PennState
College of Engineering

Productivity 4.0 – Measuring and Managing Productivity in Industry 4.0

Vignesh Ramanathan, Graduate Student & SEE Fellow
Naren Gursahaney, President and CEO (retired), The ADT Corporation
Paul Funari, Vice President, UPS Security Services (retired)
James Wilk, Senior Manager, Accenture, LLP
Charles R. Schneider, Founder and Chairman, U.S. Security Associates, Inc.
Vittal Prabhu,* Professor and Director, SEE 360 Initiative

Service Enterprise Engineering (SEE 360) Initiative
Engineering the 21st Century Economy
see360.psu.edu

*Corresponding Author prabhu@engr.psu.edu

What Is Productivity? Why Is It Relevant?

Utilization and Productivity are metrics often used in industries, especially when trying to arrive at aggregated trends for an entire industry or across different industries. Improving productivity is key for businesses to gain a competitive edge. This has been true since the days of Frederick Taylor, one of the pioneers of Industrial Engineering.

The context in which these metrics are applied has changed greatly over the past few years. The advent of technologies and related phenomena such as cheaper computing, cloud computing, cheaper sensor technology, the internet of things and 3D Printing have led to new possibilities for the factories of the future. ***These new systems rely on much greater flows of information than was previously possible and is often described under the umbrella term – Industry 4.0.*** Some (particularly members of the European Union Commission) argue that the service industry needs to rethink its strategies in a similar fashion and that a Service 4.0 needs to be defined too. Given this greater reliance on information, crucial metrics that define productivity are only more important than ever.

Better technology today facilitates greater information flow, crucial to improving productivity

It is important that productivity should keep growing constantly, as this helps improve both the quality and quantity of goods and services available to the consumer. In addition, improving productivity helps promote competitiveness. In the late 1980s, Walmart controlled only 9% of the retail sector by market share. Yet, their productivity was 40% higher than their average competitor's, fuelling their rapid growth that led to their market dominance in the retail sector today.

From an Industrial Engineering perspective, productivity is usually measured by means of resource utilization in a system, especially when the objective is to reduce cost. This represents the ratio of the time for which the server in a system is engaged in busy to the total time over which the server utilization is measured.

In Economics, two of the most commonly used metrics to determine productivity are labor productivity and multi-factor productivity. Surprisingly, productivity growth has been falling recently. This may change very soon, due to new technologies available.

Baumol's Cost Disease

William Baumol, an American economist, described a phenomenon commonly known as Baumol's cost disease in the 1960s. When looking at wages for musicians, he found that they had been steadily increasing over time. Yet, if one looked at the musician's output per unit time produced, it wasn't drastically higher. Music that could only be produced by an orchestra was not magically producible by just one person at his time. Why was this increase in wages occurring? His explanation was that rising productivity in other fields (such as manufacturing at the time) were increasing wages. To keep up with this, musicians had to be paid more for retention. Eventually, this would lead to higher and higher costs for goods and services that were very intensive on human labor and that could not be replicated easily.

One finds that a similar phenomenon is presently occurring today in the United States with human-intensive services such as education and healthcare. For manufactured goods such as automobiles and consumer products, prices have mostly remained steady after adjusting for inflation.

Curiously, while prices for such services have kept up, real wages for the average American after adjusting for inflation have stagnated since the late 1990s. Also, some measures of well-being such as this coincides almost exactly with a drop in productivity.

Measuring Productivity

Labor Productivity (LP) is defined as the ratio of Output Volume to Input Labor Use. Output Volume may be expressed as either the Gross Domestic Product (GDP) of the product or service that is offered or the Gross Value Added (GVA) of the same. Labor Input Use may be expressed as the total number of hours worked by the labor in question or the total employment related to this product or service by the business in question.

Total Factor Productivity, also known as Multifactor Productivity, also considers input from working capital, apart from pure labor input. The Cobb-Douglas production function gives a convenient function to express this. Assuming the production of a single good or service –

$$Y = A.\ K^{\alpha}.\ L^{\beta}$$

Here, Y represents total output, A is total-factor productivity, K is capital input and L is labor input. α and β are output elasticities (percentage change in output per percentage change in input). Standard factors used are 0.7 for α and 0.3 for β. A generalised form of this equation can be extended for a system with more than two goods or services.

The Recent Decline in Productivity Growth: Solow's Paradox

Baumol's Cost Disease has held true for a significant part of developed world economies after World War II. Between 1973 and 1995, annual labor productivity growth rate in the United States was about 1.3%. Some theorise that initial adoption of IT systems that were new and complex to learn at the time, may have caused this.

From 1995 to 2000 however, Labor Productivity growth rate doubled. This Labor Productivity acceleration also contributed to a growth in Multi-Factor productivity from 0% between 1973 and 1995 to 1.4% between 1995 and 2000. This growth seemed to be attributable to gains realised across many industries. Triplett and Bosworth analysed 22 different industries in 2003, of which gains in 15 industries were found. What caused this growth? Some of this growth, especially with respect to MFP, was associated with the maturing and widespread adoption of Information Technology systems. The general boom in Information and Communications technology (ICT) coincides with countries that saw the greatest growth in productivity.

Towards the latter half of the 2000s, there has been a slowdown with respect to growth in productivity. McKinsey estimates that aggregate labor productivity growth averaged only 1.3% annually from 2005-2016, when measured across the United States. According to their findings, two-thirds of all industries in the late 2000s saw a decline in productivity, with manufacturing seeing some of the largest falls. 28 other OECD member countries also saw similar declines in productivity figures. Average annual labor productivity growth from 1995 to 2004 for these countries was 2.3%, but only 1.1% from 2005 to 2016. It is especially drastic from 2010-2014, with an average growth of just 0.5% for these economies. In addition, real median income (after adjusting for inflation) has not grown since the late 1990s.

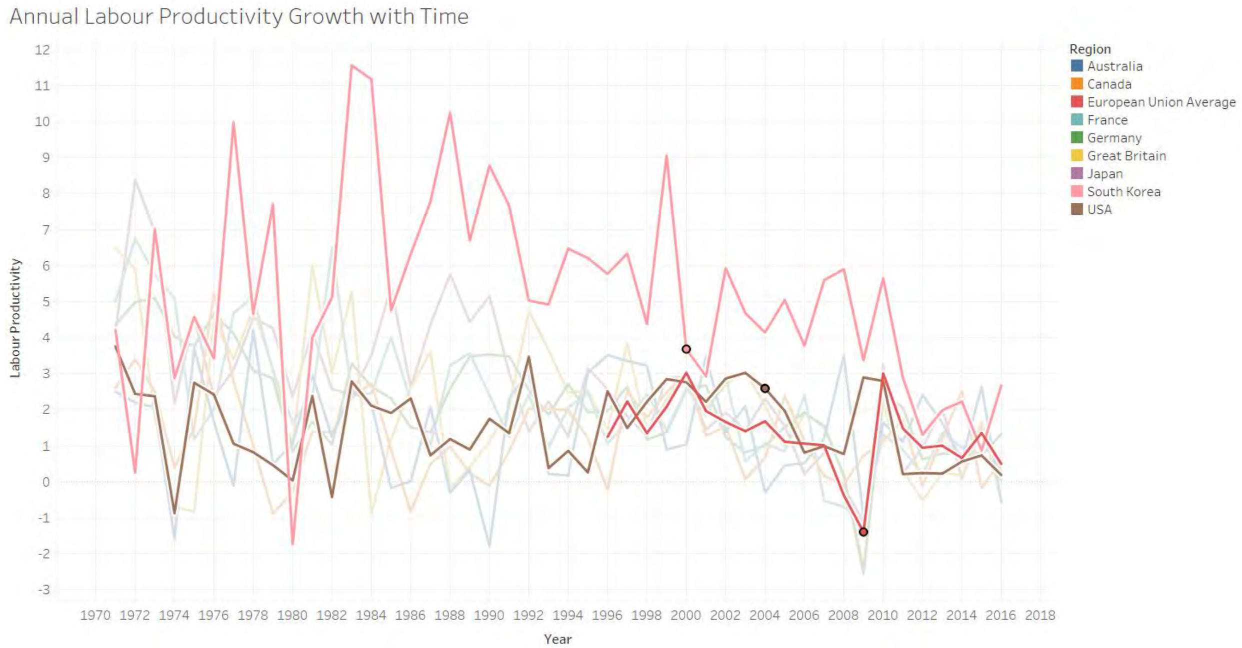

*Labor Productivity Growth Rate for developed countries around the world has varied. USA (highlighted in brown) saw a boost in the 1990s thanks to the ICT boom. The post-recession decline is quite apparent, as is the LP boost immediately after the recession thanks to reduced hours. It is a similar story for the EU (red). South Korea (pink) had enormous growth thanks to manufacturing and globalization from the 1970s to the early 2000s. Data sourced from OECD Data. (*https://data.oecd.org)

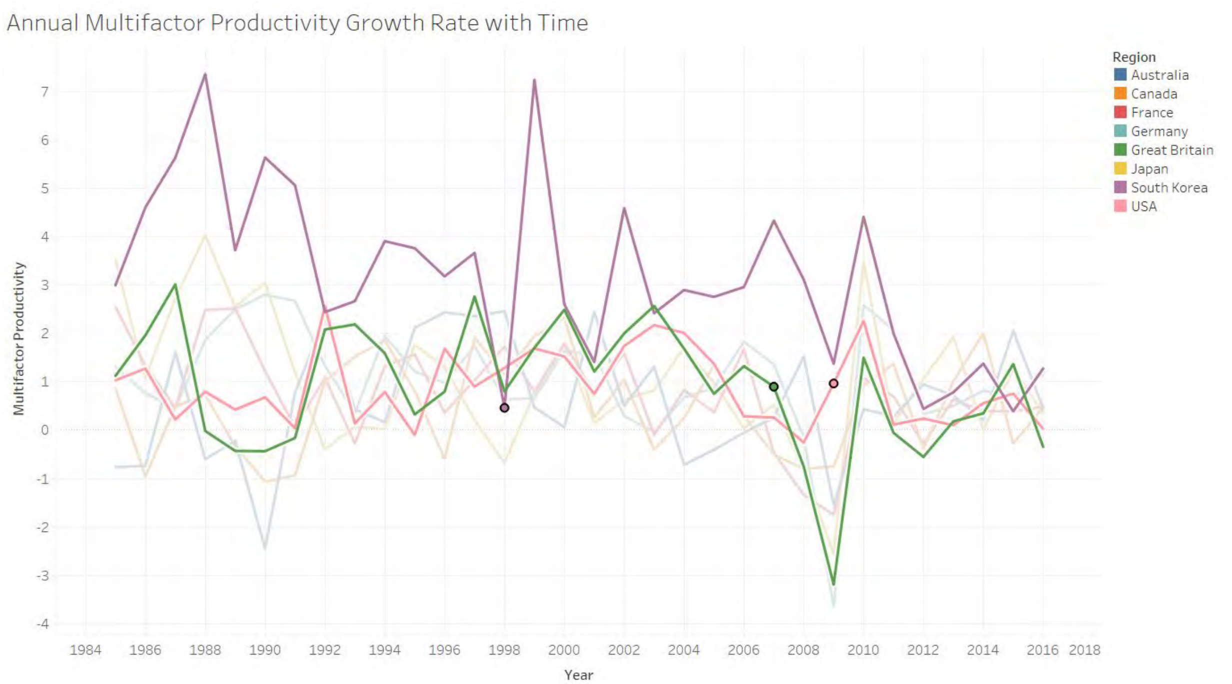

Multifactor Productivity Growth Rates present similar trends. USA (highlighted in pink) sees moderate fluctuations, as does the UK (green). The UK was chosen as aggregated data for the EU here was not available. South Korea's Multifactor Productivity Growth rates (purple) were lower in magnitude *compared to their Labor Productivity Growth rates. However, they still far outpaced the rest of the world from 1984 until the global recession in 2008. Data sourced from OECD Data. (*https://data.oecd.org)

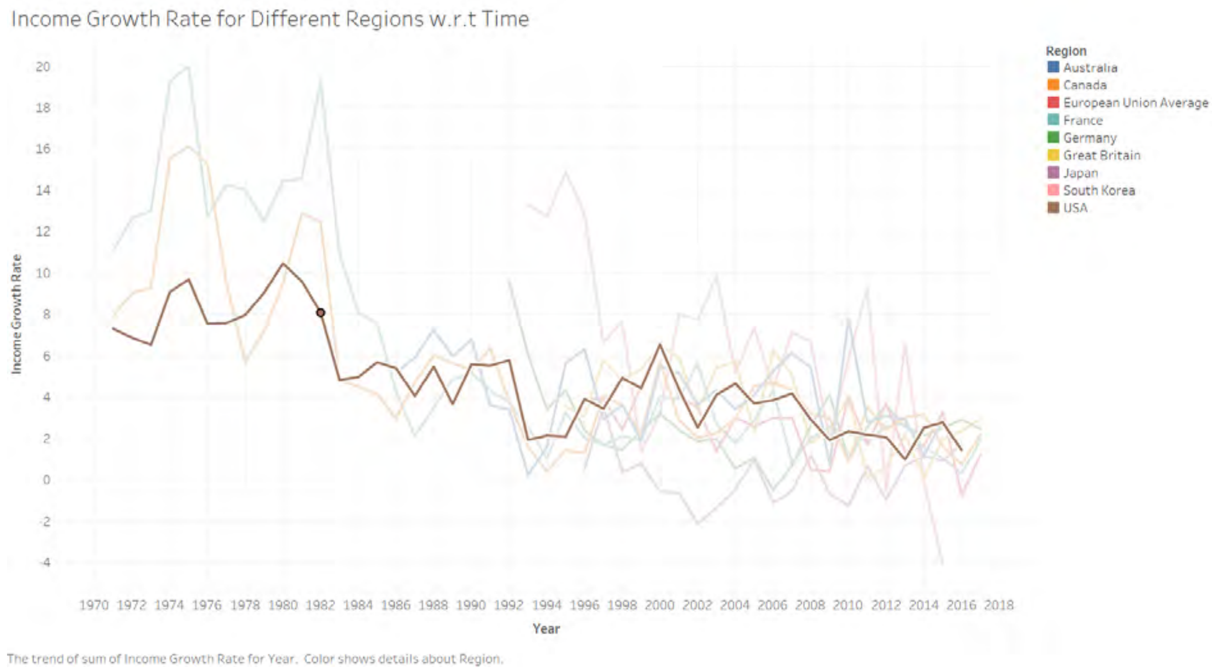

*Income Growth Rate for developed countries around the world has been plummeting since the 1970s. Highlighted is yearly income growth rate of the US. Data sourced from OECD Data. (*https://data.oecd.org)

Karim Foda of the Brookings institute had 4 key findings regarding this slump in productivity –

i. Productivity has slowed for both advanced as well as emerging and developing markets.
ii. TFP and LP growth have declined in all countries, although LP growth reduction has been lesser in emerging and developing countries.
iii. Productivity slowdown has not been limited to a few sectors, but is widespread across multiple industries, service and manufacturing alike.
iv. Productivity may have peaked and is being hampered by the 2009 recession. Studies show that TFP peaked in 2004 in the US and earlier in some European countries.

This is very surprising, giving the ever-increasing IT adoption, as well as the emergence of newer technologies such as machine learning and artificial intelligence. Such a phenomenon isn't entirely new. Upon similar productivity growth slowdown in the late 1980s, ***the American economist Robert Solow famously quoted "You can see the computer age everywhere but in the productivity statistics." This is known as Solow's Paradox.***

The McKinsey Global Institute performed research across 7 countries (France, Germany, Italy, Spain, Sweden, United Kingdom and the United States) and 6 sectors (Automotive, Finance, Retail, Technology, Tourism and Utilities). According to them, the following factors may have led to the recent decline in productivity –

i. **The continued general reduction in productivity that has been occurring since the 1990s** – This is coincident with the wane of the ICT boom seen in developed economies. In addition, supply chain maturations, especially in the retail sector, provided for little further opportunity in productivity gains. The effects of globalization and offshoring had also consolidated, leaving little room for more improvement. The technology industry also became heavily dependent on highly skilled laborers rather than automated machines, making potential productivity gains low. Manufacturing in the technology industry, one of the catalysts behind the ICT boom, has largely been relocated to Asia, Mexico and Eastern Europe. As a result, the technology industry productivity has growth declined by 14 percentage points in the past decade.

ii. **The after-effects of the 2008 recession** – The financial crisis affected certain sectors such as financial services and real estate very badly. The bursting of the housing bubble was one of the root causes of the recession. As a result, there was very little market for real estate, meaning that output volume was slashed drastically. This led to very low productivity. A similar trend was seen in the financial sector too. Crucially, this crisis reduced demand appetite for all products and services across the world. Even for the products in-demand, household began to opt in for lower value and cheaper goods to the extent possible. This lack of demand led to layoffs and slashing of worker hours across every industry. In fact, worker hours reduced to such an extent that in the middle of 2008-2009, there was a brief uptick in productivity thanks to reduced hours. Beyond 2009 however, the demand reduction was too severe to prevent productivity from plummeting.

iii. **Large scale digitization has occurred, but the effects of it seem to not have manifested yet** – MGI believes that Solow's Paradox is in effect with respect to the digital age. This is in accordance with AI being a General-Purpose Technology (GPT), where time lags are often felt.

iv. **Shift from a manufacturing economy to a service economy** – As more and more economic activity in OECD countries shifts towards the service sector, Baumol's Curse comes into effect here. Improvements in these fields simply aren't as easy as improving a process in manufacturing. Of course, it is to be noted that there are concerns that mismeasurement of productivity exists with service economies.

v. **Consolidation of businesses and decreasing competitiveness in industry** – Some think that the increasing consolidation of companies across every industry may be contributing to this productivity slowdown. This is because competitive industries usually involve each individual company trying as hard as possible to gain an advantage to help increase market share. Improving productivity would mean that more output would be possible per unit time, hence improving sales volumes. This incentive isn't quite as strong in a consolidated industry that is dominated by only one or two players. However, evidence for this isn't particularly strong. Productivity losses are dropping uniformly for all players across industry. It was found that for frontier firms across industry, the late 1990s saw an average annual productivity growth of 4-5% for frontier firms that were industry leaders, and 1% for other players. These figures have decreased to 1% for industry leaders and 0% for other players post-recession. This suggests that market consolidation may be independent of productivity loss. Moreover, some fields such as the automobile sector actually saw gains for leading German companies, but not so for leading American companies. Yet, other sectors such as the technology manufacturing sector and the finance sector saw industry consolidation coincide with reduced overall productivity growth.

There are some characteristics of the post-recession plunge unique to this period. ***This period has been 'job-rich, but productivity-poor'.*** Productivity declined uniformly across all industries due to the demand slump. This resulted in layoffs. As the economy began its recovery, jobs were slowly added back. Yet, conditions for these workers were not as friendly as before. ***The average time spent per worker has increased across all industries and OECD countries.*** This has increased the denominator in any productivity metric. Yet, gains in output have not been commensurate to this gain in worker time spent. ***This suggests that the value added by workers in these fields have decreased overall.*** It very possibly may be that the longer hours are now inducing some fatigue among workers, thereby decreasing the marginal productivity gain per extra hour spent working.

The retail sector is a microcosm of recent changes in productivity. The ICT boom led to new innovations such as better Vendor Management and Warehouse Management systems and improved forecasting. Consumer spending was best represented through the retail boom - especially with the surge of malls - and supply chain changes saw great productivity boosts. The impact of changes caused by all these factors have drastically lessened of late.

There are many potential areas for improvement that would translate across most industries –

a. **Increase in operational efficiency through higher levels of automation**: While this is already occurring heavily across manufacturing industries, many opportunities for this would exist in the service sector as well. One example of this has been the increase of automated tellers for banks (ATMs), retail stores (self-checkout counters) and airline operations (check-in counters). There are concerns that this may lead to displacement of human workers. However, that may not necessarily be the case and could lead to the mere shifting of job functions. This has been seen with bank tellers.
b. **Introduction of new business models to boost productivity**: The emergence of new business models thanks to the internet, platforms and the sharing economy presents potential productivity improvements. This would especially be so if consumer productivity as well was captured in addition to producer productivity.
c. **Reduction to barriers of entry**: Newer business models also give chances to cut out the middleman from the process who wouldn't add much value otherwise. This has already been occurring with services such as eBay and Uber. Services such as the internet also provide greater information to the end consumer, thereby promoting price transparency. Lower regulation and bureaucracy would also help improve productivity in many processes.
d. **Supporting worker training and transition**: One of the biggest concerns with automation is that workers would be instantly displaced. Yet, automation is most effective when it is used to augment a process rather than completely displacing a worker. For this, workers would need retraining to effectively be able to work with AI and ML. In addition, retooling a worker's functions to assume more of a supervisory role would boost productivity, more so than either a completely manual process or a completely automated process.

The recent slowdown in productivity growth perhaps portends an incoming boom

However, MGI estimates that this slump is only a temporary blip and that very soon, one would be able to see sustained periods of growth with annual rises in productivity at about 2%. The chief trigger for this, according to MGI, would be benefits realised by companies' heavy investments in Artificial Intelligence and Machine Learning.

Machine Learning, Artificial Intelligence & Productivity

The number of systems involving the use of machine learning and artificial intelligence has mushroomed in the past decade. ***When correctly trained and implemented, such systems can surpass human performance in both accuracy and speed.*** There has been large growth in systems implementing these technologies especially thanks to the open source nature of a lot of ML implemented nowadays. Yet, this seems to have had little impact so far on productivity.

Many argue that we are on the cusp of a tremendous breakthrough. This sentiment is shared among many technology leaders. Ray Kurzweil, presently the director of engineering at Google propounded in the early 2000s that by 2045, a 'technological singularity' will have occurred. This is proposed as the tipping point where machines and artificial intelligence would be better than humans at a wide variety of tasks and that these intelligences would themselves be capable enough to develop other systems. This hypothesis was arrived at as the logical result of what he termed as 'The Law of Accelerating

Returns'. By this, he states that every technological advancement induces a positive feedback loop that makes future advancements possible faster than previously thought. With time, these gains add up to the extent that growth is no longer a linear phenomenon, but an exponential one. This view has gained traction especially after many saw the duration for which Moore's Law was valid.

One could argue that this held valid for IT during the 2^{nd} half of the 20^{th} century and a case could be made for AI and ML in the first 2 decades of the 21^{st} century. One moment that sparked interest in Machine Learning was the ImageNet Classification by Krizhevsky, Sutskever and Hinton. This was one of the first occasions where convolutional neural networks, a form of deep reinforcement learning was used for a more widespread application such as image classification. This algorithm performed extremely well, obtaining a test error rate of 15.3% in the 2012 ImageNet Large Scale Visual Recognition Challenge as compared to the next best entry that had a 26% error rate. Neural Networks weren't completely new technologies. The first instance of a Neural Network was propounded by Rumelhart, Hinton and Williams in 1986. What made this algorithm feasible to classify 15 million images was the drastic improvement in computational power. After 2012, many other researchers attempted similar approaches to the problem. By 2017, 29 of the 38 competing teams in this competition obtained a misclassification error rate of below 5%.

Machine Learning has also yielded promising results in other tasks previously deemed possible only for human brains. There has been much improvement in voice recognition, to the extent that voice-recognizing AI is widely deployed today by companies such as Google and Amazon. Many complex games have also seen computers beat some of the best human players in the world in their respective fields. Examples include IBM Watson participating in Jeopardy in 2005, IBM Deep Blue defeating grandmaster Garry Kasparov in 1997 and Google's DeepMind program AlphaGo defeating professional South Korean Go player Lee Sedol in 2016 in a 5-game series.

Much potential lies with commercial applications, be it the manufacturing or service industry. For instance, there were 3.5 million individuals employed as motor vehicle operators (including truck and cab drivers) in the United States. Given that level 4 autonomous vehicles from Waymo (where cars are fully autonomous within their operational design domain) have already been operational in Phoenix since 2017, it is not unreasonable to see mass implementation of autonomous vehicles within the next decade or two. Such vehicles are incapable of handling extreme emergency scenarios by themselves at present. On similar lines, it would be very possible to automate farm equipment. This would be a great boost to productivity measured in terms of utilization, given that conventional automobiles can be idle for about 90% of their life time. Other benefits such as increased safety and lower costs would also accompany self-driving cars.

Much potential lies in productivity improvements for the transportation industry

The service industry is also starting to see AI-enabled applications. For instance, Google unveiled in May 2018 a demonstration of a virtual assistant setting up an appointment with a hair salon completely by itself. If successful, these could possibly substitute call centres in the future. Another interesting application was in in 2016 when DeepMind was used to monitor power consumption with Google's Data Centres. With this implementation, the power consumption required for cooling these data centres was reduced by 40%.

An increasing number of devices are beginning to incorporate sensors. Given the amount of data that this generates, there is much potential for machine learning applications with these. This, along with the Internet of Things (IoT), holds promise to transform manufacturing. As of 2018, 23 billion devices are connected to the IoT, with this projected to only grow exponentially. ***One of the biggest advantages of some reinforcement learning techniques in use today is that these techniques are usually self-improving.*** Hence, updating and maintaining such systems (if successfully implemented) would require less work than the status quo. Once devices are operational, they continuously feed back data into a central database, which is further used as training data for future use cases. This is akin to the law of accelerating returns mentioned by Kurzweil. Decreasing costs behind sensors and computing power would only further accelerate this. The open source nature of many technologies today only serves to hasten the process of advancement. This also incorporates the benefits of network effects. For instance, Waymo has 25,000 vehicles operating on the road. Any new incident encountered by even one of these cars would be instantly accessible to other cars, making information transfer much faster.

Tasks Suitable for Machine Learning

It is important to note that Machine Learning is not a means to solve any sort of problem. In addition, the most crucial aspect of any ML application is the data. Even a simple model would outperform a deep-learning based model with good data. ***Hence, it is paramount that good features to train an algorithm be identified and that appropriate use cases that cover as many scenarios as possible be involved in the training data for any algorithm.***

In addition, tasks such as classification still require baseline rules known as 'the ground truth' be established by a human before training is performed. A simple example is as follows – Let us assume that an algorithm is being trained to classify comments on a video as either spam or not spam. While an algorithm can be used to instantly classify millions of comments, it still needs to learn from training data, that depends on humans classifying these comments manually. Even if this classification process in itself were automated, it would still be dependent on rules that were constructed by the human at hand prior to automation. This is applied for a separate dataset known as training data, which the algorithm learns from. If this initial classification is itself faulty, the algorithm can never perform well. If one were to try and extend this concept to productivity, one would find that machine learning models are very unproductive initially, with lots of input time required to produce output that is not wholly reliable. With time however, as one scales up with the amount of input data fed, one would be able to improve productivity using AI.

From the above example, it is clear that training an algorithm is highly specific to the task at hand. Keeping this in mind, there are types of tasks for which Machine Learning excels. Brynjolfsson and Mitchell identified some key criteria that could be used to identify tasks suitable for Machine Learning applications. Some of these include –

- **A function that maps well-defined inputs to well-defined outputs** – For a certain input, there exists a clear output which cannot be disputed. Examples of this include predicting whether a person is at risk of heart disease by looking at associated historical attributes of people with and without heart disease and predicting no-shows for an airline using historical data. In each of these cases, the output is clearly established for historical data. In such cases, it may be so that only the correlation between input and output is captured by the algorithm as opposed to the causation.
- **Large datasets can easily be created or already exist for the given problem** – The performance of an algorithm on new data depends on the amount of data used to train this algorithm. The more the data present, the better the algorithm performance. The more reliable the algorithm performance, the more reliably can one apply these techniques to large-scale problems, thereby improving productivity.
- **The process provides clear feedback with well-defined goals and metrics** – Extending the first point, ML works at its best when performance is easily quantifiable. If different scenarios exist for a task where adjusting values of a certain parameter result in different values of a KPI, this task could be improved using ML. If these metrics are defined for a large system, ML works very well as there would now be several different parameters to examine. An example of such a metric would be system utilization in case of a large traffic system or house prices keeping in mind different criteria such as economic, geographic and social factors. These are prime use cases to boosting productivity as there is a clear metric that can be improved upon.
- **The output does not depend on long chains of logic or diverse background knowledge** – ML models are best at predicting output given a certain input state. If a logical step is involved in processing this data, the algorithm usually fails to learn this. In addition, if appropriate conclusions can be made of a field only after applying a reasonable amount of background knowledge to a task, ML will usually not be able to outperform a human. Some exceptions exist, such as chess and Go, as these games involve simulating a large number of potential outcomes that require a lot of processing power. These are tasks which computers excel at. These are tasks which computers excel at. Such complex tasks usually aren't perceived as those which

need improvements in productivity metrics, for a large amount of input time and effort is required to produce some form of meaningful output.

Machine Learning cannot automatically be applied to every task. Certain forms of activities are more conducive to improvements through such techniques

- **An explanation behind the decision-making process is not essential** – One of the biggest problems behind complex models is that it is not easy to discern *how* a decision was arrived at. Models that do provide transparency such as linear regression or classification trees are usually not very powerful or accurate. This is one of the arguments by many that ML should not be used as a tool by itself but should be used to support the decision-making process. Some state that this would be so only for complex tasks such as detecting potential heart-disease within patients. As these are complex tasks in themselves, the addition or subtraction of AI would not impact productivity.
- **Error tolerance and imperfection in solutions is permitted** – Very few models can ever be a 100% accurate. This is thanks to the heavy reliance of ML on probability and statistics, inducing a certain amount of variance into the process. For problems where only improvements are required over the status quo as opposed to a perfect, optimal solution, ML solutions would be appropriate.
- **The task involved should not change over time** – Unless the algorithm itself were continuously retrained over time with new data, it is important that the basic task should not change over time. For instance, the same definition of spam should hold when trying to classify new input emails. Hence, we can see that ML has the potential to transform productivity of only those tasks that are relatively time insensitive in nature.

Concerns Regarding the Impact of Technology & Automation On Productivity & The Workforce

Economic reasons dictate that large-scale automation of tasks is a near-certainty in the future. Another catalyst for automation would be shrinking labor supply in developed economies across the world. This would be a cause for concern for productivity growth. A fall in labor growth could lead to a fall in growth of the economy as well. According to a Bain report in 2017, the United States labor force growth is forecasted to slow down to 0.4% by 2020. Estimates in this report state that of the 3.1% total GDP growth annually among OECD member nations between 1961 and 2015, 1.1% was thanks to labor force growth. In the US, 1.4% of the total 2.8% GDP growth annually in this period, could be attributed to labor force growth. Mass retirements of the baby boomer generation would reverse this trend. The potential large-scale elimination of jobs (assuming no worker retraining occurs) would also affect labor force growth.

One way in which this could be reversed were if productivity gains due to automation outpaced contributions by labor force growth. For this to occur, productivity growth would have to increase at a rate of 54% compared to growth rate between 1995 and 2015. Given the increasing shift of activities from the manufacturing sector to the service sector, this will not be easy to achieve. Service economies provide lower potential to improve productivity (by existing metrics). Bain compares the automobile manufacturing sector and hospitals to illustrate this. From 1993 to 2014, the automotive industry saw a 28% growth in productivity per worker, but a 28% decline in labor participation rate. Conversely, hospitals saw only a 16% growth in productivity per worker, but a 28% growth in labor participation rate.

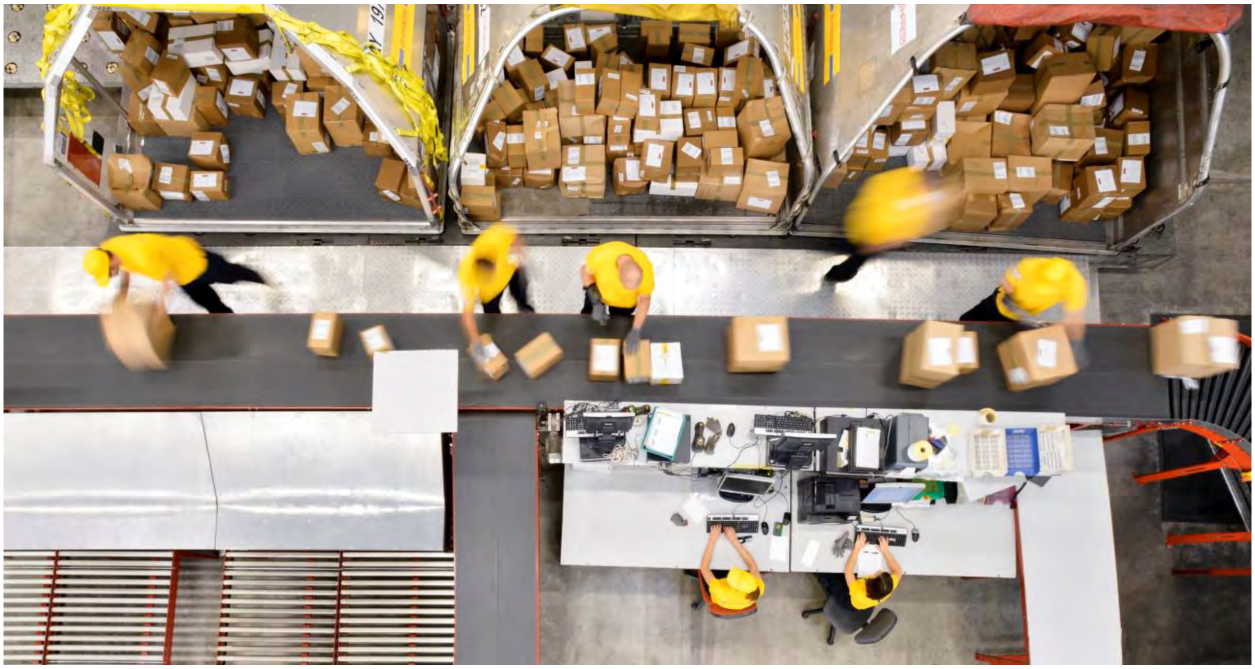

Automation potentially could disrupt the lives of millions of workers across the world

The manufacturing sector has seen success with automation fuelling productivity gains. In 2016, the Chinese firm Foxconn was able to replace 55% of its workforce in one factory or about 60,000 workers with robots. Automation is predicted to affect 80% of existing workers through factors such as wage-suppression and job loss. In the present scenario, the benefits of automation would be disproportionately enjoyed by highly skilled, highly paid workers. This may be alleviated through a slow transition to automation to allow time to retrain workers. Presently, the transition is occurring rapidly and would not allow for this.

To alleviate this toll on labor, some countries are considering the introduction of a Universal Basic Income for all of its citizens. In addition, the manufacturing sector is seeing the emergence of ‘co-bots’, devices that require human supervision and handling to operate effectively. Newer categories of jobs are also emerging, such as data analysts and social media marketing managers. Yet, it is to be noted that these workers tend to be highly skilled and highly paid.

According to a literature survey performed by Yang and Brynjolfsson, some studies indicated that the productivity gains in the 1990s were associated with transactional types of Information Technology (such as data processing), but not due to strategic systems (such as sales support) or informational investments (such as email infrastructure). Yang, Rock, Syverson and Brynjolfsson suggest that there are some general issues to be kept in mind while assessing the impact of technology on productivity –

- **Mismeasurement of inputs and outputs** – It may be that input and output simply aren’t measured effectively by existing metrics especially after the introduction of IT systems. This is exacerbated by the fact that IT systems are used extensively in service industries, where measurement of productivity is not easy. IT systems usually improve aspects of a process such as customer service, service quality, service convenience process speed and process responsiveness. These are aspects not typically measured in productivity statistics. For example, the opening of a 24-hour ATM is a massive improvement in convenience and quality as opposed to a human bank teller. Yet, this would not be measured by productivity statistics that merely look at factors such as the number of transactions. Other examples also include the ease with which information can be transmitted thanks to IT advancements and more recent advancements such as smartphones. Thanks to all of these factors, mismeasurement could potentially be dragging down productivity statistics quite severely.
- **Redistribution and dissipation of profits** – It may be that recent advancements are beneficial only for the advancement of productivity within a firm and to the industry as whole. Often, a firm’s priorities seem more towards taking profits away from competitors rather than reducing costs. This is especially the case with resources such as information, whose dispersal is proliferated with newer systems. Intense IT usage for activities such as market research and general marketing do not advance industry productivity.
- **Concentrated Distribution** – It is also seen that the benefits of many cutting-edge advancements are seen only by a small number of industry-leading firms. Given that most industries are heading towards increased market consolidation and lesser competition, there is little incentive for firms to share information. In addition, upon obtaining novel technologies, many firms seek to hide and obstruct these technologies, preventing competitors’ access. While this may lead to short term benefits for the firm, it does not help serve long-term industry advancements.
- **False Hopes** – It simply may be that these technologies are not as revolutionary and have as much impact on general productivity as per expectations. There have been technologies in the past such as nuclear power, supersonic travel and interstellar exploration that promised much, but have not delivered on their expectations.

- **Lags due to Implementation, Learning, Adjustment and Restructuring** – One of the most widely shared opinions is that large scale implementation of General Purpose Technologies (GPTs) simply take a lot of time to install and develop. Thanks to the learning curve associated with these technologies, it takes a while before potential productivity benefits associated with these technologies can be realised. IT is considered as a GPT and Artificial Intelligence seems to tend towards the same as well. Studies considered in the late 1980s showed that the effects of IT implementation usually took 2-4 years to manifest themselves. The reasons behind this lag are twofold. Firstly, it simply will take time before effects of any large-scale system change takes place. This is true of any GPT. In addition, many of these technologies require complementary infrastructure and training to be set up to ensure maximum utilization of this new resource. For IT, it implies hardware upgrades and training, while extensive training and new software is required with Artificial Intelligence and Machine Learning. Other instances of complementary technology requirements include a charging station network for electric cars and an intricate warehouse and supply chain network for an eCommerce business. An example of this time lag is the growing prevalence of the usage of artificial neural networks for different applications post-2012, despite the first appearance of these algorithms in 1986.
- **Past Productivity is a Poor Predictor of Future Long-term Productivity** – It is also to be kept in mind that historically, past productivity growth metrics have been poor predictors of future productivity growth metrics. While for a very short year span, past performance is significant, there is little correlation between productivity metrics (both LP and TFP) that are measured ten years apart.

Acknowledgements

The authors appreciate the feedback provided by Chintan Patil, Graduate Student at the Pennsylvania State University Department of Industrial & Manufacturing Engineering. He is also an SEE Fellow at the Penn State SEE 360 Initiative.

References

1. Herman M., Pentek T., Otto B. (2016, January). Design Principles for Industrie 4.0 Scenarios. In *2016 49th Hawaii International Conference on System Sciences (HICSS).* Retrieved from https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7427673
2. Brynjolfsson E., Rock D., Syverson C. (2017, November). Artificial Intelligence and rhe Modern Productivity Paradox: A Clash of Expectations and Statistics. In *NBER Working Paper Series.* Retrieved from http://www.nber.org/papers/w24001.pdf
3. Brynjolfsson E., Mitchell T. (2017, December). What Can Machine Learning Do? Workforce Implications. In *Science Magazine.* Retrieved from http://science.sciencemag.org/content/sci/358/6370/1530.full.pdf
4. Bosworth B., Triplett J. (2003, September). Productivity Measurement Issues in Services Industries: "Baumol's Disease" Has Been Cured. In *FRMBY Economic Policy Review.* Retrieved from https://www.newyorkfed.org/medialibrary/media/research/epr/03v09n3/0309trip.pdf
5. Yang S., Brynjolfsson E. Information Technology and Productivity: A Review of the Literature. In *Advances in Computers*, Academic Press, Vol. 43, pages 179-214, 1996. Retrieved from http://ebusiness.mit.edu/erik/itp.pdf
6. Foda K. The Productivity Slump: A Summary of the Evidence. In *Global Economy and Development at Brookings*. Retrieved from https://www.brookings.edu/wp-content/uploads/2016/08/productivity-evidence.pdf
7. Remes J., Manyika J., Bughin J., Woetzel J., Mischke J., Krishnan M. (2018, February). Solving the Productivity Puzzle: The Role of Demand and the Promise of Digitization. By *McKinsey Global Institute*. Retrieved from https://www.mckinsey.com/~/media/McKinsey/Featured%20Insights/Meeting%20societys%20expectations/Solving%20the%20productivity%20puzzle/MGI-Solving-the-Productivity-Puzzle-Report-February-22-2018.ashx
8. Harris K., Kimson A., Schwedel A. (2018, February). Labor 2030: The Collision of Demographics, Automation and Inequality. *Bain & Company Inc.* Retrieved from http://www.bain.com/Images/BAIN_REPORT_Labor_2030.pdf
9. Lee T. (2017, May). William Baumol, whose famous economic theory explains the modern world, has died. From *Vox.* Retrieved from https://www.vox.com/new-money/2017/5/4/15547364/baumol-cost-disease-explained
10. Kurzweil R. (2011, October). Kurzweil Responds: Don't Underestimate the Singularity. In *MIT Technology Review.* Retrieved from https://www.technologyreview.com/s/425818/kurzweil-responds-dont-underestimate-the-singularity/
11. Kurzweil R. (2001, March). The Law of Accelerating Returns. In *Kurzweil Accelerating Intelligence Essays.* Retrieved from http://www.kurzweilai.net/the-law-of-accelerating-returns
12. Evans R., Gao J. (2016, July). DeepMind AI Reduces Google Data Centre Cooling Bill by 40%. In *DeepMind Blog.* Retrieved from https://deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-40/