

Adoption of New Information and Communications Technologies in the Workplace Today

Timothy Bresnahan, *Stanford University and NBER*

Pai-Ling Yin, *Stanford University*

Executive Summary

The invention of new applications based on information and communications technologies (ICTs) has had two economic effects up to now. These applications have transformed production, creating value for applications-inventing companies and their customers and increasing economic growth through quality improvements. The same applications have shifted the relative demand for different kinds of labor, raising the demand for already highly compensated managers and professionals relative to other workers. This paper considers the likely impact of new ICT technologies coming into application in the workplace today in light of the economic and technical forces behind ICT application up to now.

I. Introduction

For decades, new information and communications technologies (ICTs) have been coming into the workplace. The application of ICTs represents most of modern technical progress in the services sector (most of modern employment) and in the management and marketing functions of the rest of the economy (most of the rest of modern employment). Today, a new wave of novel ICTs are moving into the workplace, both replacing and complementing existing technologies. This paper takes up two simple to state, if hard to answer, questions: What valuable impacts will come from application of these new ICTs, and what impacts will these have on labor demand, and therefore on income inequality, going forward? Today there is a large academic and public policy debate

about these questions, based on an enormous amount of speculation and uninformed by any examination of the actual application of these new ICTs in the workplace. In this paper, we examine that actual application, expose flawed analysis in the current debate, and emphasize the most likely impact of ICT adoption in the workplace today.

In the language of engineering, ICTs are “enabling technologies”: they permit, but do not direct, the invention of applications. Instead, a complex process of “coinvention” (in the language of economics) creates applications that address market demands and organizational supply processes. ICT coinvention is defined as the product and process improvements created by industries as they apply new ICT. One driver of coinvention is the ICT advances themselves (supply), such as cheaper storage, faster networks, or more capable software. ICT advances produce a large scope of feasible opportunities. The other driver is the industry circumstances (demand) of firms trying to use ICT: competition, customer demand, and the production processes already in place. The industry circumstances define what is profitable for firms. These firms define the direction of technical change by narrowing the scope of feasible opportunities to only those that serve their strategic interest. These applications define the economic value and impacts on labor demand of the enabling technologies. As a result, the bulk of ICT innovation is generated *outside* the tech sector, by firms that are users of ICT. For example, the Internet is a wonderful group of inventions. It permitted, but did not direct, valuable inventions in e-commerce such as online stores. The value-creation and labor-demand implications of online stores depend mostly on the economics of stores (online or not) and only indirectly on the fact that the Internet uses computers and communications equipment.

Three aspects of ICT coinvention warrant emphasis in light of the current concerns over the impact of ICT progress on the workplace:

1. ICT coinvention is very difficult, requiring both brainpower and experimentation to create value. Simply seeing the overlaps between what is newly feasible and what is desirable is very difficult and involves a great deal of invention. ICT coinvention is not just a matter of buying servers, PCs, and phones. Instead, it involves the incremental improvement of applications at a particular using firm. The applications grow steadily more complex and valuable because they build upon the learning that results from their earliest, simple, variants.

2. To the extent that ICT coinvention changes production processes, it changes them at an organization-wide level, not at the level of the

individual worker. ICT coinvention often changes communication and/or incentives throughout an organization or even an entire supply chain rather than just automating individual tasks. A familiar example is Uber, which has a different form of communication between driver and passenger than hailing or calling for a cab and changes incentives by requiring drivers to rate passengers and passengers to rate drivers. Less familiar corporate applications often change workers', suppliers', or customers' incentives and communication.

3. ICT coinvention often increases product quality rather than just simply reducing the costs of production. The automatic teller machine (ATM) network is a familiar example. ATMs did not merely replace human tellers. They allowed consumers to make deposits and (especially) withdrawals when the bank was closed. Once ATMs were networked, they allowed cash access even while far away from any branch of the bank. These increases in convenience increased the quality of banking services. Figure 1 summarizes the ICT coinvention process and implications.

As new waves of ICT have come into the workplace, from mainframes to mobile phones, they have enabled new rounds of ICT coinvention. In this paper, we examine the round of coinvention that is just now beginning. A large number of firms are engaged in ICT coinven-

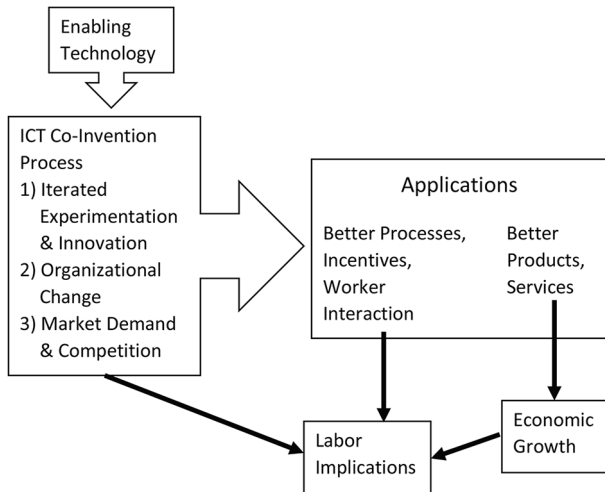


Fig. 1. ICT coinvention of applications involves sustained experimentation to produce organizational change and quality improvements that have economic impact with labor implications.

tion today, drawing on such new technologies as big data, analytics, mobile, and the cloud. What is striking in our investigations is how much coinvention today looks like earlier rounds of coinvention. In short, while there has been terrific technical progress in ICT, there has been little change in the ICT coinvention process. Coinvention still requires considerable brainpower and experimentation. Coinvention still looks for ways to change whole organizations. Indeed, modern coinvention often looks for ways to change whole supply chains. That is still very difficult. And coinvention still seeks to improve product quality in a way that aligns the new applications with the firm's strategic goals. That, too, is still difficult. The main features of coinvention over the last 50 years are still present today.

The three key features of ICT coinvention have important implications for value creation and for labor demand:

1. The need for brainpower and experimentation to translate ICT into new and valuable products and services provides a very simple explanation of why ICT coinvention has raised the demand for smart managers and professionals. Despite all the advances in artificial intelligence (AI), there does not yet seem to be a program that can figure out how to change the incentives of a large number of workers, customers, and suppliers in order to make a new organization that can supply improved product quality as defined by market demand.
2. The organizational change involved in ICT co-invention increases demand for workers with "organizational participation skills (OPS)" in all wage brackets.
3. The complexity of ICT coinvention renders it a long and sustained process, so these labor demand implications will sustain many of the important disparities in the labor market over the last 50 years (changes that many observers have found alarming).

These implications of ICT expose the overly simplistic and thus misguided view among scholars and other observers that improvements in AI will accelerate the substitution of computer intelligence for human intelligence.¹ Less extreme arguments predict that computer intelligence will only replace some tasks of highly skilled work (e.g., Remus and Levy 2015) or that substitution will occur, but it will take a long time (e.g., Autor 2015). We think this debate simply misses the point about the future of ICT-based applications in the workplace by incorrectly focusing on the potential replacement of highly skilled, highly paid intel-

lectual work by ICT in the future. First, the complexity of the coinvention process sows the seeds of its own sustained labor demand. There will continue to be increased demand for managers and professionals whose skills are complementary to ICT coinvention. Second, coinvention instead leads toward a continued expansion of variation in wages. Incomes of the less educated, less established, less successful workers are rising very slowly or falling relative to their ICT-complementing counterparts.² Both the ongoing economic growth and the increase in income inequality that comes from ICT coinvention will, in short, likely continue into the future.

II. ICT Coinvention: Lessons Thus Far

Today, we are familiar with the ideas behind “click here to check the status of your order” or an e-mail that says “please check in for your flight tomorrow.” The first offers information to the customer about the transaction. The second gathers information about the transaction from the customer. Each of these very simple ICT applications is an example of a broad class of applications that add value in parallel ways. We summarize these two classes as “product quality improvements” and “organizational change.” These improvements engender a chain of repeated coinvention, leading to a long and sustained process of experimentation to extract value from ICT.

A. *Product Quality Improvements*

Adding information to transactions can increase product quality and set off a chain of coinventions. We start with a simple example in which a buyer can check the status of an order. Information about the status of an order is a *product quality improvement*. We use a simple consumer example because it is easy to understand. The consumer is better off knowing when the ordered goods will arrive. If the goods will arrive after the desired date of consumption it may be desirable to cancel this order, for example. Adding the information to the transaction, in the language of the economics of technology, increases product quality.

That product quality improvement can be particularly important for transactions that are farther away or between less connected transactors. The ability to provide information about transactions themselves enables transactions between more distant or less familiar buyers and sellers.

A series of famous consumer banking innovations associated with demand deposits illustrate the product quality point and its ongoing importance. All make it more convenient for the customer. The first was the ATM, which let the customer withdraw, check balances, or deposit even when the bank was closed. The next was the ATM network, which let the customer do those activities even if away from any of her bank's branches. Both of these increase convenience.

The increase in convenience is a product quality improvement, and should expand output. As Bessen (2016) has pointed out, the ATM also represents automation of a function formerly performed by human tellers if the customer uses an ATM for deposits or withdrawals that otherwise would have been done in the bank. Some replacement of human work, and some expansion of output through the convenience/product quality improvements. But as Bessen also points out, bank teller demand did not fall. What happened is that tellers moved to more valuable services, like selling other products to the customers. While the example may be extreme, it points out an important feature of product quality improvement: expanding the demand for labor while potentially changing the nature of the ideal worker.

Of course, as ICT-based production has spread, some tasks formerly undertaken by people are indeed now done by machines. This *limited substitution* of capital for labor is most pronounced in routine work in white-collar bureaucracies.³ Few large firms today employ people to look up an account balance in a paper register and write bills. Instead, the system generates the bill automatically. The demand for billing clerks, at least billing clerks who do no more than process bills, is reduced. And, of course, this is not just the billing clerks, but the clerks in a number of different functions in white-collar bureaucracies.

What we have learned from the literature on coinvention, however, is that the business value from automation of such routine white-collar work functions is not the cost savings from removing the (few, and comparatively cheap) humans from the production process. Instead, the automation opens opportunities for improving the process itself to create better-quality products and services. For example, once there is a billing database, smart billing can produce information for a "decision support" system that improves quality.⁴ Getting from a simple transactional system to a decision support system requires managers and professionals to continue the process of coinvention.

This distinction between two kinds of invention is central to the economics of technology. Process innovations permit production of the

same output at lower cost. Product innovations permit production of a new and better kind of output. Process innovations that substitute ICT for labor have not been as important as product innovations that raise the demand for labor. Process innovations that substitute ICT for labor in very particular tasks—the assumed type of innovation in much of the debate about thinking machines—are rare.

In the business-to-business context, there has been large scope for increased product quality improvements as a result of ICT coinvention. To see this, consider ICT applications with the same functionality as “check the status of your order.” This information can be valuable to a business buyer for a number of reasons. Intermediate-input transactions can be large, and intermediate inputs can be complementary to other schedulable inputs, so that better information permits better cost minimization. Of course, an ICT application is systematic, so that after a time buyers have accumulated a database of past order status communications from sellers. The creation of such a database enabled the buyer to coinvent analytical ICT applications. An analytical application that provides answers to such questions as, “Which suppliers deliver fastest?” “Which suppliers deliver at the promised time?” or “Which suppliers inform us when there is a problem?”

Such analytical applications, in turn, enable the invention of decision support applications. To continue on the simple order-status example, a buyer decision support application might bring together information about different sellers to support the decision about from which one to buy. Seller’s behavior in delivering on time (and so on) is entered into a calculation that advises the employee who chooses sellers. Through this path from information about the transaction, analytical applications, and decision support, much value has been created.

Of course, valuable information about transactions, analytical applications, and decision support applications are much broader than the “status of your order” example. Just as buyers want to know whether the goods will arrive on time, sellers want to know whether a particular buyer will pay, will pay in full, will pay on time, and so on.

At this point, we can pause and make an important observation about value creation from ICT coinvention. Trust can arise in markets if buyers know which sellers deliver on time and sellers know which buyers pay on time. Improved information enables the invention of ICT-based applications such as electronic commerce. Many economists have noted the value of repeated play in creating value in dynamic relationships. Our point here is that repeated play needs an information system.

The critical enabling technologies for each round of ICT coinvention are not only the new technologies from the tech sector, but the systems that were built in previous rounds of coinvention. What can be coinvented? This is partially influenced by the supply of new technologies from the tech sector, partially by the value that can be created in the using firm's environment, and partially by the existing stock of applications and data assets already in place. Thinking that the nature of the new technologies will define the subsequent applications and how they will affect work is a mistake of technological determinism. Significant portions of coinvention have very little to do with technology—the portions about the firm's strategic goals, the market's needs, and so on. The task of coinvention has been rapidly increasing the demand for the well-paid workers who understand those issues, and it will continue to do so—even more so as we add valuable new technologies to the list, of which AI and machine learning form a small but not trivial part. The mistake of technological determinism ignores both the historical experience ICT coinvention process and ignores the most important actors that determine the direction of technical change, value creation, and labor market change: the firms that apply ICT.

B. Organizational Change

Sometimes ICT innovation changes production processes. When it does, the minimum divisible unit at which we can analyze technical change is at the level of the *organization*. ICT coinvention involves brainpower and experimentation at the organizational level.⁵ The literature on coinvention emphasizes the cognitive difficulty of inventing new business computer systems, and their complementarity with changes in organization and new services at the firm level.⁶ Typically, the coinvention process is directed by managers (decision makers with control-authority at various levels of leadership in a firm) because of its high potential value, that is, changing the nature of the organization or improving product/service quality. Managers and professionals do more research as a result, and turn their results into operations more systematically. This calls for new cognitive skills, having a deep understanding of one's own organization and one's customers' needs. It calls for brainpower in finding ways to adapt to change.⁷ Thus, the real driver of spreading income distributions is the increase in the demand for smart coinventors in the form of managers (leaders) and professionals who complement ICT.

The adaption at the organizational level to exploit ICT means that skills related to organizational activity become complementary. One systematic use of ICT is to measure and monitor work and the output of work so as to improve incentives and communications within firms.⁸ This has systematically changed the demand for workplace social skills.⁹ Demand is up for workers who can work effectively in teams, operate within the incentives of the organization, function effectively if their manager is offsite, and so forth. Demand is down for workers who want to do their task and who tend to drift off task if not overseen by a manager. We call these “organization participation skills” (OPS) to reflect the importance of efficient interaction and communication between members of the organization. All of these OPS are demanded as part of a bundle—a worker has to be able to perform a particular task or tasks and engage OPS. Note that workers in any level of the organization that can exploit OPS in conjunction with the ICT coinvention process will experience increased demand for their labor.

Managers, for their part, have new demands for OPS as well in ICT-based production. They need to be able to work with teams, to design formal incentive systems, to get effective work out of people at far off locations, and so on. On top of increased brainpower and increased knowledge (e.g., of customers), managers in ICT-based production often need new social skills. Here, too, we see an important element of bundling of different skills in the same person and integration of those skills with organizational activity.

It is worth highlighting again the divergent implications that arise once one recognizes the ICT coinvention process. In contrast to a focus on the organization, Reich (1992) has had great influence with his suggestion that we focus on the individual worker when we think about the impact of ICT on work. Reich drew attention to “symbolic analysts” or “knowledge workers” who use computers, especially PCs. An individual-level analysis misguides the discussion away from the complex integration of workers in an organization transformed by ICT coinvention and toward whether the computer enhances the productivity of the particular worker in a way that leads to declines in employment (substitution) or to expansions in employment by raising the marginal value product of the worker. In short, the analysis mistakenly focuses on the (net) substitutability or complementarity between the computer and the worker who uses it, rather than on the wider implications of organizational changes resulting from ICT coinvention. This misdirection has led almost all of the labor economics literature on skill-biased

technical change to measure “technical change” as “computer use by the individual.”¹⁰

While substitution of computer for human work in specific tasks does occur, the weakness of individual computer use as a metric for technical change is most prominently exemplified by computer programmers: “Now we don’t have to program computers, they program themselves” (Domingos 2015). Tasks that were once done by computer programmers are now done by software.¹¹ As a result, almost all of the engineering work in programming today’s computer, network, Web, and mobile applications, whether in the workplace or not, is done by computers. This substitution has not led to a crash in the demand for computer experts, indeed, we have had the reverse!

Programmer-productivity software is, like scientific and engineering applications generally, significantly easier to invent than other kinds of applications used in the workplace. There is little problem of communication between inventor or customer—both are programmers—about what to invent or how to use it. This difference leads to a systematic effect in the direction of technical progress: scientific and engineering software advances more rapidly than organizational applications do (Bresnahan and Greenstein 1996). Scientific and engineering software, however, is used only by a small subset of the labor force, while organizational software affects a far broader portion of work.

C. A Long and Sustained Process of Experimentation

This increase in the demand for managers and professionals flowing through coinvention is ongoing and perennial. One might have thought that coinvention is a one-time thing, and that firms needed a one-time burst of brainpower to figure it out. On the contrary, two different forces have worked to create a sustained, increasing shift in demand for smart managers and professionals over decades.

One force is the nature of individual “projects” of coinvention in individual firms. Much of coinvention is experimental, spread out over a large number of rounds of improvements, and thus ongoing. Experimentation and improvement in the business aspects of a particular ICT-based system can go on for years or even decades.¹²

Another important force is that new opportunities for coinvention have been created by invention, especially by the improvements in ICT brought into the workplace. In figure 2, we show a number of technol-

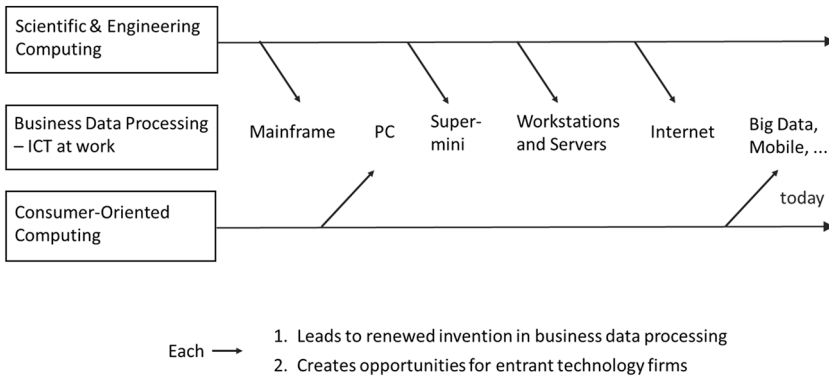


Fig. 2. Cycles of coinvention

ogies that were originally developed outside the workplace, most but not all in scientific/engineering computing, which have migrated to use in the workplace. Each wave of migration creates new opportunities for ICT coinvention in the workplace (as does the ongoing improvement in ICT, though typically not in such dramatic fashion). There is, in short, a pattern of renewal of enabling technologies setting off new rounds of ICT coinvention.

As these cycles of renewal have brought a wider range of ICT technologies, what has been enabled is not just new applications, but also the possibility of larger and more complex applications. Early applications were typically at the level of the firm. Over time, applications have expanded beyond firm boundaries to the level of the firm-and-its-customers, the-firm-and-its-suppliers, whole markets, or whole supply chains.

The analytical literature on coinvention suggests that the complexity brought by increasing potential scope of the application (firm-to-market-to-supply chains) will slow the advance of the most complex apps,¹³ leading to a renewed cycle of experimentation in the ICT coinvention, raising the demand for skilled engineers, and especially for managers and professionals. ICT coinvention has also grown more difficult as it has spread out to organizations such as those in health care, where the problems of changing decision rights in the organization are daunting.¹⁴

Of course, part of the punchline of figure 2 is that this cycle has continued into the present.

III. CT Coinvention Today

We turn now to the ICT technologies that are coming into use in the workplace today. Many of these, such as mobile apps and devices, cloud computing, big data, and analytics, have been widely used in consumer-facing computing, but are still new enabling technologies in the workplace. While these technologies are new to the workplace, new technologies have come into workplace computing before. Our question in this section will be whether there is a radical transformation underway or if these new technologies are setting off a new round of ICT coinvention like the earlier waves of technology shown in figure 2. We will conclude that the current round shares the three key features of early waves. The need for brainpower and experimentation to (co) invent new applications is slowing and directing the coinvention process. The locus of change is the organization, not the individual worker. And product/service quality improvements are central. This is a story of taking up another important opportunity, not a story of radical change.

We start with the World Economic Forum's survey of multinational enterprises, "The Future of Jobs Report" (Schwab 2016). We start here for two reasons. First, the WEF report concludes that we are living in a time of radical change in the workplace. Second, the actual results of the survey are broadly similar to those of other contemporary primary research.

The results of interest to us are reported in table 1. Briefly, the survey asked HR or strategy executives to respond to a list of potential areas of coinvention, called "drivers of change," ranking each area as a "top trend" or not and specifying to time frame in which it will be relevant to application. While the WEF has a longer list, we restrict attention to the subset of "drivers of change" that have an ICT basis. Following the WEF, we rank these by the percent of respondents who feel that the particular the area is a "top trend."

The first thing to note about this table is that the two technologies that are getting the most discussion—robotics and AI—are far down the list ordered by how many users think of the technology as a top trend, with 9% and 6% of users agreeing to the "top trend" designation, respectively. These two technologies are significantly less likely to get that designation than the technology-based reorganization of work, mobile Internet and cloud, or big data, all rated as "top trends" by between one-quarter and one-half of respondents. Second, those same

Table 1

Results from the World Economic Forum Survey of Multinational Enterprises

| “Driver” ^a | Percent Rating as “Top Trend” | Timing |
|---|----------------------------------|-----------|
| Changing work environments and flexible working arrangements ^b | 44 | Already |
| Mobile Internet and cloud technology ^c | 34 | 2015–2017 |
| Advances in computing power and big data ^d | 26 | 2015–2017 |
| The Internet of Things ^e | 14 | 2015–2017 |
| Crowdsourcing, the sharing economy, and peer-to-peer platforms ^f | 12 | Already |
| Advanced robotics and autonomous transport ^g | 9 | 2018–2020 |
| Artificial intelligence and machine learning ^h | 7 | 2018–2020 |
| Advanced manufacturing and 3D printing ⁱ | 6 | 2015–2017 |

^aSource: World Economic Forum, “The Future of Jobs Report,” January 2016. This survey was filled in by senior “talent and strategy” executives at 371 companies in the first half of 2015.

^bNew technologies are enabling workplace innovations such as remote working, coworking spaces and teleconferencing. Organizations are likely to have an ever-smaller pool of core full-time employees for fixed functions, backed up by colleagues in other countries and external consultants and contractors for specific projects.

^cThe mobile Internet has applications across business and the public sector, enabling more efficient delivery of services and opportunities to increase workforce productivity. With cloud technology, applications can be delivered with minimal or no local software or processing power, enabling the rapid spread of Internet-based service models.

^dRealizing the full potential of technological advances will require having in place the systems and capabilities to make sense of the unprecedented flood of data these innovations will generate.

^eThe use of remote sensors, communications, and processing power in industrial equipment and everyday objects will unleash an enormous amount of data and the opportunity to see patterns and design systems on a scale never before possible.

^fWith peer-to-peer platforms, companies and individuals can do things that previously required large-scale organizations. In some cases the talent and resources that companies can connect to, through activities such as crowdsourcing, may become more important than the in-house resources they own.

^gAdvanced robots with enhanced senses, dexterity, and intelligence can be more practical than human labor in manufacturing, as well as in a growing number of service jobs, such as cleaning and maintenance. Moreover, it is now possible to create cars, trucks, aircraft, and boats that are completely or partly autonomous, which could revolutionize transportation, if regulations allow, as early as 2020.

^hAdvances in artificial intelligence, machine learning, and natural user interfaces (e.g., voice recognition) are making it possible to automate knowledge-worker tasks that have long been regarded as impossible or impractical for machines to perform.

ⁱA range of technological advances in manufacturing technology promises a new wave of productivity. For example, 3D printing (building objects layer-by-layer from a digital master design file) allows on-demand production, which has far-ranging implications for global supply chains and production networks.

two technologies are placed firmly in the future by these respondents, with the timing of their impact three years off as of the survey.

The many public statements that these new technologies are changing everything contrast with the empirical fact that these technologies are, in round numbers, not changing anything yet.

Second, the same survey reveals that there are *other* important ICT technologies that are having an impact on the workplace or that respondents see as about to have an impact (versus being years off). The important technologies enabling coinvention today are ones that support ongoing organizational change and the new imports from consumer computing—mobile, cloud, and big data.

With this influx of new applications and this increase in the complexity and scope of the applications of ICT has come a remarkable development. Rather than becoming simpler, the invention of the applications of ICT has retained the strategic, organizational, and creative aspects of coinvention first noticed decades ago: (a) innovation and experimentation are required to apply ICT technologies in the workplace, (b) organizational adaptations are required to both create and implement ICT applications in the workplace, and (c) the resulting new services and products from this coinvention process represent simultaneously higher-quality and lower-cost changes to the economy. There has been little maturation of the coinvention process. Instead, the need to change complex organizations, align the new application with firm goals, and have an implementation that specifically achieves what the grand strategy expected remain centerpieces of coinvention. These three properties all enhance and increase demand for human labor.

A. Innovation and Experimentation

A cluster of technologies has emerged associated with big data. “Big” data sets—sometimes called data streams—are typically gathered automatically. Many ordinary consumer activities contribute to these data streams already, such as searching on the Web, looking through an online store before buying things, online social interaction or other communication, mobile app usage, and so on. Big data techniques have been pushed very far forward in the largest consumer-oriented ICT firms. Many of these technologies have so far had their main value-add as specialization recommendation-engine technologies. Google, for example, uses complex algorithms both to decide what search results to show a particular person and what advertisement to show a par-

ticular person after a particular search. Amazon does not recommend the same products to everyone, nor Netflix (nor Pandora) the same entertainment, just as Google returns different search results to people it “decides” have different interests. A great deal of machine learning is involved in making that “decision.” Note, however, that these technologies enhance, rather than replace, human decision making by providing information in a useful manner to the human decision maker.

These consumer-facing applications have led to a large volume of technical progress—the new technologies that are now coming into the workplace. It has proved difficult to manage these data streams with traditional relational data base management software and associated programming tools, both because the scale of “big” data can be quite large and because it is difficult to impose definitions on large, cheaply gathered data streams. (In contrast, traditional DBMS software dealt with transactions data sets, where there can be a known and enforced definition of all the fields associated with a transaction). Not surprisingly, “predictive analytics,” which are statistical techniques with a data-mining flavor, are closely associated with big data. Some of these techniques involve machine learning, while others do not.

The potential application of these techniques in the workplace today is generating tremendous interest among employers. A number of surveys report that large US firms fall into three roughly equally sized groups: those that already have a big data coinvention project under way, those that have identified and are capturing big data streams (but don’t yet have a coinvention project under way), and others (Dresner Advisory Services 2015). Penetration of the big data techniques out of a coinvention project phase into production applications is perhaps a bit slower: in a recent survey, Dresner Advisory Services reported that 83% of 3,000 firms were not currently using big data, despite 59% of the firms responding that application of big data is “critically important” (Maguire 2016). The technical disciplines associated with big data and analytics, such as data scientists, are in very hot demand, and the supply-side response of trying to train more data scientists in statistics, computer science, and economics departments is very large.

The coinvention wave associated with big data involves, like the earlier waves of coinvention, substantial amounts of experimentation, exploration, and the deployment of managers with the brainpower to understand management, statistics, and their customers, workers, or suppliers.¹⁵

The pattern of experimentation is taking time and multiple rounds

to lead to value creation in this round as in earlier rounds of coinvention. This pattern has not changed in modern times. Harnessing the benefits of new technology is hard, so the realization of its benefits is often delayed far beyond the date when its applications are envisioned. This has been observed historically, and the impending technologies of today are no different. David Court wrote in the *McKinsey Quarterly*, January 2015:

[W]e recently assembled a group of analytics leaders from major companies that are quite committed to realizing the potential of big data and advanced analytics. When we asked them what degree of revenue or cost improvement they had achieved through the use of these techniques, three-quarters said it was less than 1 percent. (Court 2015)

Even more simplistic data involved in back-end production processes poses application challenges. The measurement and networked communication opportunities from ubiquitous sensors have created excitement over the possibilities from an Internet of Things (IoT). However, as a June 2015 *McKinsey* report discovered, “Most IoT data are not used currently. For example, only 1 percent of data from an oil rig with 30,000 sensors is examined. The data that are used today are mostly for anomaly detection and control, not optimization and prediction, which provide the greatest value” (Manyika et al. 2015, 4).

The ultimate value of analytics based on big data is clear to these adopters at an early stage. What is not clear is how that ultimate value will be realized. Only after the initial applications have been built, and their potential role in the organization assessed based on evidence, will the value proposition and the labor demand impact become clear. Until then the process of innovation and experimentation to discover those applications will require more human labor, not less.

Big data applications development hits two very different bottlenecks when it moves into more ordinary businesses. It can be difficult to determine the profitable applications for big data. Those organizations that have gathered big data, but are not yet using it, typically report that as the bottleneck. A Dell survey reported that for organizations that have big data but no applications using it, “the top barrier is not knowing if the benefits are worth the costs.” (Dell 2015, slide 4). Typically it is a business person, not an IT person, who leads the effort to define big data projects. Those organizations that are using big data typically identify the costs of the IT infrastructure for big data processing. This is much like the early stages of many important ICT technologies of

the past. Inventing the uses is a bottleneck at the beginning, when the costs (data scientists are expensive as well) make the risk/benefit/cost calculation difficult.

Many big data applications go to places where risks are low because coinvention costs are low. One current example of ICT coinvention is security applications in finance, such as fraud detection in credit cards and related systems. These are low-risk applications for big data because they are already statistical—the output of a fraud-detection system is a probabilistic assessment of a particular transaction—fraud or legit? Thus a big-data-based system need only provide a better prediction of fraud in a statistical sense, and it can be plugged into the existing people- and computer-based system for fraud detection without altering it fundamentally. PayPal, for example, uses big data and machine learning for fraud detection (Knorr 2015). PayPal uses a linear fraud-prediction algorithm, a neural network-based algorithm, and deep machine learning and data mining—alongside human intelligence. Dr. Hui Wang, senior director of risk sciences at PayPal, notes:

I never worry that these machines will replace humans. Yes, we can add layers, but you can talk to any machine learning scientist and they will say that the algorithm is important, but at the end of the day what really makes the difference is that a machine cannot find data automatically. . . . There is so much data, so much variety, but the flip side is: What is useful? We still rely on human oversight to decide what ingredients to pump into the machine. (Knorr 2015)

Another current day example of ICT coinvention is risk analysis in insurance, which was, of course, already analytic, so adding big data analytics is not a radical change. Practitioners confirm that the path of invention in this area is to undertake the low risk, low change part of the application first. Only later do more profound changes follow.¹⁶ As the applications mature, they tend to enhance rather than replace the efforts of workers: insurance industry analytics modelers work with claims staff to design effective models.¹⁷

Related examples can be found throughout the current coinvention wave. Ram Narasimhan of GE emphasized the change from consumer applications—coinvention—when he said “What data scientists do at Google or Yahoo or Facebook is a little different from what we at GE do” (Gage 2016). But he also emphasized the continuity with previous applications, such as detecting and predicting engine failures. These were already statistical; early big data applications permit doing a better job because there is more data about more GE-manufactured en-

gines and “[t]hese engines are becoming smarter, and they generate a lot more data” (Gage 2016). Like fraud detection in finance, failure prediction in operations was already statistical and it is a natural move to go to machine learning on top of big data—statistical analysis that, because of the larger and more detailed data streams, can be faster and more detailed.

This pattern of undertaking experiments where it is comparatively easy to see that they might go well is an element of coinvention that is familiar from earlier rounds of coinvention (see Bresnahan and Greenstein [1996] for much earlier rounds). The next round of coinvention will be based on the early experiments in their specific business context. How that next round will create value will depend on what is learned (by humans) from the experiments, not on some predetermined substitution of ever smarter computers for the next-smarter kind of human worker.

Indeed, like traditional database applications before them, big data applications are not primarily based on substituting new and improved machine intelligence to do something that human intelligence did before. Instead, there is *expansion*. People are using computers to do many things, and machines have more and more embedded computers, so more data streams are feasible to capture cheaply. Cheap computers, cheap communications, and especially cheap storage make it cost effective to make data streams that capture information that it was never cost effective to capture before.

The increase in demand such that there is now a shortage of data scientists and programmers is the most obvious labor market effect of this new coinvention. However, the human element is not limited to those technical specialties to develop the new applications. Instead, there are a large number of managers who decide on the experiments to do in the course of coinventing ICT applications based on big data. This creates opportunities for managers who understand not only their business, but also some statistics and machine learning—a new bundle of skills—to create value.¹⁸ Like earlier rounds of coinvention, this round is creating new valuable bundles of different human skills and rewarding them highly.

This effect of increasing the value of human work need not be limited to managers and professionals, though if the earlier pattern of coinvention’s impact on the workplace is followed, it will be stronger there. Firms that effectively execute strategies today have organizations comprised of integrated departments and workers, so the critical

knowledge and decisions necessary for a successful ICT application are found in consumer-facing staff. Midrange participants in that staff that embrace the opportunity to make decisions based on machine learning and have the OPS to participate effectively in the coinvention of new systems will find the demand for their services raised.

B. Organizational Adaption

What we learned in the review of coinvention above is that organizational adaptations have long been required to both create and implement ICT applications in the workplace. This is also true in the present. The wave of new ICT coinvention set off by today's new technologies has set many companies off on a path of organizational change. As in earlier waves of coinvention, the ICT often augments how people in firms interact with one another or with suppliers and customers, rather than replacing those humans. The replacement of ever-smarter humans by ever-smarter machines does not appear as an important part of this dynamic, much less as its central thrust.

One well-known contemporary example is Uber. Compared to a traditional taxi company, Uber does indeed involve some limited aspects of automation. The task of taking customer orders and dispatching vehicles was traditionally done by (human) dispatchers on the telephone, more recently by order-taking websites plus human dispatchers on the telephone. It is a critical element of this part of Uber, however, that the automatic scheduling, combined with the flexible work schedule of Uber drivers, raises the efficiency of drivers and cars. Ubers wait less time between rides for their next fare. This increase in efficiency raises marginal physical product of an hour of driver time (more of it is spent productively) and is one key to the labor-demand expansion effect of Uber.

Uber also involves a significant change in organizational incentives compared to traditional taxi companies. Uber drivers, unlike cab drivers, are rated by their customers—Uber is set up to make rating nearly universal. This changes the incentive of a driver to be on time, courteous, and safe (not all of which are famous features of cab drivers.)¹⁹ Uber drivers who are rated as weak on these attributes soon leave the service. Similarly, drivers rate passengers on Uber, changing passenger incentives. Passengers who are abusive in the car, are late to come out to the street, and so on, are rated down. Uber also makes it easy for driver and rider to communicate by cellphone to, for example, settle on an exact pickup location. These coordination and incentive terms cut

across rider, platform (Uber), and driver—changing the organizational structure of the market. Einav, Farronato, and Levin (forthcoming) have a very interesting analysis of Uber and other “peer-to-peer” networks from a market-design perspective, emphasizing the changes in the organization of the entire market and its new incentives and information structures.

Bank of America recently adopted wearable technologies for its workers with the intent to monitor and improve productivity. They successfully achieved their goals. However, the way in which these productivity gains were achieved was based on changes in organizational structure to enhance human interaction within the firm, not outside of the firm:

When they began studying the reps they figured the key indicator would be how they talked to customers, but they found that it was actually how they talked to one another that was most important because employees shared information and techniques.

They learned that employees talked most during the 15 minutes their lunches overlapped, so they tried giving one group lunch all at the same time and let the other group continue to have lunch according to the old staggered schedule. What they found surprised them.

Network cohesiveness, which measures how well they communicate went up 18 percent. This reduced stress (as measured by tone of voice) by 19 percent. All of this led to happier employees and lower turnover rates, which went down 28 percent. The key metric though, call completion time improved by 23 percent. These are numbers that on a scale of Bank of America could translate into billions in savings. (Miller 2015)

The process of coinvention requires innovation and experimentation by the entire organization, so it is no surprise that the first step in coinvention is the enhancement of communications within the organizations to spur innovation and experimentation. As coinvention successfully creates workplace applications, the organization must also adapt to implement those applications. Current explanations of the coinvention processes written by leading practitioners, relating to current applications of the newly imported consumer-oriented technologies, could easily have been written when servers were coming into business data processing in the early 1990s. Here is a representative discussion of the difficulties of coinvention in big data and analytics (a close complement) applications:

[C]apturing the potential of data analytics requires the building blocks of any good strategic transformation: it starts with a plan, demands the creation of

new senior-management capacity to really focus on data, and, perhaps most important, addresses the cultural and skill-building challenges needed for the front line (not just the analytics team) to embrace the change. . . . Existing organizational processes are unable to accommodate advancements in analytics and automation, often because protocols for decision making require multiple levels of approval. (Court 2015)

These large organizational changes go beyond just needing to invent new applications software. More critically, the cost of fully integrated technology adoption requires adaption by the entire organizational structure, from incentives, to routines, to culture, to redefining job descriptions. Thus the impact on labor demand is more or less the opposite of the simple task-substitution model. As the practitioner we just cited wrote, “Automating part of the jobs of employees means making a permanent change in their roles and responsibilities” (Court 2015). What is particularly interesting here is that the practitioner is discussing, not the use of big data to have better data, but the use of analytics running on top of big data to make better decisions. It is precisely this substitution, of machine for human decision making, that makes the application the most difficult to implement organizationally.²⁰

D. Improving Product Quality Creates More Demand

The resulting new services and products from this coinvention process represent simultaneously higher-quality and lower-cost changes to the economy. Higher quality implies that demand for these innovations will increase, potentially increasing the demand for complementary human labor. Lower cost may lead to lower prices, which would also lead to higher demand for these innovations, again increasing the demand for complementary human labor.

For example, one aspect of consumer-oriented computing has been rapidly taken up in the corporate sector, the mobile app. Originally introduced as an entrepreneurial technology to replace existing firms and markets (and with a few important successes in that arena out of the many millions of mobile app entrepreneurs), the technology for making and distributing mobile apps was quickly adapted by existing companies.²¹ The corporate mobile app—whether from an airline, a retailer, or Starbucks—quickly took on one of the classical attributes of computing at work, product quality improvement. The implication of technical change that leads to product quality improvement is different from automation. Automation substitutes capital for labor (and the resulting

cost rise may lead to some price falls and thus output expansion). Product quality improvement leads directly to output expansion.

In the case of the corporate mobile app, firms typically adapted the technology to deepen customer relationships. Airlines were able to give travelers more information about their flight, give them a boarding pass that didn't need to be printed, and let them check the upgrade list as often as they liked. Rather than replacing the functions of gate agents and customer service agents, the app lets the traveler interact with the airline more frequently. Starbucks was able to provide a convenient way to store loyalty rewards, which increased customer retention and necessitates more, rather than less, baristas.

ICT innovations are more likely to redefine jobs to higher-quality services and activities at lower costs, as some tasks are automated.²² "Even where automation has made significant progress, its impact has been far less than the headlines would have us believe" (Remus and Levy 2015, 5). In the legal profession, automation of scanning and preparing legal documents means that those services are cheaper and scalable, so demand could increase among previously unserved markets.²³ In the medical profession, a robot pharmacist at UCSF allows the current 100 staff pharmacists to improve services:

[N]ow nearly all have been reassigned to different parts of the hospital, where they make IVs, help adjust patients' drug regimens, and perform other tasks that had been neglected when they were simply filling prescriptions.

There is a theory among pharmacists that robots will ultimately benefit the profession. The more automation that enters the field, the more pharmacists can focus on uniquely human tasks: counseling patients and working with doctors to ensure the proper medicine gets prescribed. This is especially true of pharmacists who work in hospitals. UCSF hasn't laid off any human pharmacists since installing its robotic dispenser. (Manjoo 2011a)

Next was the mobile banking app, which let the customer deposit checks and check balances from anywhere—and which, when coupled with a smartphone payment system (Apple Pay or its many competitors) lets the customer do transactions out of her demand deposit account without cash. At each stage, convenience increases. Of course, the same can be expected of the modern, mobile-based "tellers in your pocket." They increase convenience, substitute some routine human tasks, and create opportunities to move the human worker to higher-value tasks.

Other corporate mobile apps have similar flavors. "Multichannel marketing" apps from retailers permit the customer to buy on her mo-

mobile device, as well as online or in the store. This has some elements of automation—a purchase made online or on the mobile device does not require an in-store salesperson. Yet it also has elements of complementarity. If the purchases suggested online or on the mobile app take advantage of the information gleaned about the customer from an in-store visit, the efforts of the store sales rep to learn the customer's size or tastes are thereby spread out over more transactions than the salesperson personally attends. That leverage raises the demand for particularly effective salespeople.

By far the most important lesson of the product-quality point is that ICT-based production processes are not merely automated versions of the previous processes. Just as with the organizational-change point, the mechanism by which ICT-based production processes have saved more low-wage labor than high-wage labor is not one of automation.

E. Labor Market Outcomes

One implication of a risky and uncertain process of ICT coinvention is that some firms should be more successful than others. The more successful firms should not just be more profitable, but should show the major indicia of coinvention, such as changes in skill demand and higher ICT use. This offers an opportunity to test the implications of our model of ICT coinvention's impact on the labor market by looking at labor market studies that are carried out at the level of the organization/firm-level studies. By now, there have been a large number of these studies.²⁴ We discuss a couple of the newer ones here.

Bessen (2016) has provided valuable new information on the implications of computer use in organizations. He predicts employment growth from 1980 to 2013 in an occupation in an industry (accountants in insurance). He includes two predictors relevant to our enquiry: (1) to what extent do people in this occupation in this industry report that they use a computer at work, and (2) to what extent do people in other occupations in this industry report that they use a computer at work? Use of computers within an occupation in an industry predicts *higher* employment growth in that occupation/industry. It is use of computers in *other* occupations in the industry that predicts lower employment growth. Since the "computer use" variables are the only ICT measures included, this points to an organizational, rather than a task-substitution, model of computer-based productivity change.²⁵ More interesting is the pattern of impacts of organizational computing that this suggests. There is

little evidence that either own-computer-use or others'-computer-use impacts employment growth in routine-intensive occupations, an unhappy fact for the computer-human substitution hypothesis. However, when Bessen looks across the wage distribution, he finds that others'-computer-use lowers employment growth in the bottom three quartiles of the wage distribution significantly, but has a positive (if imprecisely estimated) impact on the top wage quartile. This is consistent with the idea that organizational changes in computing allow better or smarter management of the overall organization, raising the productivity of comparatively rich workers.

Another implication is that firms should be spreading out in the wages they pay; successful coinventing firms expand, typically because of the increase in product quality, and demand more labor, especially more capable labor. This implication, too, is borne out in the data. Song et al. (2015) examine the sources of the substantial increase in earnings dispersion across workers in the United States from 1978 to 2012. They find that "virtually all" of the increase in earnings dispersion between workers is accounted for by an increase in the dispersion across employers (firms). In a study looking at the same phenomenon, but on a narrower range of firms (only manufacturing firms), Dunne et al. (2004) find that "the bulk of overall wage dispersion is accounted for by between-plant dispersion, and the contribution of this component has been growing over time." An advantage of the narrower frame of the Dunne et al. (2004) work is that they have ICT variables, and find that the increase in wage dispersion across plants is highly correlated with ICT adoption.

The right level of analysis is not the individual job, but the entire organization, including today parts of the organization that extend beyond firm boundaries to customers or suppliers. The impact of ICT adoption and ICT coinvention at the firm level has been to raise the demand for workers and their wages. The story that ICT is substituting for smarter and smarter human jobs as computers grow smarter is just wrong. ICT capital-intensive production replaces a wide range of workers, but is complementary to skilled workers, both in the sense of smart people and people with OPS.

IV. Concluding Remarks

The use of ICT in white-collar organizations and in markets has, over the last 50 years, had two important economic consequences. It has

been valuable, particularly through restructuring organizations and markets and through improving product quality. It has also contributed to income inequality because higher-paid workers are complements to the adoption of ICT in firms. Task-level substitution of capital for labor has been an unimportant force.

Today, firms are importing new ICT technologies such as big data and mobile from the consumer computing world. Our examination of the new wave of ICT coinvention suggests that the important economic consequences of current and medium-term technical progress will be largely the same as over the past 50 years. There will be continued value creation and changes in organizations and markets. There will continue to be increased shifts in labor demand toward highly compensated workers, especially managers and professionals, whose work is complementary to computer- and network-based production.

Technological developments have led some to predict a sea change in the world of work as machines are substituted for humans wholesale, and others to predict a continuation of the (supposed) past trend of substituting ever-smarter machines for a wider range of human tasks. These perspectives fail to account for the basic economics of ICT applications coinvention, resting instead on an incorrect assumption of technological determinism and an incorrect assumption that substitution of computer for human at the task level is an important force.

The radical forecasts embody a strong form of technological determinism that is not appropriate for analysis of the adoption of ICT in the workplace. The claim about new ICT technology, that “this changes everything,” typically denotes an important technology whose eventual applications have not yet been invented. The point is serious, not sassy. New enabling technologies do not set the direction of technical change. Technical change is defined by applications of the technology that respond to market demands. Furthermore, the applications are hardly limited to the automation possibilities of the new technology, but instead are diverse and span a wide range of opportunities to change organizations and markets. The vast bulk of ICT innovation occurs in the many industries outside the tech sector attempting to harness the power of new ICT. That innovation is influenced not only by cutting-edge developments in computer science but also (and much more) by the business, market, organizational, and strategic concerns of those firms. Any particular advances in ICT—including “smart machines,” “robots,” or AI, will itself be transformed by coinvention as it is brought into productive use as much as it transforms production. The

coinvention process is consequently renewed, sparking another round of complementary labor demand.

Our examination of the facts surrounding ICT coinvention in the workplace today highlights the complementarity in quality improvements, organizational change, and technical and managerial brainpower. The spreading income distribution arises from complementarity between ICT and highly remunerated workers, not from the substitution for human work on a task level. The thing to like about ICT—valuable technical progress—and the thing to worry about—spreading out of the income distribution—will be the same in the near future as in the recent past.

Endnotes

We gratefully acknowledge the research assistance of Navya Maitri Konda, Jessica Ou, and Nafia Chowdhury. For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see <http://www.nber.org/chapters/c13762.ack>.

1. Consider headlines such as, “Why the Final Game between AlphaGo [an AI machine] and Lee Sedol [a human grandmaster in the game of Go] Is Such a Big Deal for Humanity” (Metz 2016), Klaus Schwab’s argument at the World Economic Forum that we are at the start of the “fourth industrial revolution” and that “[t]here has never been a time of greater promise, or one of greater potential peril” (Schwab 2016), or the perception of the driverless car as the sign that there is about to be an acceleration of substitution of computer intelligence for human intelligence.

2. Compare Council of Economic Advisers, chapter 5, “Technology and Innovation,” *Economic Report of the President*, February 2016.

3. See Bresnahan (1999) for details.

4. Keen (1981) lists 11 different kinds of decision support systems coming into widespread use decades ago. Rather than reducing cost, he writes, they tend to add value. This is typical of assessments of DSS over a long period.

5. A summary of the management implications of the organizational focus can be found in Hammer (1990).

6. See Barras (1990), Bresnahan and Saloner (1996), Bresnahan and Greenstein (1996), Brynjolfsson and Hitt (1996), and Friedman and Cornford (1989).

7. Bartel and Lichtenberg (1987) suggest that high levels of cognitive skills may be particularly important in creating and adapting to change, notably in implementing new technology. The managerial side of computer-based production processes is an excellent example of this story.

8. See, for example, Hubbard (2003), Athey and Stern (2002), or Bartel, Ichniowski, and Shaw (2007).

9. Social skills are very important in labor markets, though they are badly measured by education, occupation, or income, the three most common skill proxies in the empirical labor economics literature.

10. Card and DiNardo (2002) point out that an organizational-computing perspective exists, but draw no implications from this.

11. An impressive list of technologies, from the multiprogramming computer and the RDBMS to today’s frameworks and APIs, has contributed to this. Friedman and Cornford (1989) made the useful observation that moving much of the engineering work of ICT coinvention from human to machine had changed the “bottleneck” in computer systems

invention from managing the hardware to managing the software to deciding what the system should do, that is, to coinvention.

12. The airline reservation system is a famous example.

13. This appears to be entirely correct (see McElheran 2015).

14. Dranove et al. (2014) discusses the difficult but ultimately successful efforts to bring Electronic Medical Records to the hospital sector.

15. Susan Athey has suggested a set of implications for the education of managers rooted in a deep understanding of the possibilities for new work in the “big data” era (O’Toole 2013).

16. “In the initial phase, carriers develop models that demonstrate early evidence of success” (Chester, Clarke, and Libarikian 2016).

17. “As the analytics function matures, model builders work closely with frontline staff, who become involved in the nuts and bolts of building the model” (Chester, Clarke, and Libarikian 2016).

18. Athey describes managers that have this bundle of skills as “rock stars” on the labor market (O’Toole 2013).

19. Wallsten (2015) undertakes an empirical investigation of taxicab complaints and concludes that *competition* from Uber has a positive impact on cab driver behavior.

20. “For example, it’s great to have real-time data and automated pricing engines, but if management processes are designed to set prices on a weekly basis, the organization won’t be able to realize the full impact of these new technologies. . . . If you automate pricing, for instance, it is hard to hold the affected manager solely responsible for the profit and loss of the business going forward, since a key part of the profit formula is now made by a machine” (Court 2015).

21. A much fuller account of the origins of the corporate mobile app can be found in Bresnahan, Davis, and Yin (2015).

22. “According to our analysis, fewer than 5 percent of occupations can be entirely automated using current technology. However, about 60 percent of occupations could have 30 percent or more of their constituent activities automated” (Chui, Manyika, and Miremadi 2015).

23. “Instead of serving just a handful of high-paying clients, this maestro might be able to use machines to help serve thousands of clients over the Web, providing legal help to those who can’t access it today” (Manjoo 2011b).

24. The complementarity between ICT adoption and organizational and skills change can be found in, for example, Bresnahan and Greenstein (1996), Bresnahan, Brynjolfsson, and Hitt (2002), Forman, Goldfarb, and Greenstein (2012), and Song et al. (2015). All of these papers document the variety across firms; indeed, that variety is their measurement frame.

25. It also shows that the long literature based on individuals’ computer use at work is irrelevant to the assessment of the labor-demand impact of ICT-based production. Both the part of the literature that finds impacts of computer use on wages (starting with Krueger [1993]) and the part of the literature that casts doubt on that finding (DiNardo and Pischke 1997 “pencils”) are irrelevant to assessing the labor-demand impact of ICT in the workplace.

References

- Athey, Susan, and Scott Stern. 2002. “The Impact of Information Technology on Emergency Health Care Outcomes.” *RAND Journal of Economics* 33 (3): 399–432.
- Autor, David H. 2015. “Why Are There Still So Many Jobs? The History and Future of Workplace Automation.” *Journal of Economic Perspectives* 29 (3): 3–30.
- Barras, Richard. 1990. “Interactive Innovation in Financial and Business Services.” *Research Policy* 19 (3): 215–37.

- Bartel, Ann, Casey Ichniowski, and Kathryn Shaw. 2007. "How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills." *Quarterly Journal of Economics* 122 (4): 1721–58.
- Bartel, Ann P., and Frank Lichtenberg. 1987. "The Comparative Advantage of Educated Workers in Implementing New Technology." *Review of Economics and Statistics* 69 (1): 1–11.
- Bessen, James E. 2016. "How Computer Automation Affects Occupations: Technology, Jobs, and Skills." Law and Economics Research Paper no. 15-49, Boston University School of Law, (January). <http://dx.doi.org/10.2139/ssrn.2690435>.
- Bresnahan, Timothy F. 1999. "Computerization and Wage Dispersion: An Analytic Reinterpretation." *Economic Journal* 109 (7): 390–415.
- Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." *Quarterly Journal of Economics* 117 (1): 339–76.
- Bresnahan, Timothy F., Jason P. Davis, and Pai-Ling Yin. 2015. "Economic Value Creation in Mobile Applications." In *The Changing Frontier: Rethinking Science and Innovation Policy*, edited by Adam Jaffe and Benjamin Jones. Chicago: University of Chicago Press.
- Bresnahan, Timothy F., and Shane Greenstein. 1996. "Technical Progress and Co-Invention in Computing and in the Uses of Computers." *Brookings Papers on Economic Activity, Microeconomics* 1996:1–83.
- Bresnahan, Timothy, and Garth Saloner. 1996. "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*, edited by David B. Yoffie. Cambridge, MA: Harvard University Press.
- Brynjolfsson, Erik, and Lorin Hitt. 1996. "Paradox Lost? Firm-Level Evidence on the Returns to Systems Spending." *Management Science* 42 (4): 541–58.
- Card, David, and John E. DiNardo. 2002. "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles." *Journal of Labor Economics* 20 (4): 733–83.
- Chester, Ari, Richard Clarke, and Ari Libarikian. 2016. "Transforming into an Analytics-Driven Insurance Carrier." *McKinsey Quarterly*, February. <http://www.mckinsey.com/industries/financial-services/our-insights/transforming-into-an-analytics-driven-insurance-carrier>.
- Chui, Michael, James Manyika, and Mehdi Miremadi. 2015. "Four Fundamentals of Workplace Automation." *McKinsey Quarterly*, November. <http://www.mckinsey.com/business-functions/business-technology/our-insights/four-fundamentals-of-workplace-automation>.
- Council of Economic Advisers. 2016. "Technology and Innovation." *Economic Report of the President*, chapter 5, February. <https://www.whitehouse.gov/administration/eop/cea/economic-report-of-the-President/2016>.
- Court, David. 2015. "Getting Big Impact from Big Data." *McKinsey Quarterly*, January. <http://www.mckinsey.com/business-functions/business-technology/our-insights/getting-big-impact-from-big-data>.
- Dell. 2015. Global Technology Adoption Index, Slide 4. https://powermore.dell.com/wp-content/uploads/2015/10/GTAI_2015_Results_Deck_Final_for_Web_101115.pdf.
- DiNardo, John E., and Jörn-Steffen Pischke. 1997. "The Returns to Computer

- Use Revisited: Have Pencils Changed the Wage Structure Too?" *Quarterly Journal of Economics* 112 (1): 291–303.
- Domingos, Peter. 2015. *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. New York: Basic Books.
- Dranove, David, Chris Forman, Avi Goldfarb, and Shane Greenstein. 2014. "The Trillion Dollar Conundrum: Complementarities and Health Information Technology." *American Economic Journal: Economic Policy* 6 (4): 239–70.
- Dresner Advisory Services. 2015. Big Data Analytics Market Study, November. <http://www.bigdataanalyticsreport.com/>.
- Dunne, Timothy, Lucia Foster, John Haltiwanger, and Kenneth R. Troske. 2004. "Wage and Productivity Dispersion in United States Manufacturing: The Role of Computer Investment." *Journal of Labor Economics* 22 (2): 397–430.
- Einav, Liran, Chiara Farronato, and Jonathan Levin. Forthcoming. "Peer-to-Peer Markets." *Annual Review of Economics* 8.
- Forman, Chris, Avi Goldfarb, and Shane Greenstein. 2012. "The Internet and Local Wages: A Puzzle." *American Economic Review* 102 (1): 556–75.
- Friedman, Andrew L., and Dominic Cornford. 1989. *Computer Systems Development: History, Organization, and Implementation*. New York: Wiley.
- Gage, Deborah. 2016. "What Data Scientists Do All Day at Work: Ram Narasimhan of GE Talks about the Importance of Curiosity and What Makes His Day." *Wall Street Journal*, March 13.
- Hammer, M. 1990. "Reengineering Work: Don't Automate, Obliterate." *Harvard Business Review* 68 (4): 104–12.
- Hubbard, Thomas N. 2003. "Information, Decisions, and Productivity: On Board Computers and Capacity Utilization in Trucking." *American Economic Review* 93 (4): 1328–53.
- Keen, Peter G. W. 1981. "Value Analysis: Justifying Decision Support Systems." *MIS Quarterly* 5 (1): 1–15.
- Knorr, Eric. 2015. "How PayPal Beats the Bad Guys with Machine Learning." *InfoWorld*, April 13. <http://www.infoworld.com/article/2907877/machine-learning/how-paypal-reduces-fraud-with-machine-learning.html>.
- Krueger, Alan B. 1993. "How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984–1989." *Quarterly Journal of Economics* 108 (1): 33–60.
- Maguire, James. 2016. "The Surprising Truth about Big Data." February 2. <http://www.datamation.com/applications/thesurprisingtruthaboutbigdata.html>.
- Manjoo, Farhad. 2011a. "Will Robots Steal Your Job?" *Slate*, September 26. http://www.slate.com/articles/technology/robot_invasion/2011/09/will_robots_steal_your_job_2.html.
- . 2011b. "Will Robots Steal Your Job?" *Slate*, September 29. http://www.slate.com/articles/technology/robot_invasion/2011/09/will_robots_steal_your_job_5.html.
- Manyika, James, Michael Chui, Peter Bisson, Jonathan Woetzel, Richard Dobbs, Jacques Bughin, and Dan Aharon. 2015. "The Internet of Things: Mapping the Value beyond the Hype." *McKinsey Quarterly*, June. <http://www.mckinsey.com/business-functions/business-technology/our-insights/the-internet-of-things-the-value-of-digitizing-the-physical-world>.
- McElheran, Kristina. 2015. "Do Market Leaders Lead in Business Process Innovation? The Case(s) of E-Business Adoption." *Management Science* 61 (6): 1197–216.
- Metz, Cade. 2016. "Why the Final Game between AlphaGo and Lee Sedol Is

- Such a Big Deal for Humanity." *Wired* (online), March 14. <http://www.wired.com/2016/03/final-game-alphago-lee-sedol-big-deal-humanity/>.
- Miller, Ron. 2015. "New Firm Combines Wearables and Data to Improve Decision Making." *TechCrunch*, February 24. <http://techcrunch.com/2015/02/24/new-firm-combines-wearables-and-data-to-improve-decision-making/>.
- O'Toole, Kathleen. 2013. "Susan Athey: How Big Data Changes Business Management." *Insights by Stanford Business*, September 20. <https://www.gsb.stanford.edu/insights/susan-athey-how-big-data-changes-business-management>.
- Remus, Dana, and Frank S. Levy. 2015. "Can Robots Be Lawyers? Computers, Lawyers, and the Practice of Law." December 30. Available at: <http://dx.doi.org/10.2139/ssrn.2701092>.
- Reich, Robert. 1992. *The Work Of Nations: Preparing Ourselves for 21st Century Capitalism*. New York: Alfred A. Knopf.
- Schwab, Klaus. 2016. "The Future of Jobs Report." World Economic Forum, January. <http://reports.weforum.org/future-of-jobs-2016/>.
- Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter. 2015. "Firming up Inequality." Working Paper no. 21199, NBER, Cambridge, MA.
- Wallsten, Scott. 2015. "The Competitive Effects of the Sharing Economy: How is Uber Changing Taxis?" Unpublished Manuscript, Technology Policy Institute.

Copyright of NBER Innovation Policy & the Economy (University of Chicago Press) is the property of University of Chicago Press and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.