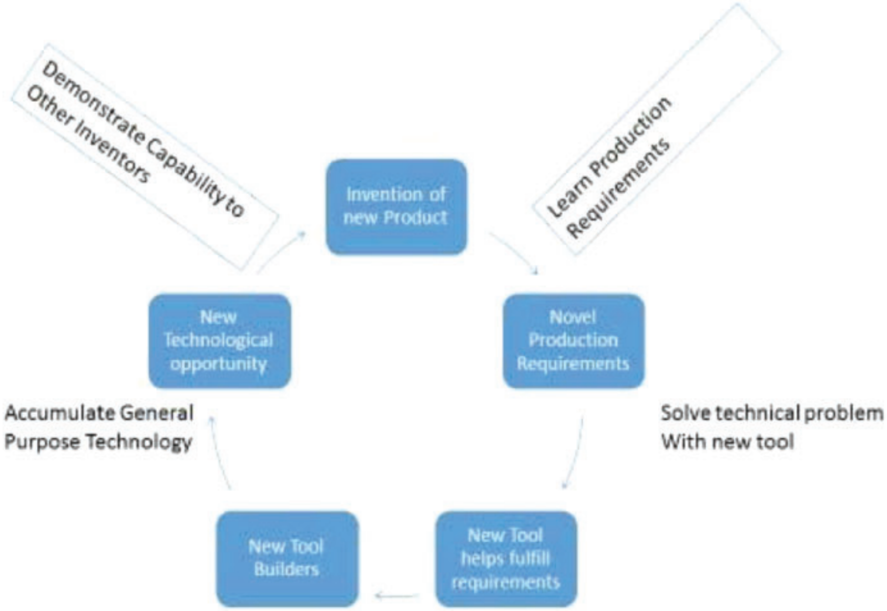

Figure 3. Schematic of BDP Analysis.

coordinate ICT and applications innovation. Second, and less obvious so we shall come back to it, ICT learning at the boundary between the technical and the commercial is indirect and to some degree inverted. This aspect of the individual company learning process has profound implications for *social* (multi-company) learning. There are two points in Figure 1 at which the Rosenberg innovation cycle crosses the technical/commercial boundary. They both have changes from the MT case.

One change is to the idea, so important in MT innovation, that a new industry or a new product would lead the MT-using company to "learn production requirements" which could be fulfilled by a new MT-using step. In the ICT era, for the most valuable uses, this is not an accurate description. Instead, of "novel production problems" being "posed" and then "solved," we see the joint invention of new products, new technical requirements and new production processes over a period of years within the firm.[14] In Figure 3 this is reflected by making "learn production requirements" into a tilted line—this step, because so different, represents a gap between private (single firm) and social (multi-firm) learning in the ICT era. Part of the gap is that much of the knowledge accumulation stays in the single firm. Indeed, using firms have tended to treat their commercial innovations as competitively sensitive and protect them from imitation if possible.

Another change from MT innovation is in the "demonstration" of a new capability invented in one place to other inventors. In the high-value uses of ICT, when the ICT infrastructure products that supported valuable innovation at one firm are offered to another firm, acceptance typically follows visibility rather than being "direct" or guided by value. Indeed, new opportunity often sets off a reverse product life cycle in the new firms using it, in which the technologies which will ultimately support a valuable product are available, but the product itself is not visible. The path to that product may, at the firm where the spill-in might occur, be long and indirect. Accordingly, I have marked the "demonstration step" in Figure 3 as also out of parallel and out of sync. Above and beyond the problem that demonstrations may not be informative is another problem that we will turn to below, the demonstrating firm may not want its innovation to be demonstrated to competitors.

---

[14] These are not "novel production problems" which are "posed" by a new product or industry and then "solved" technically. At the closest, there is the "requirements specification" for a new ICT system—one of the most notoriously difficult parts of the ICT innovation cycle.

To capture the differences, I have made make a new version of Figure 1, in Figure 3. The logical sequence of steps is the same. However, the difficulty in seeing the application of general components at each new application innovation step and the difficulty of seeing the general inventions that are called forth by new applications are important changes. The two steps which have been identified as different are marked now by having a slanted entry in the cycle—they are bottlenecks, waiting for commercial innovation and waiting for innovation that sees the overlap between the technically possible and the commercially valuable. That different figure means to capture the essential differences between high-value MT inventions and high-value ICT inventions a century and more later.

When I reused the Rosenberg logic, essentially without alternation, in the narrowly technical ICT segments, I kept the same figure and changed the nouns. When, however, I tried to apply it to the same phenomena in the later era—thus able to keep the same nouns—I needed to change the logic a bit. This illuminates the Rosenberg model as well as the world, I hope.

5.1.1 Implications

These characteristics of the innovation process in the business data processing part of ICT have led to a number of industry institutions and market outcomes. I begin with the progress of learning, at the applying firm level and then at the economy-wide level.

One implication, implicit in the examples above but quite general in the process that has led to valuable ICT-based applications, is that the path to a valuable application at the firm level is often indirect. The first stages of this path often lead to creation of a system inside a company for which *both* the value proposition and the mode of implementation are visible *ex ante*. As is the case with accumulative learning generally, more will be visible *ex post* the creation of the first, partial system. Learning is accordingly slowed; slowed mechanically by the time taken passing through the *ex ante* visible stage, and slowed by smaller incentives for applications firms to innovate. (The incentives will be smaller whenever it is uncertain that the *ex ante* visible stage will lead to follow on innovation.) The slow firm-level learning, especially at the early stages, is an important difference from the Rosenberg model of MT. It is not the same as visible requirements eliciting solutions.

A related implication arises in the diffusion of new ICT across firms and industries within business data processing (BDP). Bresnahan and Greenstein (1997) examined the path of diffusion of a new technology into enterprise ICT sites. The path did not take the form of high-value-first (as in the model of Griliches 1960) but of high visibility first. That was a systematic statistical investigation, but historical investigations have also shown that the areas—functions, production processes, industries, etc., which have taken up ICT for BDP have tended, generally, to be the ones where there is first visibility both that the new ICT might create value and about how to implement to achieve that value.[15] This slows economy-wide learning. The learning about the most valuable areas is delayed. It is delayed mechanically by waiting for learning in other areas and by the lower incentive to innovate in the earlier, more visible applications (unless the most visible and the most valuable have been the same, which so far has not been systematically true).

It is important to note that this cause of slowness is above and beyond the externality across innovators we always see with regard to any widely shared technical progress. Rosenberg noted the externality with regard to MTs; my point here concerns the additional problems associated with identifying the most valuable applications of new ICT waves and with discovering the way to implement those applications. The particular difficulty of making these commercial innovations is a consequence of the need, within BDP, for learning to cross the technical/commercial boundary. In turn, it causes problems with the rate of commercial innovation and value creation.

In short, at the early stages of the application of most of the important rounds of ICT in BDL, there has been sharing but not convergence. This flows primarily from the incentive implications of the problem of visibility.

---

[15] See Cortada (2006), p 123-4, for an example. He discusses early applications of ICT in insurance companies, noting that merely replacing clerks with machines came first, after which “managers were also coming to realize that computers could also make information become more conveniently available” and that a decade later manages came to understand that “implementing computers could also lead to changes in the way their firms functioned.” Variations of this pattern of visible first, more complex later, fill the history of ICT applications.

5.2 Convergence at the very end of the process

I do not mean that there has been no element of convergence in business data processing. In fact, there have been two important elements of convergence. Each illuminates the applicability of the lessons of the NR model to modern times, if we only keep our eye on the conditions required for the model.

First, there has been a great deal of convergence in *programmer tools* used by developers who work in and for enterprise computing sites undertaking BDP. Indeed, it is very hard to think of another brain-work job that has had as much automation as that of the computer programmer. The mechanism is much like that that worked in MT. A particular problem in coding, in building systems, in maintaining systems, etc.—all the purely technical tasks of running an ICT department—raises visible needs for new tools which are seen as requirements by the using engineers. The requirement is fulfilled by the creation of a new tool, and the tool comes to be supplied by the ICT industry to programmer sites in all the customer industries. A learning cycle and a positive feedback loop breaks out within the purely technical parts of computing, excluding the commercial application. This has been important technical progress,[16] though the structure of ICT supply, to which we shall turn in a moment, has often meant that the applications industries have been charged hefty fees for the resulting programmer tool products.

A second area of convergence arises once a particular ICT application has diffused widely through using industries and no longer confers much competitive advantage. At that stage, ICT-using firms will more likely choose to obtain a solution from a vendor than to have a unique application in house. A wide variety of applications markets grew in BDP, accelerating somewhat in the 21st century—after many areas of applications had somewhat matured. Today, many applications sectors have a wide range of choices for applications. They can buy a software license, and run it on their own systems. They can outsource a function to an ICT services firm that will run it for them on its systems. Or they can obtain a wide variety of "cloud" or "software as a service" (something as a service, where "something" can be a wide variety of things) that lets them easily access the inputs and outputs of their BDL system from a variety of devices and locations. All of these different forms have been growing for years. What is new is the acceleration, in the late stages of diffusing of certain areas of application, of shared commercial systems as well as shared technical systems.

Again, the punch line is that within the commercial side of ICT usage, where much value has been created, there has been a great deal of sharing of technical inputs across industries and using firms, and even some convergence. That convergence, however, has not been an important part of the process by which the learning about commercial innovations took place.

5.3 Apparently shocking leaps

Visibility problems lead to accumulation of knowledge about different areas within different disciplines, sometimes with an incentive to keep knowledge secret. This process leads to the accumulation of distributed knowledge a la Hayek. Both knowledge of potential applications and knowledge of technical opportunity are advancing, but the distributed nature of knowledge means that the overlap between technical opportunity and commercial value is not known. This is an important difference from the common pool of accumulation in the Rosenberg analysis. It leads to apparently shocking leaps in the overlap between technological opportunity and commercial value creation. In the history of ICT, there have been a number of these leaps; Bresnahan and Greenstein (1999) emphasized the leaps by competitive entrants into existing categories. Here, I discuss only one leap, the transition of the Internet from a military, scientific, and engineering technology to one widely used by consumers and in commercial application.[17]

The commercial needs most quickly assuaged by the widespread use of the Internet were mass market electronic communications, mass market electronic commerce, and mass market electronic content provision. These needs had been clear for some time, as there were important existing markets in electronic communications, commerce, and content provision, and the widespread use of PCs made mass-market variants more obvious. Yet the many well-funded efforts to create these markets before the widespread use of the Internet had largely come to naught. It was only after a large mass of users moved to the Internet that a way to operationalize these markets became visible.

---

[16] "Programmer tools" sounds narrow, but is actually a very wide category. For example, the relational database management system (one of the most valuable inventions of the 20th century) is a programmer tool.

[17] See especially Greenstein (2015) for the history of this migration. For the analysis of recombination and of anticipation in this section, I draw heavily on Bresnahan (2012).

There had been a great deal of technical progress in Internet technologies, which had accumulated.[18] However, it is not only the technical level of the Internet at the time it came to be widely used, but the industrial organization of it. Because the Internet heretofore supported scientific and technical computing that general elements of its technical progress largely accumulated in an open-systems way, not as part of an upstream vendor's proprietary platform. Further, some application elements, perhaps most importantly email, also accumulated in an open-systems way. With no controlling vendor to slow it down, the open-systems internet could add new technical elements such as the world wide web (WWW) and the browser.[19] The parallel to Rosenberg's analysis of the MT industry is obvious. Software is a special kind of tool with zero marginal cost, so the "upstream industry" for Internet software was even thinner than the small firms sketched by Rosenberg. Much applications learning proved to be general and accumulated in an Internet knowledge pool rather than in the using laboratories or factories.

With email, the WWW, and the browser available, the mass market Internet quickly added a large number of users. The commercial fact of having many users, along with the openness and the high technical level of the Internet, led it quickly to fill the three longstanding needs. Interestingly, using firms have overwhelmingly succeeded in keeping the Internet open to this day—there is no I important Internet vendor with a position anything like that of IBM in the mainframe or Microsoft in the PC. That is an important difference to which I now turn.

5.4 Marketing investments on the part of ICT vendor firms in business data processing markets

One striking feature of the ICT market segments serving nontechnical customers—commercial businesses and consumers—is the emergence of large vendor firms with considerable market power. Largely absent in the MT era, it is an important difference to explain.

While the ICT vendor firms selling to technical customers, we saw in earlier sections could focus primarily on technical advance and make a technical sales pitch to knowledgeable customer, business data processing firms have needed to support customers with field sales forces and closely integrated support for the customer's innovation process. This marketing capability is complementary to technical progress made by the ICT vendor. It is also complementary to technical progress made outside the ICT vendor and captured via a "strong second" strategy. These marketing investments have been a central part of the strategy of enterprise computing vendors from IBM to EDS to Oracle. The close engagement of the vendor's marketing efforts with the customer leads both to support for the difficult process of requirement specification by the customer and learning by the seller about the customer's innovation.

5.4.1 Emergence of very successful firms/profitable—marketing investments

Those marketing investments were essential to the emergence of very concentrated industry structures in the BDP segments of ICT (Bresnahan and Malerba, 1999). The leading firms in these segments gained much of the rents that occurred as a result of the collective learning process (Bresnahan and Greenstein, 1999). These features of market structure follow from the need for successful firms to invest in close marketing connections to customers.

5.4.2 Anticipating (or not) the next innovation

Some of the most important commercial innovations by ICT vendors in BDP, that is to say, some of the most profitable innovations of the 20th century, were designed to deal with problems raised by visibility.

In the mid-1960s, IBM invented the (proprietary) modular platform. A customer buying within the S/360 family could mix and match storage, computer, and input/output. More importantly, if a customer discovered that their application needed more storage, or better Input/Output or more computation, the customer could upgrade the relevant components while preserving their investment in other components. This fit quite well with the adaptive and exploratory nature of many enterprise software applications innovation processes. The modular platform permitted a customer to undertake the foreseeable part of an application and, if there were learning about how to improve the application, to grow later. An IBM System 360 customer who discovered that an early application was successful—whether after trying some risk or doing the foreseeable part—but needed more storage to achieve more could keep

---

[18] This is an important general equilibrium point, emphasized by Bresnahan and Greenstein (1999)—there is a great deal of flow of technology, and not just the three ICT basic elements, across segments. It is often relevant to these leaps.

[19] See Bresnahan and Greenstein (1999) for a discussion of why the important leaps into commercial computing have most commonly come from more open-systems organized segments.

the existing computer and add more disk or tape storage. The (proprietary) modular platform keeps the customer for IBM while dealing with the problem that the requirements specification part of the innovation amplification cycle is very difficult to forecast or execute.

Also in the 1960s, Ross Perot invented a particular model of an independent software vendor—one that wrote applications software once and resold it many times. This is another commercial innovation that could grace anyone's resume.[20] Perot's EDS had a number of important implementation decisions, including building an organization which could be at customer's site and engaged in the customers' innovation process. For our purposes here, I want to emphasize the very early phases of EDS and the problem of who owns the innovation involved in a new application. Perot's idea was to work very closely with one customer—as a contractor, paid by the customer!—to learn the requirements for, write, and then implement, application software, and then later to resell the software to many other customers. Problems would arise with this scheme if the other customers were competitive with the first one, so Perot, aiming at financial services, sought out what were (then) local monopolies, the "Blue" sellers of health insurance.

Once EDS was up and running, it did not need the first customer to pay for the first copy of the application software, though it did need a customer to cooperate in the requirements specification. Attractive pricing of consulting services to the first customer, followed by reselling the application software to many others, often worked for the firm and, later, its competitors. Still, cutting-edge applications firms often keep requirements specifications and application software in house to prevent it from being resold to competitors.

Both software reselling and (proprietary) modular platforms were so profitable for sellers that they became institutionalized in ICT supply. That occurred despite the fact that these institutions reflect another important element of the innovation process, *goal conflicts* between vendors and customers. Customers do not seek to be locked into incumbent proprietary platforms, nor do they necessarily want their difficult-to-specify requirements benefitting their competitors. This underlines important differences from the MT era. First, applications firms in the ICT era—leading edge applications inventing highly valuable new applications—sought to *prevent spill outs*. Second, upstream vendors came to have market power and sought to partially *capture the rents* from their customers' innovations. These incentives were not fully internalized by contract. Instead, the goal conflicts became an impediment to simple economy—wide learning. These goal conflicts are particularly an impediment to the accumulation of learning in the upstream industry.

5.5 Punch line

Examining the (typical) innovation process in high-value business data processing has proved remarkably useful to our inquiry. Because that process necessarily economizes on information about the overlap between commercial value creation and technological opportunity, it is typically incremental and involves much recombination *within the firm*. Because the most innovative applications are difficult to invent, user firms view them as competitively significant and thus seek confidentiality. These basic structural facts about the commercial implications have four implications:

1. There was a powerful incentive to invent the proprietary modular platform and thus to have large, highly profitable upstream firms.
2. There was a powerful incentive for using firms to limit the degree to which learning accumulated upstream.
3. Spill outs from innovation at one user firm or in one user industry to others were slow/
4. Spill-ins to particular firms and industries were slow

All these effects—the key differences between ICT innovation and MT innovation in the Rosenberg account—were caused by the problem of vision and visibility in ICT-using-firms' innovation. Commercial innovation, or perhaps more precisely, commercial innovation that takes advantage of new technological opportunity, was systematically hard to see *ex ante*. This explains the different relationships between upstream vendor firms and using firms, the limits on the process of accumulation or learning, and much of the pace of innovation. In the ICT ear, it is as

---

[20] Capital's share of GDP has grown dramatically in our century. Capturing the rents to the proprietary modular platform and to the software that is written once, then sold many times (plus their follow-on inventions) accounts for a large portion of that increase.

important to think about the as yet-unlearned as a bottleneck as to focus on the learned as an opportunity. Rosenberg never really emphasized the slowness arising from the externality across MT users, and rightly not.

6. The near future: scale as a "visibility"

The problem of *ex ante* visibility, limited foresight of commercial innovations, intermittent great leaps—is there nothing analytical that can be said *in media res?*

Technical progress in our young century, as measured by aggregate output growth in rich countries, has been slow by post enlightenment standards. Existing clusters of invention, some in their later stages, have continued: what little remains of work involving physical activity by humans in factories, farms, etc. goes on being automated away; the ICT-based innovation waves just discussed, reinforced by the widespread use of the Internet, has continued, among many others. The question on the table is whether the *new* technological areas opened up by *new* ICT are going to make large contributions to economic growth.[21]

To begin to frame this question, we need to think about what those new areas are. We have seen an enormous growth in new consumer-oriented ICT-based services. The most successful of these so far have been from Google, Facebook, Amazon, Apple, Baidu, Tencent, and Alibaba but there are a number of others. Much of the money—so far—has come from matching a consumer buyer to a seller, as when a seller pays Google or Facebook to place an advertisement, or a storefront is on Amazon, or Uber matches a driver to a passenger, or an ad runs in an app on an iPhone. Another important output growth area lies in extending the definition of media, news, and entertainment. While there is significantly broader growth in the use of ICT, and while there is a great deal of diversity in innovation among the new consumer-facing areas, if we are to look for new general-purpose ICT, it is here that we will find it.

A cluster of technologies linked to the label "big data" have originated in these new consumer-oriented businesses and may be being "demonstrated" to the rest of the economy. Consumer-oriented web and mobile applications generate enormous data streams. The label "big" correctly suggests large scale, a topic to which we shall turn in a moment. But "big" also connotes speed—a classic big data application might serve an advertisement on a website or in a mobile app while the consumer is still looking at it. Also, "big" suggests heterogeneity. In contrast to a traditional (e.g.) transactions data set with known data definitions, "big" data applications often put together much less structured information from consumer-typed text to clickstreams and drawn from a wide range of sources. Consumer-oriented ICT-based firms have built technologies associated with capturing, storing, managing, and processing these data. These technologies have gained enormous interest outside the firms that pioneered them, with many ICT-using firms noting that they, too, have large data streams and perhaps could profit from them.

Closely complementary to these data technologies are a cluster of analytical technologies. "Big" data are a mess, and many of the analytical technologies have a strong statistical flavor. Prediction—e.g. predicting the probability the user will click on an advertisement—is the center of the new statistical techniques, not inference. Learning about differences across consumers and updating the prediction model is an important current application of machine learning. These new underlie the enormous growth of consumer-oriented sectors of ICT-based, including search, social media, mobile applications (including the largest, search and social media) and a number of new buyer/seller matching technologies, such as online stores (whether accessed by PC, app, or voice, or by any of these or in a physical store), short-term room rental services, ride hailing services, music and other entertainment delivery services, and the like. While these categories contain a great deal of highly diverse technical progress, any enquiry about their general importance must look for their common themes. I am strongly drawn to two of these: consumer orientation and scale.

"Scale" in these applications means more than just volume, but it certainly means volume. Volume can most clearly be explained from the perspective of the underlying ICT basic elements. During the time, I took to type the

---

[21] A related but distinct question, too large for this brief essay, is whether they are going to make large contributions to shifts in factor demand. ICT has contributed to a large increase in capital's share (significantly larger than any seen in centuries) and in the share of labor income going to the higher-paid. Most forecasts on this front are unimpressive, wrapped around the twin observations that productivity growth is slowing and all human work is going to be replaced by new forms of robots are AI. (Note that those cannot both be correct.) If modern AI ends up replacing half as much human thinking over the next 35 years as did applications running on top of relational database management systems over the last 35, it will be quite the surprise and a very good thing.

preceding paragraph, uninterrupted, Google Search has done more calculation, stored more data, and communicated more information than did all of humankind up to 1980. This is enabled, of course, by the fact that computation, storage, and transmission have become very inexpensive—and continue to do so, so more scale is possible. But "scale" also means big data (not just large but complex), artificial intelligence (AI), analytics, and many technologies for capturing, (partially) controlling, visualizing, and acting upon very large bodies of information.

The first, visible applications of these high scale ICT methods has been to such simple, consumer-facing systems as advertising servers and recommendation engines ("consumers who bought this book also bought . . ."). These systems use approximately zero labor input, are clearly economic at very large scale, and are profitable because the stakes from making an error are small. By having a large volume of data on who clicks on advertisements, Google or Facebook can show an ad that has a very slightly higher probability of getting a click. Very slightly higher—if human operators had to examine 1/1000 of these ad placements, they would be unprofitable.[22] Thus, "visible," in the sense of this paper, has meant "low stakes" for the high scale applications of 21st century ICT.

The potentially very valuable cluster of technologies associated with 21st century "scale" may well diffuse from their current low-stakes application to much more valuable uses. This will involve technical progress in applications sectors paralleling that needed to make database management systems valuable—learning what are the costly exceptional cases, learning how to avoid or over-ride them, and so on. Many technologists and business people in ICT-using companies are working on such problems today. They are also working on importing other important broadly useful technologies from the consumer-facing sectors, including cloud, mobile devices, chat, and so on.

At this stage, we do not know what they will invent.

Many observers, particularly those focused on AI, believe they do know what those technologist and business people will invent. They use what strikes me as an incorrect model of the diffusion of ICT. They use an argument that might well have been right for MTs, but that is quite wrong for ICT. Technologists and technology enthusiasts are suggesting that AI, predictive statistical methods, big data management methods, and so on are about to revolutionize everything. AI can drive a truck! Surely it can do everything humans do at work! Surely there will be no human work at all before long! These unwise forecasts should be put next to the actual use of these technologies in business (see Bresnahan and Yin, 2016) and be confronted with a realistic analysis of AI diffusion.

The public debate about the future of AI is entirely confused about what are the economically important examples of its application. Do not think of AI-in-use as a driverless truck. Think of AI-in-use as a machine deciding what advertisement to show about a particular customer within a fraction of a second, creating and updating a model of whether the individual consumer will respond to the ad, and automatically learning over time about the individual consumer's advertising response function. Just to get you thinking like an economist, targeted advertising is up over a hundred billion dollars in revenues, while we hope to soon see economically meaningful applications of the driverless truck. The other economically important application to AI is in user interfaces, especially voice interfaces. Efforts to have voice interfaces lead to productive, rather than consumer-oriented, applications are largely in the future.

The public debate is equally confused about how the technologies underlying the early ICT uses will diffuse. Using companies are building "data marts" internally to organizationally manage big data, and inviting managers to find projects that will work. Who among our customers can we sell more to? Whom can we predict might defect to a rival? Can we sell people more online (in our stores) if we know what they buy in our stores (online)? When will giving people to opportunity to buy on their mobile device annoy them, and when will it move merchandise? These are extensions of the online giant firm big data methods to interesting—and now visible—alternatives. Business people will invent more and more of them—perhaps enough to turn these technologies into something really important.

These new inventions by ICT-using firms, to underline the message of this paper once more, are not at all the same as "well, what do people do around here that is like driving a truck?" Over the long and adaptive invention cycle within using firms, more and more complex processes will come into big data and associated analytics techniques. There are no real signs that today, it is particularly easy to leap to the end of that process with a visionary step. It is entirely possible that these may *become* valuable technologies. That awaits, however, a great deal of

---

[22] Set aside that many of these applications only work if they are near-instantaneous. Even if we humans could decide quickly, the cost, at 1/1000 human interventions, would look like the human cost plus epsilon.

learning.[23] As of today, the right choice between "we are in an era of slow growth of labor productivity" and "all human work is being automated away very rapidly" is the former.

The critical point looking forward is that future stages of exploitation of this technical and market opportunity, like others in ICT, depends on future innovation and learning, especially learning about what applications the technology will serve well. Needing to forecast that learning is entirely congruent with Rosenberg's analysis; seeing that the need for learning and discovery, especially the commercial/applications learning, is the bottleneck to new innovation and continued value creation, is what I think is the contribution of this paper.

7. Conclusion

Nathan Rosenberg gave himself a very difficult task when he sought to explain a large piece of technical progress in an important era by the microeconomics of the innovation process. He saw the task as "understanding historical events which otherwise appear to be random or capricious." The unifying forces that got rid of the caprice were a strong form of sharing—"convergence"—and the emergence of an upstream industry that advanced and spread the shared technologies.[24] Many arguments about cumulative technical change and about the difficulty of foresight give in to the dark side of this argument, over-embracing the random and capricious. This is an error, even though the importance of limited foresight and the problem of visibility of commercial value mean that there is even more difficulty in forecasting technical and market events in the ICT era.

Instead, we should embrace the wisdom of Rosenberg's approach and seek to understand the economics of the process guiding linked technical advances. The examination of ICT segments with largely technical customers shows that Rosenberg's analysis is timeless. The examination of ICT segments with enterprise and commercial customers shows that the analysis does not require as strong a concept of sharing as "convergence." However, the problem of visibility of commercial value and visibility of the overlap between technological opportunity and commercial value has meant that the social learning process emphasized by Rosenberg for MT has been limited in ICT—using firms have an incentive to keep innovations proprietary, which limits spill outs, and sellers of general components have an incentive to capture the returns to users' innovation, which further limits sharing and exploration.

The innovation process itself limits spill-ins.

The essential problem appears to lie in the difficulty of seeing overlaps between technical opportunity and commercial value creation at the early stages. As we have seen, that one problem, of "visibility" explains the differences in the nature of the upstream industry between ICT and MT, the limits on sharing and coordination, and the different structure of accumulation of general versus firm specific knowledge. Recognizing the gap between commercial and technical knowledge in ICT modifies the Rosenberg analysis and leaves it quite useful in our era. It also highlights the differences between a GPT where all of the invention is technical and contained within a single sphere of knowledge versus a GPT where applications invention lies in a different, commercial sphere of knowledge.

Acknowledgments

I am very grateful to the editors of this Issue, to the participants in the Conference on "Beyond Technological Innovation and Diffusion," and to two referees for valuable comments. More than anything, I am grateful to Nathan Rosenberg for a lifetime of his collegiality and his example in scholarship.

---

[23] At this writing, the most systematic focus of applications invention for the AI cluster is in health care. A broad array of tasks where a reasonable person could say AI has long been in use, such as in error detection systems (e.g. fraud detection in payment cards) is deepening with the new techniques—visibility in another dimension.

[24] Paul Strassmann (1963), in his Discussion of Rosenberg's paper, made the canny observation that the late 19th century saw the acceleration of "a new way of thinking" in manufactures, and that the learning essential to Rosenberg's story involved an overhaul of the "inter industrial communications system." Since another overhaul appears unlikely, "a similar phenomenon of convergence will not again be observed."

References

Atkeson, A. and P. J. Kehoe (2007), 'Modeling the transition to a new economy: lessons from two technological revolutions,' *The American Economic Review*, **97**(1), 64–88.

Bresnahan, T. (2012), 'Recombination, generality, and re-use,' in J. Lerner and S. Stern (eds), *The Rate and Direction of Inventive Activity Revisited*, University of Chicago Press for the NBER.

Bresnahan, T. and S. Greenstein (1997), 'Technical progress in computing and in the uses of computers,' *Brookings Papers on Economic Activity*, Micro., pp. 1–82.

Bresnahan, T. and F. Malerba (1999), 'Industrial dynamics and the evolution of firms and nations competitive capabilities in the world computer industry,' in D. Mowery and R. Nelson (eds), *The Sources of Industrial Leadership*. Cambridge University Press: Cambridge, pp. 79–132.

Bresnahan, T. and P.-L. Yin (2016), 'Adoption of new information and communications technologies in the workplace today,' in S. Greenstein, J. Lerner and S. Stern (eds), *Innovation Policy and the Economy*, Vol. 17. University of Chicago Press for the NBER, pp. 95–124.

Bresnahan, T. F. and S. Greenstein (1999), 'Technological competition and the structure of the computer industry,' *Journal of Industrial Economics, Wiley Blackwell*, **47**(1), 1–40.

Bresnahan, T. F. and M. Trajtenberg (1995), 'General purpose technologies "Engines of Growth?,"' *Journal of Econometrics, Elsevier*, **65**(1), 83–108.

Copeland, D. and J. McKenney (1988), 'Airline reservations systems: lessons from history,' *MIS Quarterly*, **12**, 3.

Cortada, J. W. (2006), *The Digital Hand: Volume II: How Computers Changed the Work of American Financial, Telecommunications, Media, and Entertainment Industries*. Oxford University Press: Oxford.

David, P. A. and G. Wright (2005), 'General purpose technologies and productivity surges: historical reflections on the future of the ICT revolution,' *Working Paper*, Stanford University.

Field, A. J. (1992), 'The magnetic telegraph, price and quantity data, and the new management of capital,' *Journal of Economic History*, **52**(2), 401–413.

Farrell, J. and P. Klemperer (2007), 'Coordination and lock-in: competition with switching costs and network effects,' in M. Armstrong and R. Porter (eds), Ch. 31 *in the Handbook of Industrial Organization*, Vol. **3**, Elsevier,

Foray, D. (2004), *The Economics of Knowledge*. Boston, MA: MIT Press.

Freeman, C. and F. Louçã (2001), *As Time Goes by: From the Industrial Revolutions to the Information Revolution*. Oxford University Press: Oxford.

Greenstein, S. (2015), *How the Internet Became Commercial: Innovation, Privatization, and the Birth of a New Network*. Princeton University Press: Princeton, NJ, USA.

Malerba, F. and L. Orsenigo (2000), 'Knowledge, innovative activities and industrial evolution,' *Industrial and Corporate Change*, **9**(2), 289–314 [InsertedFromOnline].

Nelson, R. and S. Winter (1982), *An Evolutionary Theory of Economic Change*. Cambridge, MA: Harvard University Press.

Rosenberg, N. (1963), 'Technological change in the machine tool industry, 1840–1910,' *The Journal of Economic History*, **23**(4), 414–443.

Rosenberg, N. (1969), 'The direction of technological change: inducement mechanisms and focusing devices,' *Economic Development and Cultural Change*, **18**(1), 1–24.

Simon, H. A. (1962), 'The architecture of complexity,' *Proceedings of the American Philosophical Society*, **106**(6), 467–482.

Strassmann, W. P. (1963), 'Technological change in the machine tool industry, 1840–1910: discussion,' *The Journal of Economic History*, **23**(4), 444–446.

Von Hippel, E. (1986), 'Lead users: a source of novel product concepts,' *Management Science*, **32**(7), 791–805.