The Effect of Artificial Intelligence Implementation on Total Factor Productivity

by

Matthew B. Toy

* * * * * * * * *

Submitted in Partial Fulfillment
of the requirements for
Honors in the Department of Economics

UNION COLLEGE
March, 2021

ABSTRACT

TOY, MATTHEW B. The Effect of Artificial Intelligence Implementation on Total Factor Productivity. Department of Economics, March 2021.

Investment in and availability of artificial intelligence has become a central concern for most developed economies because of is expected positive impact on an economy. Unlike other forms of capital investment, investment in AI may lead to innovative products and processes that should increase productivity. However, AI's overall effect on productivity remains largely unknown. Adopting AI replaces labor with capital, which will have a positive effect on labor productivity, but overall productivity may remain the same or even decrease. I look at the impact of AI implementation on Total Factor Productivity (TFP) in order to assess its effect on the economies of the developed world. The data on AI use is from the Stanford Human-Centered Artificial Intelligence database, which provides comprehensive measures of a country's adoption of AI. Utilizing the methods set forth by Letta and Tol (2018), I perform a cross-country comparison of AI's effect on TFP. I use OLS to estimate a model of national productivity which controls for country specific factors that would drive TFP and is focused on productivity growth due to the implementation of AI specifically. My findings suggest that more investment in AI implementation does not increase overall productivity. However, I do find that the number of startups focused on AI cause an increase in TFP. These findings contribute to the discussion of the productivity paradox and support the justification that an implementation lag may play a substantial role in the limited short-term productivity seen as a result of AI implementation.

TABLE OF CONTENTS

CHAPTER 1

INTRODUCTION

Countries throughout the world are rapidly investing in artificial intelligence. Its implementation is costly, yet countries large and small desire this new technology and are willing to pay the price. However, the effect of artificial intelligence on overall productivity is largely unknown. Economists have demonstrated that improving technology in a given nation should increase the productivity of that nation as a whole, but artificial intelligence implementation itself has not been shown to boost productivity statistics. This productivity paradox has been justified with many potential explanations (Brynjolffson, Rock and Syverson, 2017, Aghion, Jones and Jones, 2017, Syverson, 2017). Ultimately, the most common explanation is based on the time lag between when artificial intelligence technology is initially implemented and when its effects can actually be seen in the productivity statistics. The research throughout this paper reinforces that conclusion and provides evidence supporting the argument.

The anticipated effect that artificial intelligence will have on productivity makes investment in the new technology attractive. When superior levels of technology are implemented throughout a given country, that country is able to produce more output with the same levels of labor and capital inputs. Referring back to the basic production function, where output is a function of both labor and capital, the impact of technology can be measured by the residual difference between the actual output produced and the expected output produced. Implementing artificial intelligence would theoretically boost the impact of technology on a given country's output and therefore increase productivity of that country as a whole.

By nature, artificial intelligence is not only used in one industry and its spillover effects make its development even more advantageous. Technology that uses artificial intelligence has crossed industry verticals and its application has been widespread. As a result, positive externalities are created, and the impact of artificial intelligence implementation becomes more than simply developing a new machine. From autonomous automobiles to automated surgeries, the positive implications of artificial intelligence implementation can be seen in almost any industry.

In an attempt to capture the effect of artificial intelligence investment specifically, I use five independent variables that measure artificial intelligence implementation. These variables measure the quality of research focused on artificial intelligence, overall interest in artificial intelligence technologies, private investment in artificial intelligence and the number of startups focused on artificial intelligence. The data comes from the Stanford Human-Centered Artificial Intelligence database and contains 27 countries over four years. For my dependent variable, I use total factor productivity. Total factor productivity is commonly used by economists and is an indirect measure of the productivity growth due to technology.

Previous research suggests that there will be no visible impact on productivity as a result of the implementation of artificial intelligence. While this may come as a surprise, economists offer different perspectives on the productivity paradox and why it occurs. The most common and widely held belief is that there is a productivity paradox because of the time delay after implementation. When artificial intelligence implementation occurs, it takes time for its widespread use and effects to be seen in the overall productivity statistics. As a result, the productivity we might expect to see as a result of the implementation is not there. This effect is not unusual and largely resembles the Solow Productivity Paradox of the 1970s and 1980s during which productivity largely remained the same even with rapid increases in information technology. My research supports these conclusions and provides further evidence for why the time lag argument is the superior reasoning for this paradox.

I use multivariable ordinary least squares regression models with both one and two-way fixed effects to develop my conclusions. In these models, I use fixed effects for time and country to account for potential endogeneity issues. I also use per capita estimates of my variables to deal with the large differences in country and economy sizes. I use a logarithmic construction similar to Letta and Tol (2018) to achieve reliable results with vastly different sized nations.

As expected, my findings are largely consistent with previous research. Research, overall interest, and private investment in artificial intelligence have no immediate impact on total factor productivity. These findings uphold the productivity paradox theory and are unsurprising. If there was a positive effect of artificial intelligence implementation in the short run, I would see an increase of total factor productivity due to private investment and other measures that seemingly boost artificial intelligence use in a given country. I believe that the lack of positive productivity effects in the short run is true and due the fact that artificial intelligence is largely in a research and development state and the technology needs to be put into use. Although many of these variables offer no implications for total factor productivity, the variable on the number of startups does have a significant positive correlation. This is inconsistent with previous research and its impact on total factor productivity is largely unexpected. However, this positive coefficient upholds the typical economic belief that artificial intelligence implementation increases overall productivity.

The statistical significance on the number of startups focused on artificial intelligence supports the time argument of the productivity paradox. Unlike other variables I measure, the number of startups in a given nation should have a relatively immediate effect on the economy and therefore the productivity. Forming a startup does not occur overnight – there must be some investment before the startup is launched. This earlier investment contributes to the impact of startups on total factor productivity and may be partly responsible for their fast-acting effect. The other variables I use measure investment (of money, research, or education) that is expected to have a future impact instead of an immediate one. At this point in time, the majority of artificial intelligence investment across countries is in areas that are expected to have less of an immediate impact on the overall economy. Since the one variable that represents an immediate application of artificial intelligence implementation is positive and significant, it supports the claim that there will be a positive effect on productivity from artificial intelligence, it just has not happened yet.

The main target of my research is to determine the effects of artificial intelligence implementation on overall productivity of a given nation. The rest of the paper is organized in the following way. The first section contains background and a literature review of the history of technology implementation, measuring the impact of artificial intelligence, the productivity paradox, and implications on foreign and economic policy. The second section will describe my data sources and variables, followed by a description of my economic model and econometric relationship. The paper will conclude with a section about my regressions and results and finally a formal conclusion.

CHAPTER 2

BACKGROUND AND LITERATURE REVIEW

This chapter contains relevant background on artificial intelligence implementation as well as a literature review. I discuss the potential advantages of the new technology and why firms across the world are making investments. Understanding potential uses and implications of the artificial intelligence is critical to appreciate the impact of the technology on productivity. I provide insight on how the impact of artificial intelligence is measured and discuss the productivity paradox that has occurred with artificial intelligence technology. Finally, I include a section regarding the future policy implications of artificial intelligence implementation and complete this chapter with a brief discussion of the United States and China.

2.1 Artificial Intelligence Implementation

Societies have developed and implemented new technologies in order to increase their productivity for hundreds of years. From steam engines to computers, technology has been at the heart of both scientific improvement and productivity growth. Over the past 150 years, our society has been shifting toward automation (Aghion, et al., 2017). Since the industrial revolution, we have been automating processes that have now led to the introduction of self-driving cars and robotic "dogs" used to identify dangerous issues in factories and warzones[1]. Why are we so fascinated with these automation technologies and artificial intelligence? Because they are projected to drive the economy forward and change the world through the next industrial revolution (Mims, 2018).

[1] Tesla, Boston Dynamics

The automation of various procedures using artificial intelligence is advantageous because it not only decreases the overall cost of production of existing goods but also makes new products possible that could not have been produced without it. Over the past twenty years, United States manufacturing employment has declined by millions even as United States manufacturing productivity has increased over the same time period[2]. At a fundamental level, artificial intelligence using automation does one thing: replaces existing labor with capital investment. This phenomenon motivates firms and businesses to invest in artificial intelligence because it not only decreases their labor costs but also increases their productivity. It is seen as a win-win for firms and has been recognized by major consulting firms McKinsey and PricewaterhouseCoopers as the direction of the future (Cameron, Andrews, and Gillham 2017, McKinsey&Company, 2017).

In more recent years, artificial intelligence, which is widely considered a largely technological and engineering phenomenon, has bled into other industries and is becoming widely available. In the auto industry, Tesla has launched its beta software and cars are on the road with no driver intervention. In 2019, Amazon, a largely automated platform, acquired Whole Foods and brought artificial intelligence into the grocery market space. The increasing implementation in different industries throughout the world demonstrates the beginning of a global shift in adaptation to the newly available technology (Naqvi, 2017). Moreover, the definition of "innovative" has changed for companies in ways that they have not expected. For example, information technology is at the core of self-proclaimed innovation for companies like Pfizer, which works with IBM Health Watson to fuel their immune-oncology research, and the

[2] Bureau of Labor Statistics

United Postal Service which now develops service-oriented technologies using artificial intelligence[3].

The spillover of artificial intelligence into many different industries reflects the following trends. The first is that the increased use of artificial intelligence in both technologically and non-technologically focused firms is causing industry verticals to become increasingly broken and industries themselves to become more interconnected. Industry verticals refer to a group of companies which occupy a specialized market. Naqvi (2017) states that, "now, tech-firms can as easily build cars as they can make software; they can create trading platforms or develop new drugs, or trucks that drive themselves, or drones that fly with little human intervention, or robots that can be nurses or guides or housecleaners or companions." Artificial intelligence has largely created and driven a marketplace with technology firms acquiring non-technology firms and vice versa. The second trend is that artificial intelligence is not simply a fad but instead a major aspect of foundation for the future. In other words, artificial intelligence is here to stay. When new technology is created by firms, it is typically marketed to other firms for use. Artificial intelligence, however, is not marketed for external use but instead is designed for the designer (Naqvi, 2017). As a result, firms in all industries are starting their own artificial intelligence programs or acquiring technology firms in order to gain proprietary access to the new platforms (Naqvi, 2017). This interconnection and large-scale adoption of firms in all different industries has created a shift in technology use that would be extremely difficult to leave behind. Firms throughout the world are investing in artificial intelligence and making critical decisions based on artificial intelligence use today and well into the future.

[3] Pfizer 2016 Financial Report, UPS Pressroom January 2019

The increased investment in artificial intelligence among nations has been referred to as a new arms race (Santos & Qin, 2019). In 2017, Vladimir Putin stated, “Whoever becomes the leader in this (artificial intelligence) sphere will be the ruler of the world (Horowitz, 2018).” While China dominates global artificial intelligence funding, the United States and France have both made artificial intelligence investment a priority (Agrawal, et al., 2019). Elon Musk echoed the seriousness of artificial intelligence progress in 2017 when he said that growth in artificial intelligence technology could spark World War III if it is left unchecked (Horowitz, 2018). This sentiment is held by countries throughout the world and has become a driver of global investment in artificial intelligence.

Although the immense power that artificial intelligence may yield is undeniable, it is important to understand that countries adopting it see it as a general-purpose technology (Agrawal, et al., 2019). With the adaptation of the new technology across many different industries, artificial intelligence allows for different technological advances in a wide variety of spaces. The idea of technology as an enabler has a stronger foundation in healthcare and the auto industry than it does in weaponization and military strategy (Horowitz, 2018). Across the world different countries are exploring potential applications of artificial intelligence in health care systems and treatment (Panch, et al., 2018). The impact of these new technologies on the healthcare industry, for example, could benefit a country tremendously.

Although artificial intelligence is not a military innovation in itself, it could potentially enable a number of military innovations and become weaponized. Currently, the technology remains largely driven by private, commercial, companies and could easily fuse to different countries around the world (Horowitz, 2018). As a result, proprietary knowledge that comes through artificial intelligence research and investment is available to all countries, and all militaries, that are willing to invest and are looking to utilize artificial intelligence in their own way. It would be naïve to think that heavy investment in artificial intelligence and growth of the industry was not tied to national security in some aspect.

Besides industry spillovers and potential military interest, it has become advantageous for countries throughout the world to increase their artificial intelligence development because of the substantial impact it has on the economy as a whole. As a result, we have seen growing investment in the artificial intelligence space in both advanced and emerging economies (Mou, 2019). The United States has allowed the private sector to largely control artificial intelligence development (Agrawal, et al., 2019). This has led to the commercial driven growth that can be seen in private sectors throughout the country. Recently, the United States has begun a more "hands on" approach that includes a plan to double federal research and development spending on Artificial Intelligence[4]. In contrast, China's strong relationship between artificial intelligence and Chinese government funding from the beginning places them on track to lead with the new technology in several sectors (Mou, 2019). This could be reflective of the differences between a capitalist economy (United States) and a socialist economy (China). However, emerging market economies as a whole have invested significantly less in artificial intelligence production even though they may see greater returns from its implementation than more developed economies (Mou, 2019).

Ultimately, the choice to invest in artificial intelligence results from the analysis of a variety of different factors. The technology spillover effects throughout different industries allow nations to see benefits of both governmental investment in artificial intelligence as well as

[4] https://www.whitehouse.gov/briefings-statements/white-house-launches-national-artificial-intelligence-initiative-office/

private and commercially driven growth. An increase in artificial intelligence growth would improve productivity in each industry individually, and as a result the entire economy as a whole.

The anticipated spillovers due to artificial intelligence implementation have high expectations. As stated previously, artificial intelligence is considered a general-purpose technology. General-purpose technologies generate substantial spillovers to the rest of the economy (Hogendorn and Frischmann, 2017). Government policies and regulations should enhance these spillovers and may justify greater public subsidy of artificial intelligence research. However, in a small case study of artificial intelligence spillovers in Sydney, Australia, findings suggest knowledge spillovers are limited and are not in line with its high potential (Cetindamar, Lommers, Zhang, 2020).

With artificial intelligence, the spillover effects that guide economic policy both nationally and internationally are largely focused on the diffusion of ideas locally and abroad. In an extreme example presented by Goldfarb and Trefler (2018), suppose Canadian researchers and scientists controlled all of the knowledge surrounding artificial intelligence and gave it to the United States and China for free. Then, the Canadian subsidy would help the world but would not give Canada any edge over the competition. Understanding the balance between sharing information to improve the worldwide landscape and keeping it locally to give one nation a competitive edge is at the crux of policy development surrounding artificial intelligence. Goldfarb and Trefler (2018) argue that local, domestic spillover effects have larger positive externalities than foreign spillovers. Under growth models presented in their research, they also indicate that the artificial intelligence industry may move to the country with the most relaxed policies (Goldfarb and Trefler, 2018). The intense privacy policies that are seen in Europe may be a poor market for artificial intelligence development when compared to the United States. For countries looking to gain an edge in the artificial intelligence industry, government involvement in these areas may boost the overall effect of the new technology.

2.2 Artificial Intelligence, Measurement, and Total Factor Productivity

Economists have developed ways to measure the impact of new technology on productivity long before artificial intelligence. For years the primary measure of the growth of an economy was gross domestic product (GDP). However, measuring economic growth and improvement with GDP has been refuted for a variety of reasons and is a flawed measure of economic welfare (Jones, et al., 2016). One reason is that GDP growth understates the true gains in output per person that has occurred over the last 200 years (Hulten, 2000). This is largely due to its lack of adjustment for quality improvement of new goods. Another critique is that GDP overstates the true improvement of economic welfare because it does not measure the negative spillover externalities associated with increased economic output properly (increasing unemployment, depletion of natural resources, etc.) (Hulten, 2000). Since economies can grow for reasons such as increased productivity, higher levels of labor, increased capital, etc., economists have developed a different approach to separate economic growth due to productivity from economic growth due to labor and capital inputs.

To understand the effect of a technological change on a given economy, we need to revisit the foundation of growth economics. The most basic production function that is presented in economic literature is presented below:

$$Y_t = F(K_t, L_t) \tag{1}$$

Where Y is output, K is capital services and L is labor services, all at a given time t. Put differently, total output is a function of both the output of labor services and capital services utilized at a given time period. However, this model of output assumes equal productivity in all time periods from labor and capital. In order to account for the changes in productivity that are associated with technical progress, we must adjust the equation slightly.

To quantitatively assess the contribution of each source of growth to the actual growth of output, we need to find a measurement of technical progress. Technical progress in itself is not directly observable but is important for explaining the increase in overall standard of living (Rossana, 2011). As a result, economists have modeled an equation that provides an indirect measurement of the extent of technological progress:

$$Y_t = A_t * F(K_t, L_t) \tag{2}$$

Where A is a measure of technical progress, L is labor services provided and K is capital services provided in a given time period. By modelling economic growth in this way, we can measure technical progress because as A increases, the economy is technically able to produce more output from a given resource base (Rossana, 2011). The difference between the projected level of output and the actual level of output is the result of the improvement in technical change. This is known as the Solow residual which, in other words, is the residual growth rate of output that is not explained by the growth in inputs (Hulten, 2000). Today, the variable A is known as Total Factor Productivity and is an indirect measurement of productivity growth that is due to technical progress.

Total Factor Productivity (TFP) has become the standard measurement of productivity growth due to the implementation of new technology. However, modern economists argue that

TFP may not be the best measure to capture the effects of these technical implementations. Carlaw and Lipsey (2003) argue that TFP does not measure technological change at all. Instead, they believe that TFP is correctly interpreted as, "contemporary returns that are in excess of the normal rate of return on investing in new technologies" (Carlaw and Lipsey, 2003). They provide six examples of how TFP fails to correctly measure the excess output that is not captured by a country's inputs. Antonelli and Quatraro (2009) explain that TFP provides a biased measure of technological change unless it coincides with the effects at local markets. As a result, TFP increases due to technology that are seen in one country may not have the same effect in other countries which have different local conditions. Finally, Bar-Shira, Finkelshtain and Simhon (2003) find that TFP misrepresents productivity due to technological change because it captures other changes in the economy such as savings rate changes or other economic boosters. Because TFP captures other factors that contribute to growth, measuring productivity due to technology with TFP may create skewed results that could be inaccurate. Additionally, they present a similar argument to Antonelli and Quatraro (2009) and suggest that cross-country comparison of productivity growth may be biased due to each country's own characteristics. The impact of artificial intelligence on an economy may not be appropriately measured using the TFP model presented earlier.

2.3 Artificial Intelligence: A Productivity Paradox

The introduction of personal computers in the 1980's was thought to boost productivity across the world (Triplett, 1998). However, in July 1987, economist Robert Solow is famously quoted, "We see the new technology everywhere but in the productivity statistics (Hulten, 2001)." In 1987, it appeared to ring true. Since 1973, average total factor productivity growth rate was 0.2% and the labor productivity growth rate was only 1.1% (Triplett, 1998). This low growth rate was coming at a time when new investment in technology was flourishing. Businesses were investing heavily in computers and the age of information technology as we know it today was beginning; but productivity growth was significantly lower than it was in years past. To put it in perspective, the growth rate from 1948 to 1973 was 1.9% for total factor productivity and 2.9% for labor productivity (Triplett, 1998). How could the decline in productivity growth occur after the introduction of personal computers? The new technology that would eventually change the world was resulting in significantly lower productivity growth. A few explanations followed suggesting productivity measurement issues, or the effects of the new technology were not modeled correctly, but ultimately maybe computers just were not as productive as people thought. Today, we have seen the large-scale effect of computers and their implementation throughout the world. It is safe to say that yes, they do improve our well-being and our productivity. Also, it is worth noting that the sharp rise in oil prices caused by OPEC exerting its cartel power was the main reason for the decrease in TFP growth during the 1980s; it was not computers that caused the fall.

Brynjolfsson, Rock and Syverson (2017) believe that the same productivity paradox that was witnessed before is occurring again, but this time with artificial intelligence implementation. That is, the overall productivity growth that we expect to see as a result of artificial intelligence investment may not actually be there, or worse, even result in a productivity decline. Although systems using artificial intelligence can surpass human capabilities, “measured productivity growth has declined by half over the past decade, and real income has stagnated since the late 1990s for a majority of Americans” (Brynjolfsson, et al., 2017). Again, we must ask how this is possible. The existing literature focuses on three primary reasons.

The first is an issue regarding “false hopes.” This explanation suggests that the expected productivity growth and excitement surrounding artificial intelligence is misplaced. Even though some sectors may show productivity growth with artificial intelligence implementation, it may not be widespread throughout the different industries and therefore may not affect aggregate productivity growth (Brynjolffson, et al., 2017). Aghion, Jones and Jones (2017) attempt to find the linkages between artificial intelligence and economic growth, yet their conclusions do not offer any concrete explanation for this paradox (Aghion, et al., 2017). It is entirely possible that artificial intelligence is not as productive as we think, and the decline in productivity statistics supports this hypothesis.

The second, more widely accepted, explanation for the paradox is that there is a measurement problem. Many economists believe that a traditional TFP calculation that has been done in the past does not offer the best measurement of productivity growth due to artificial intelligence implementation. This idea is supported by many economists and they have attempted to develop new ways to measure not only the effect artificial intelligence implementation on TFP but also the effect of new technologies in general. Antonelli and Quatraro (2010) attempt to create a new methodology that untangles the effect of technology bias on the standard change of the aggregate production function. Also, Seamans and Raj (2018) suggest that the current data is not enough, and it is necessary to acquire data from individual firms. Fox (2012) argues that the aggregation of firm productivity may yield inaccurate results of aggregate productivity due to the current method used. Combined with research mentioned previously by Carlaw and Lipsey (2003) and Bar Shira, et al (2003), it can be concluded that economists are searching for an improved measure of TFP and the effect of technology on aggregate productivity in general.

Brynjolfsson, Rock and Syverson (2017) present the mismeasurement explanation of the paradox in a different light. They claim that the productivity benefits of new technology have been enjoyed but not accurately measured. More specifically, new technologies such as smartphones, tablets, and online social networks are relatively cheap and require minimal cost. However, consumers spend a lot of time using these new technologies which suggests that they receive substantial utility even if they account for a very small share of GDP (Brynjolfsson, et al., 2017). With all of these arguments for a measurement explanation, it is important to note that Syverson (2017) and Byrne, Fernald and Reinsdorf (2016) find that mismeasurement may not be the main cause, if any cause at all, of the productivity paradox. While many of the products today may be offering benefits that are not recorded properly in GDP, the same could be said about technology implementation in any time period (Brynjolfsson, et al., 2017).

The third economic justification for why we may see this productivity paradox with artificial intelligence is that we have not given it enough time. Implementation and restructuring lags of new components and improvements in technology may be having a substantial effect. The diffusion of new technology across industries has presented an opportunity for productivity growth that we have not seen before, but at the same time have not given enough time to fully capture the effects. Unlike the previous two explanations, the argument that there has not been enough time for the full effects of artificial intelligence to come to fruition is in line with the current TFP results. In other words, this explanation does not attempt to justify the unexpected impact of artificial intelligence on aggregate productivity growth by neither claiming that investors and creators were incorrect, nor the measurements are inaccurate. This justification suggests that there are reasons to be excited about the future while taking into account the low productivity growth that we have seen in recent years (Brynjolfsson, et al., 2017).

The gap between implementation and productivity growth that is seen in artificial intelligence is consistent with the Solow productivity paradox surrounding computers. It took a substantial amount of time for the widespread implementation of computers to have a positive effect on productivity growth. Triplett (1998) explains that one of the most common justifications for the Solow productivity paradox is a time lag. In his paper, he analyzes the argument, "you don't see computers in the productivity statistics *yet*, but wait a bit and you will" (Triplett, 1998). Ultimately this is fundamentally the same justification presented by Brynjolffson, Rock, and Syverson (2017).

The third explanation suggests that it takes a considerable amount of time for the effects of artificial intelligence on productivity to be captured. The amount of time that it takes for the technology to reach its full potential and start to make a serious impact is longer than is expected. This occurs for two primary reasons. The first reason is that it takes a significant amount of time to build enough of the new technology that it will affect the aggregate productivity of a nation. In order to make an impact on aggregate productivity, the diffusion of the technology throughout the nation must be widespread and enormous. The second is that the productivity increase may not be captured until investment in complimentary technology takes place on a large scale (Brynjolfsson, et al., 2017). For example, consider the computer productivity paradox that occurred in the late 1900s. The widespread usage of computer networks provided a substantial positive impact on productivity statistics (Atrostic & Nguyen, 2005). Van Ark (2016) supports this idea and argues that the new digital economy is still in its "installation phase" and productivity effects may only occur once it has been deployed on a much larger scale.

2.4 Cross-Country Comparison Issues with Artificial Intelligence

The effects of artificial intelligence implementation may be different depending on the size of a country's economy. Nations which invest large quantities of capital on artificial intelligence should yield higher levels of artificial intelligence implementation. However, measurement issues occur in cross-country comparison because smaller countries may have lower levels of total investment yet have equal or even greater levels of artificial intelligence adoption when compared to their larger counterparts. Therefore, it is important to account for the different sized economies in the analysis of artificial intelligence implementation.

Economists have devoted a substantial amount of effort to evaluating aggregate productivity in different countries since 1975 (Lafuente, et al., 2019). Since every country is different, and every country has their own economic policies, TFP discrepancies must be evaluated based on each individual country. As a result, the cross-country comparison becomes more difficult. However, economists have utilized different methods to overcome the cross-country analysis issues.

Albarran, Inklaar, and Woltjer (2019) find an issue with the assumption of homogeneity that is made in aggregate cross-country comparison. Their research focuses on the idea that different types of capital have different marginal products and cannot be assumed as homogenous. More specifically, they suggest that wealthier nations acquire capital with a higher marginal product than nations with less wealth. Put differently, wealthier nations are more likely to invest in artificial intelligence than developing nations and therefore spending on artificial intelligence is not random. They overcome this issue by creating a slightly adjusted measurement of TFP from the Penn World Tables that can account for more cross-country variation in income levels.

Eberhardt and Teal (2010) run into a similar issue with cross-country comparison. Specifically, when comparing aggregate productivity averages across countries, TFP does not account for the heterogeneity that exists between them. Once again, treating all of the countries in the same way is does not appropriately account for their individual differences. Eberhardt and Teal (2010) attempt to take into account parameter differences among countries and are able to obtain a more accurate measure of technology's effect on productivity.

Ultimately, potential issues may arise because countries are not randomly assigned to large investment in artificial intelligence. To illustrate this, assume a country would normally have productivity growth of 2%. That country chooses to invest heavily in artificial intelligence and only has 1% productivity growth in the next year. We would conclude that the investment in artificial intelligence was presumably unwise, but we can only draw that conclusion if we are certain that the country would otherwise have been on the same growth track as others. This last assumption cannot be made because it is not random which countries invest heavily in artificial intelligence. Instead, we have to determine an appropriate measurement for cross-country comparison that provides reliable results. The $RTFP^{NA}$ variable is a modified version of TFP that is used by Letta and Tol (2018). Their findings suggest that this $RTFP^{NA}$ variable is more reliable than typical TFP and can be used in cross-country comparison.

My research differs from previous studies for a number of reasons. Albarran, Inklaar, and Woltjer (2019) look at the cross-country productivity comparisons with different compositions of capital. Although my research utilizes a similar methodology, I focus primarily on artificial intelligence whereas Albarran, Inklaar, and Woltjer (2019) examine capital stock estimates. Eberhardt and Teal (2010) focus on the manufacturing industry in their cross-country TFP comparisons instead of country adoption of new technologies as a whole. My subject matter is closest to the paper by Brynjolfsson, Rock, and Syverson (2017), but instead of focusing on the productivity paradox my research is concerned with the short-term effects of artificial intelligence on TFP for a given country.

2.5 Further Implications of Artificial Intelligence Investment on Foreign and Economic Policy

Economic literature has shown that policy will shape how artificial intelligence impacts society (Agrawal, et al., 2019). Revisiting the idea of artificial intelligence as a general-purpose technology, policy decisions will have large scale implications for artificial intelligence impact across a variety of different industries. Optimists believe that artificial intelligence will change the world for the better, and the positive impacts of the implementation of new technologies greatly outweigh any potential negative spillover effects. Pessimistic viewers believe that artificial intelligence could lead to a handful of companies dominating society with increased inequality and few jobs left for human workers. These contrasting opinions both have weight, and policy appears to be both the connection and the guide for the path artificial intelligence will take.

The two dominant investors in the artificial intelligence space are the United States and China. While the United States has taken a largely "hands off" approach to artificial intelligence development, directing most of the development to private companies, China has made it a priority and even provided government funding for the new technology (Agrawal, et al., 2019). These two different approaches may have alternative effects on the success of artificial intelligence technology. Santos and Qin (2019) claim that China dominates global artificial intelligence funding and look like they will lead the artificial intelligence space in several areas.

This is largely due to the strong relationship between artificial intelligence development and the Chinese government.

The protection of Chinese firms by the Chinese government has made it difficult for major United States firms such as Amazon and Google to enter Chinese markets (Goldfarb & Trefler, 2018). This protection combined with limited intellectual property laws have allowed Chinese companies such as Alibaba, Tencent, and Huawei to dominate Chinese markets and surge to the forefront of global artificial intelligence use and development. Moreover, in 2017, the United States participated in fewer artificial intelligence conferences than they did in 2012 while China participated in 13% more (Goldfarb & Trefler, 2018). This trend is occurring when it is said that data and information are the new oil, whoever handles the data handles the emerging future of the global economy (Bonsu, 2020). The "hands off" approach that is taken by the United States may have a negative impact on their position in the global economy if Chinese artificial intelligence investment and growth continues to dominate the international sphere.

The possible spillover effects between countries are also affected by their policies. If China was more open about its technology, there would be opportunities for further innovation throughout the world and more widespread adoption. Although this may initially have a negative effect for the Chinese Communist Party as others use their technology to move their own countries forward, it may be advantageous in the long run. Today, Chinese artificial intelligence companies lack household recognition outside of China (Goldfarb and Trefler, 2018). By creating more open communication and expanding technological innovation beyond its borders, China may be able to benefit its own companies while boosting the artificial intelligence adoption in other countries throughout the world.

Agrawal, Gans and Goldfarb (2019) suggest that there are two main aspects of policy involving artificial intelligence implementation. The first is intellectual property policy. Intellectual property policy will influence the diffusion of artificial intelligence in different industries through privacy, trade, and liability. Data privacy has become a major concern for many people throughout both the United States and the world. Artificial intelligence, especially machine learning, is a largely data driven entity. Because of the need for large amounts of data, policy surrounding data privacy is brought into question. Although data privacy is important, government-mandated privacy regulation may lead to slower technology adoption and less innovation (Goldfarb & Tucker, 2011). Policy decisions regarding data privacy and corresponding regulations will have an impact on the artificial intelligence diffusion throughout an economy which will in turn affect aggregate productivity.

Intellectual property policy will have a substantial effect on global trade and the international impact of artificial intelligence. Privacy policies vary across countries. For example, Chinese data privacy policies are vastly different from those in the United States. As a result, data available to Chinese artificial intelligence researchers and producers may not be available to Americans (Goldfarb & Trefler, 2018). Therefore, the artificial intelligence improvements that occur in one country may not be possible in another. In future trade agreements, it is possible that artificial intelligence will spark a conversation regarding international privacy standards (Agrawal, et al., 2019). These trade policy decisions could change the effect of artificial intelligence throughout the world.

The effect of liability policies on artificial intelligence diffusion and implementation may be less obvious than privacy and trade. To illustrate this concept, consider a machine learning algorithm that is used to determine rates on auto insurance. This is a biased algorithm that gives more weight to credit score than driving history, and as a result, wealthy, dangerous drivers are given lower rates than poor, safe drivers (O'Neil, 2016). When this method of auto insurance calculation is uncovered and deemed unfair, who takes the blame? The insurance company that is using it, or the creators of the machine learning algorithm? Another example is the self-driving technology of automobiles. Tesla has launched completely autonomous self-driving cars that are being tested throughout the United States. If an autonomous Tesla kills someone crossing the street, who is blamed? The human "driver" or the Tesla company itself? These questions need to be answered with concrete liability policy in order for large scale diffusion to continue. Investors may be worried about these liability implications and could be deterred from the investment into artificial intelligence (Agrawal, et al., 2019). Decreased investment in artificial intelligence may slow its overall adoption in countries throughout the world.

The second main aspect of policy presented by Agrawal, Gans and Goldfarb (2019) are labor and antitrust policies. Labor and antitrust policies are important for the consequences of artificial intelligence in employment, inequality, and competition. Employment is one of the primary concerns when it comes to artificial intelligence implementation. When replacing labor with capital, specifically artificial intelligence, economies are reducing the number of jobs available to their labor force. As a result, many economists and government officials are concerned with a potential rising unemployment rate as a result of increased artificial intelligence implementation. However, artificial intelligence optimists believe that the new technology may substitute existing jobs and initially increase unemployment, before creating new jobs that were previously non-existent and offer opportunities for increased labor productivity. In other words, increasing technological efficiency should lead to greater labor demand in the long run. With this argument, nations should not be deterred from investment in artificial intelligence technology because they will increase their labor demand and decrease their unemployment. Qiulin, Duo and Yi (2019) find that current artificial intelligence development has supplemental substitution properties rather than crowding-out substitution properties. In other words, the positive productivity and labor effects that occur because of artificial intelligence implementation are greater than the initial employment loss. Economists warn that the long run positive effects on employment due to artificial intelligence may be a long way down the road, which is important for policy makers to keep in mind (Agrawal, et al., 2019).

The inequality aspect of policy making is based largely on the implementation of high-level technology in general. Utilizing artificial intelligence greatly benefits those who know how to use it. This implies that educated and already wealthy individuals will receive the greatest returns from large scale artificial intelligence implementation. This could increase the wealth gap and emphasize the disparity between an economy's wealthy and poor populations. Moreover, the jobs that will be replaced by artificial intelligence first will be low-income positions which require manual labor. Not only will the implementation of artificial intelligence benefit the educated and wealthy, but it may also force poor, uneducated members of the economy out of work and create further inequality.

Finally, policy regarding competition in artificial intelligence investment is critical to avoid the development of monopolies and few companies dominating society (Agrawal, et al., 2019). Leading artificial intelligence companies dominate the market for this new technology and have recently been under scrutiny for their immense revenues and market control. In the United States, companies such as Google, Amazon, Apple and Facebook have all been under scrutiny for their monopolistic characteristics and the markets they dominate. The policy decisions made could have substantial impacts on future artificial intelligence implementation and market performance.

CHAPTER 3

DATA SOURCES AND DESCRIPTION

In this section, I describe my data sources and how I create my dataset. Next, I provide information on my independent variables used to capture the level of artificial intelligence implementation in a given country. Finally, I discuss my dependent variable, $RTFP^{NA}$, and its construction before presenting the summary statistics.

3.1 Data Sources

In my research I attempt to measure the effect of artificial intelligence implementation on total factor productivity. The data I observe for this analysis come from two primary sources: the Stanford Human-centered Artificial Intelligence (HAI) database and the latest edition of the Penn World Tables (PWT). The Stanford HAI data was used to develop the 2019 Stanford HAI index report, which includes relevant data on artificial intelligence implementation from 27 different countries. The research is broken up into eight sections in an attempt to gauge all relevant aspects of artificial intelligence implementation. The eight sections are research and development, conferences, technical performance, economy, education, autonomous systems, public perception, and societal considerations. The Stanford HAI group uses these data to develop an index report and score for each individual country in the dataset (HAI Index Report, 2019). The data has been gathered from a variety of databases and vetted by over 150 academic and industry experts. The mission of the Stanford HAI program is to, "provide unbiased, rigorous, and comprehensive data for policymakers, researchers, journalists, executives, and the general public to develop a deeper understanding of the complex field of artificial intelligence"

(HAI Index Report, 2019). The Stanford HAI has made the data they use to build their own artificial intelligence index public and available for download.

Data on productivity comes from the Penn World Tables. The latest edition of the Penn World Tables was released on January 28th, 2021 and includes a wide range of macroeconomic variables for many different countries. The Penn World Tables has been a highly regarded database for major macroeconomic variables for over 40 years and offers information on 183 countries between 1950 and 2019 (Feenstra, et al., 2018). Run by the University of Groningen, these data are consistently updated and offer an accurate measure of country macroeconomic statistics. The data published by the University of Groningen Economic Growth Centre is publicly available for download.

The dataset I am using is panel data. It is made up of 27 countries from 2015-2018 yielding four years' worth of data and 108 data points. Ultimately, I end up with 26 countries over four years, and 104 total data points, due to incomplete data from the United Arab Emirates. These countries were selected by the Stanford HAI team because they offer a unique cross-section of countries in regions across the world. The countries also represent different income levels from high income to lower middle income, although most countries included are high income countries. The technology itself is expensive and as a result is more easily adopted by high income countries. The Stanford HAI group developed its first HAI index report in 2017 and used data dating back to 2015. The data from the Penn World Tables is also available for this time period. The Penn World Tables has more of both countries and years, but I am limited to the countries and years included in the Stanford HAI dataset.

The dataset I use for my regressions contains elements of both aforementioned data sources. In order to determine the impact of artificial intelligence implementation on TFP, I obtain variables that are related to artificial intelligence from the Stanford HAI dataset and TFP from the Penn World Tables. When combining these two datasets, I create cross-sectional data on both productivity and artificial intelligence implementation. The variables described below are available for each country and year in my dataset and combine into a complete dataset with no missing values.

3.2 Independent Variables

The independent variables on artificial intelligence I use in my regressions are captured by the Stanford HAI database. Since there is no single clear variable to represent artificial intelligence implementation, I use the number of artificial intelligence journal citations, number of artificial intelligence conference citations, share of enrollment in artificial intelligence focused courses, artificial intelligence private investment, and the number of start-ups focused on artificial intelligence. Together, these variables offer a comprehensive outlook on a country's overall artificial intelligence implementation. The rest of this chapter will contain individual variable descriptions for each of my independent and dependent variables.

The first independent variable I use to capture artificial intelligence implementation is the number of journal paper citations. This independent variable estimates the number of times a paper published in an academic journal was cited. For example, the United States had 48,175 papers that were published by an American academic journal cited by other authors in 2018. I chose this variable instead of the total number of journal papers published because I wanted to ensure that I was measuring the quality of work instead of just the quantity. With this variable, I am able to account for the quality of papers, and therefore research, published in academic journals for each given country. This falls under the Research and Development section of the Stanford HAI Index and was obtained by Stanford from the Microsoft Academic Graph.

The second independent variable used in my regressions is the number of conference citations. By estimating the quality of conference publications, this variable helps capture the value of attending artificial intelligence conferences. Artificial intelligence focused conferences are held around the world and it is believed that countries who attend more conferences and create higher quality work will move further along in artificial intelligence implementation. The United States has the highest number of conference citations with 42,160 citations in 2018. Conferences is its own section under the Stanford HAI which emphasizes its importance in overall artificial intelligence development. This variable was also obtained by Stanford from the Microsoft Academic Graph.

The third independent variable I use to capture artificial intelligence implementation is the enrollment share. The enrollment share variable was collected by Stanford from Coursera and estimates the fraction of a country's enrollments that are in courses teaching Artificial Intelligence and related skills in order to gauge the relative interest in artificial intelligence content across the world. Measured over time, this variable shows enrollment trends and where the emphasis on artificial intelligence is increasing or decreasing. The enrollment share variable is categorized under the Economy section of the Stanford HAI database because it evaluates the interest in gaining skills in artificial intelligence. France has the highest enrollment share among the 27 countries measured in my dataset.

The fourth independent variable, which also falls under the economy section of the Stanford HAI dataset, is the total private investment in artificial intelligence. This variable measures the total amount of private investment funding received for artificial intelligence focused startups in United States dollars. Capturing the amount of private investment in artificial intelligence is valuable for a number of reasons. The first reason is that it offers a strong indicator of the interest of a country in artificial intelligence implementation. Theoretically, higher private investment in artificial intelligence should be associated with an increase in overall artificial intelligence use and implementation because increased investment typically means increased research and production. However, measuring private investment in artificial intelligence also adds another level of comparison between countries. While some countries, like the United States, leave much of their artificial intelligence development and implementation to the private sector, other countries, like China, heavily subsidize companies focused on artificial intelligence implementation and therefore contribute less money from private investment. The United States has the highest value for this variable with over $18.7 billion in private investment in 2018. Stanford reports obtaining the measures for this variable from Crunchbase, CapIQ and Quid.

The fifth independent variable I use in my regression is the number of startups focused on artificial intelligence. This variable measures the total number of artificial intelligence companies founded in a given country. With more companies focused on artificial intelligence there should be more research, development, and production and therefore higher levels of overall artificial intelligence implementation. Similar to total artificial intelligence private investment, this variable also has political implications. The ease of starting a business depends on the country. Therefore, if there is a positive and significant correlation with TFP, it would show that countries which have more new technology startups are more productive. The United States has the highest number of artificial intelligence startups with 1,480 companies in 2018. In per capita terms, Israel is the highest in 2018. These values also come from Crunchbase, CapIQ, and Quid.

The Stanford HAI data also includes per capita measures for each of the above independent variables. Because the countries I am comparing can be vastly different sizes, it will be important to use per capita variables to account for the size difference during comparison. The per capita variables are all calculated with the population in millions as the denominator.

3.3 $RTFP^{NA}$ and its Construction

My dependent variable is $RTFP^{NA}$ and comes from the Penn World Tables (PWT) version 10.0. $RTFP^{NA}$ stands for Real Total Factor Productivity from National Accounts Data. Specifically, this is a measure of TFP at constant national prices. It is a country-specific index of TFP with a benchmark year of 2017 and is used to measure within-country productivity growth over time (Letta & Tol, 2018). When using the logarithmic transformation of my $RTFP^{NA}$ variable, the benchmark value for all countries is zero. $RTFP^{NA}$ is a calculated variable and is not observable. Therefore, it is possible that there is ambiguity in the root cause of a rise or fall in TFP.

$RTFP^{NA}$ is calculated slightly differently than typical TFP. Instead of allowing TFP to capture the entire residual growth rate of output that is not explained by the growth of inputs, the $RTFP^{NA}$ variable adjusts for the effect of inflation which in turn captures broader changes to the economy as well. By keeping prices constant throughout measurement, $RTFP^{NA}$ is able to measure within-country productivity growth over time while removing outside influences that may have previously been included in the TFP calculation. As a result, the productivity increase that is represented by the $RTFP^{NA}$ variable should be less biased and more accurate when evaluating the effects of technical progress. The following explanation of $RTFP^{NA}$ is based on the appendix from the paper *Weather, Climate, and Total Factor Productivity* by Letta and Tol.

They explain the $RTFP^{NA}$ calculation from edition 8.1 of the Penn world Tables. Version 10.0 has updated data and uses 2017 as the benchmark year, but otherwise measures these values in the same way.

Construction of the $RTFP^{NA}$ variable

$RTFP^{NA}$ is calculated using the real GDP growth rate as well as the growth rates of capital stock and the labor force (Feenstra, Inklaar, and Timmer, 2013). The start of this calculation comes from the general Cobb-Douglas production function:

$$Y = A * f(K, L) = A * K^{\alpha}(E * hc)^{1-\alpha} \tag{3}$$

Where *Y* is output, *K* is capital and *L* is labor. *L* is defined by a product of the number of workers in the economy, *E*, times their average human-capital *hc*. Capital is raised to the power of α and labor input is raised to the power of 1-α. α is the output elasticity of capital and 1-α is the output elasticity of labor. Since the exponents in the Cobb-Douglas function add up to one, we assume constant returns to scale. A second-order approximation to the production function *f* is represented by the Törnqvist quantity index of factor inputs Q^T, which can be used to compare inputs between *t-1* and *t* for a given country as follows (Letta and Tol, 2018). Since we know how much labor grows and how much capital grows, we attempt to find Q^T.

$$\ln Q^T_{t,t-1} = \frac{1}{2}(\alpha_t + \alpha_{t-1})\ln\frac{K_t}{K_{t-1}} + \left[1 - \frac{1}{2}(\alpha_t + \alpha_{t-1})\right]\ln\frac{L_t}{L_{t-1}} \tag{4}$$

In the above equation, Q^t represents the total inputs for a given country, $\ln \frac{K_t}{K_{t-1}}$ represents the percent growth in labor inputs from time *t-1* to time *t*, and $\ln \frac{L_t}{L_{t-1}}$ represents the percent growth in capital inputs from time *t-1* to time *t*. This equation is necessary because even though we have levels of α from our previous equation, α varies among countries and may change over time. As a result, we add a time subscript and weigh K by α and L by $1-\alpha$. In the above equation, this weighting is performed by taking the average weight of labor and capital growth between *t-1* and *t*. From equation (4), we can learn how much inputs grew in a given year.

To implement equation (4), the assumption that output elasticity of capital is approximated by the country's share of GDP is not earned by labor is made (Letta and Tol, 2018). Finally, growth in productivity over time (TFP) is given by:

$$RTFP^{NA}_{t,t-1} = \frac{RGDP^{NA}_t}{RGDP^{NA}_{t-1}} / Q^T_{t,t-1} \quad (5)$$

Where $RGDP^{NA}$ stands for real GDP at constant national prices. The right-hand side of this equation calculates how much GDP increases/decreases with a country's total inputs in a given year. $RTFP^{NA}$, therefore, is a measure of how much productivity grows over time. If outputs grow faster than inputs, than productivity increases and $RTFP^{NA}$ is greater than one. If productivity remains the same, $RTFP^{NA}$ will be equal to one. If productivity declines, then $RTFP^{NA}$ will be less than one.

Table 1 contains summary statistics for all variables I use in my regression in original units and per capita terms. Per capita statistics are indicated with the "pc" appendix on their variable name.

3.4 Summary Statistics

Table 1

	count	mean	sd	min	max
numjournalcitations	108	15130.76	27689.25	358	165576
numconfcitations	108	11857.96	28889.42	55	200790
enrollshare	108	.0599758	.0327491	.0090137	.1323639
aiprivateinvest	108	8.12e+08	2.77e+09	177496	1.87e+10
numstartups	108	77.33333	216.257	2	1480
numjournalcitationspc	108	.000368	.0004518	4.82e-06	0.0027993
numconfcitationspc	108	.0002592	.0003434	9.52e-07	.0021583
enrollshare	108	.0599758	.0327491	.0090137	.1323639
aiprivateinvestpc	108	8.815697	18.51008	.0083655	117.5569
numstartupspc	108	1.65e-06	2.26e-06	9.70e-09	.0000147
rtfpna	104	1.009796	.0515582	.865912	1.212405

Total Factor Productivity is a calculated measure for each country and does not need to be converted to per capita values for my regression results.

CHAPTER 4

ECONOMIC MODEL AND METHODOLOGY

In this chapter I present the economic model and econometric relationship I intend to use for my regressions. I explain the need for including both per capita and logarithmic models in my analysis and provide justification for utilizing fixed effects. Finally, I discuss potential issues of endogeneity and how I plan to overcome their negative impact.

4.1 Economic Model

The economic model I create follows the methodology set forth by Letta and Tol (2018). In their paper, they utilize $RTFP^{NA}$ as the dependent variable and create a regression model. I will use an OLS regression model with both country and time fixed effects in order ensure the robustness of my conclusions. The literature has suggested that there will be a productivity paradox, i.e., the anticipated positive effects of artificial intelligence implementation will not be seen in the productivity statistics. Although a productivity paradox is expected, my results may differ because of the unique aspects of artificial intelligence implementation captured by my dataset and increasing technology spillover effects.

Implementing artificial intelligence technology should increase the overall productivity of countries across the world because of the effects of new technology on productivity. Aghion, Jones, and Jones (2017) present the production function below:

$$Y_t = A_t * K_t^{\alpha} L_t^{1-\alpha} \tag{6}$$

Where Y is output, A is technology, K is capital, and L is labor. Essentially, this production function shows that as technology improves (A increases), output should increase even with existing labor and capital inputs. Since artificial intelligence implementation is seen as a technological improvement, implementing more artificial intelligence should yield higher levels of productivity. Capturing the impact of technological growth is more complicated than the equation above, and as a result total factor productivity was created to measure technical progress in a given country. A modified version of TFP is the dependent variable in my economic model.

One issue that may arise with the implementation of artificial intelligence is the dual impact that it has on both technology and capital. When artificial intelligence implementation occurs, part of the effects of the new technology will be absorbed by the increased capital that is developed as a result. However, this increase in capital is different from building more factories or roads and artificial intelligence will increase technology more than other forms of capital investment. An increase in capital that is related to technology should drive up productivity and will mostly be captured by my TFP variable even if there is some spillover into the capital portion of the production function.

I use total factor productivity as my dependent variable instead of GDP or other growth measures because it isolates the effect of technology on overall growth. When a nation is attempting to improve its productivity in the long term, it is important to measure their investment using TFP instead of their overall growth numbers. If a country is spending large amounts of money on building coal-burning factors and infrastructure, their overall unemployment numbers decrease and their output increases, leading GDP and other growth indicators to rise. However, investment in these areas do not increase productivity at the same level as technological investment in the long term. Because an increase in capital changes overall output but does not affect TFP, using TFP as the dependent variable in my model captures the way in which artificial intelligence investment changes technology, and not the way that it changes output by increasing capital in the economy. Therefore, using TFP to measure productivity growth due to artificial intelligence implementation is appropriate and will effectively estimate productivity growth.

The standard economic model for my regression is shown below:

$$TFP = \beta_1 numjournalcitations + \beta_2 numconferencecitations + \beta_3 enrollmentshare + \beta_4 aiprivateinvestment + \beta_5 numstartups + \varepsilon$$

Estimation by least squares requires assuming that the independent variables are exogenous. An increase in artificial intelligence implementation, i.e. the right hand side of my equation, should increase TFP for the reasons mentioned above. However, there is reason for some concern about endogeneity in my right-hand side variables. When estimating TFP I have attempted to include all relevant independent variables available. Because artificial intelligence implementation is difficult to measure, it is naïve to think that none of the variables included in my regression have any correlation with the error term. That is, an increase in the error term, meaning an increase in something not included in my regression affecting total factor productivity, may be correlated with the number of journal citations, etc. Also, the error term in this case is not observable which makes it challenging to know whether my independent variables are endogenous or exogenous. Although endogeneity is a potential concern, I continue with my research and include one and two-way fixed effects models in an attempt to absorb the potential endogeneity. The fixed effects absorb the endogeneity that comes from different countries making different types of investments.

Utilizing fixed effects models are one way to deal with endogeneity issues. Fixed effects models capture the differences in the dependent variable associated with each unit and/or time period (Bailey, 2017). Fixed effects also allow different units to have different baseline levels of my TFP variable. I run a one-way fixed effects models for time, attempting to control for the variation that occurs in TFP growth in different years. As a result of the fixed effects, a different value of α will be used for each year. The economic model with one-way fixed effects is below.

$$TFP = \beta_1 numjournalcitations + \beta_2 numconferencecitations + \beta_3 enrollmentshare + \beta_4 aiprivateinvestment + \beta_5 numstartups + \alpha_i year + \varepsilon$$

I also create two-way fixed effects models for both year and country as well as year and region. Using two-way fixed effects accounts for the variation in both time and location (either country or region) in an attempt to further eliminate endogeneity that may occur with my independent variables. While including fixed effects for country, I am eliminating potential endogeneity by comparing within countries and seeing if they have higher TFP in years with higher levels of artificial intelligence implementation. Including fixed effects for region, like country, allows me to account for potential bias that may occur within regions that are correlated with an independent variable. When using only region fixed effects, I am comparing countries to one another within a region but not across regions. This eliminates omitted variable bias that could come from comparing Germany to Japan, but not the bias that could come from comparing Germany to France. Fixed effects models often eliminate some significance in regressions because when we remove fixed effects from the error term, they can no longer be a source for the correlation with either the independent or dependent variables. I can use two-way fixed effects because I have panel data. The two-way fixed effects models are below.

$$TFP = \beta_1 numjournalcitations + \beta_2 numconferencecitations + \beta_3 enrollmentshare + \beta_4 aiprivateinvestment + \beta_5 numstartups + \alpha_i year + \alpha_j country + \varepsilon$$

$$TFP = \beta_1 numjournalcitations + \beta_2 numconferencecitations + \beta_3 enrollmentshare + \beta_4 aiprivateinvestment + \beta_5 numstartups + \alpha_i year + \alpha_j region + \varepsilon$$

In addition to the fixed effects approach I use to eliminate potential endogeneity issues; I also have a comparison issue between the size of different countries. The United States, for example, had a population of approximately 327 million in 2018. This is much larger than Israel, which had a population of approximately 8.4 million at the same time. The size difference contributes largely to the number of journal and conference citations, number of startups, and overall artificial intelligence investment. This size difference makes these two countries almost impossible to compare and I therefore need to utilize per capita values of each variable. A per capita approach makes it possible to see if artificial intelligence implementation is actually assisting in productivity growth instead of the size and wealth of a given country causing the increase.

Letta and Tol (2018) use an alternative approach when making cross-country comparisons. Instead of using per capita estimations, they use a logarithmic construction in which a logarithmic transformation is used on the dependent variable and the independent variables remain linear. I follow this approach for my dependent variable but also use a logarithmic transformation on my independent variables for ease of comparison and interpretation. This model allows me to estimate the percent change in TFP that is associated with a percent change in any one of my independent variables. Because the regression measures the percent change in my dependent variable that is associated with the percent change in my independent variable, the log-log construction is useful in comparing countries of different sizes. The logarithmic construction of my economic models is presented below.

$$\begin{aligned} logTFP = {} & \beta_1 lognumjournalcitations + \beta_2 lognumconferencecitations \\ & + \beta_3 logenrollmentshare + \beta_4 logaiprivateinvestment \\ & + \beta_5 lognumstartups + \alpha_i year + \alpha_j country + \varepsilon \end{aligned}$$

$$\begin{aligned} logTFP = {} & \beta_1 lognumjournalcitations + \beta_2 lognumconferencecitations \\ & + \beta_3 logenrollmentshare + \beta_4 logaiprivateinvestment \\ & + \beta_5 lognumstartups + \alpha_i year + \alpha_j region + \varepsilon \end{aligned}$$

When creating a regression model, it is important to make sure the causality is correct in its construction. My economic model suggests an increase in artificial intelligence implementation causes a change in total factor productivity and the causality only works in one direction. When artificial intelligence implementation occurs, technology as a whole improves and there is a larger residual between total output and outputs produced by capital and labor. Therefore, it is expected that an increase in artificial intelligence implementation and the improvement of technology increases TFP. It is important to ensure that the reverse causality is not true. In other words, that artificial intelligence implementation is not increasing due to a rise in TFP. Since TFP captures productivity due to technological progress, there is no reason to believe that this calculated measure of productivity causes higher levels of artificial intelligence implementation within one year. I can safely proceed with both my logarithmic and per capita economic model constructions for my research.

CHAPTER 5

REGRESSIONS AND RESULTS

Since the countries in my dataset are vastly different sizes, I attempt to compare them using two different methods. I use a per capita and logarithmic approach due to the discrepancy between country size, GDP, etc. A similar methodology was used by Letta and Tol (2018) for cross-country comparison. I present the results of both estimation methods and interpret the coefficients.

5.1 Per-Capita Regressions and Results

I estimate four different specifications. The dependent variable in each specification is $RTFP^{NA}$ (Real Total Factor Productivity using constant prices from national accounts data) as presented by the 10^{th} version of the Penn World Tables published by the University of Groningen Growth and Development Center. The number of journal citations, number of conference citations, enrollment share (fraction of students that are enrolled in courses with Artificial Intelligence), total private investment in artificial intelligence, and the number of artificial intelligence focused startups are all recorded from the Stanford University HAI Institute. I use per capita estimates because it is difficult to compare countries of varying sizes and populations. Table 3 shows the results.

In the first specification I regress $RTFP^{NA}$ on the independent variables without any fixed effects. The coefficients on artificial intelligence private investment and number of startups variables are significant at the 5% level. This indicates that countries which have higher private investment in artificial intelligence have slightly lower TFP. Although this specification presents significance in two of my independent variables, I add the time fixed effects in my second specification to reduce problems of omitted variable bias. The time fixed effects remove variation that occurs over time and instead compares country TFP within a given year. These results remove some of the potential endogeneity that may occur across years.

Once I include one-way fixed effects for year, both artificial intelligence private investment per capita and the number of startups per capita remain statistically significant. Holding all else constant, a one dollar increase in private artificial intelligence investment per capita is associated with a decrease in total factor productivity of 0.00143. The coefficient on the number of startups per capita indicates that a one unit increase in the number of startups per million people is associated 0.013928 increase in total factor productivity. These results are consistent with the previous model in which no fixed effects were used. The F-statistic for this specification is 8.39, yielding a p-value less than 0.01, and suggesting that time effects are statistically significant.

In the third specification I include two-way fixed effects for both year and country. Including the two-way fixed effects in the third specification helps avoid the risk of bias that occurs within countries that happen to be correlated with an independent variable. Using two-way fixed effects helps as long as the correlation comes from connections across countries, and not from connections across years. By including these in my model, I am able to account for historical tendencies within different countries.

Table 3

Results of Per Capita Specification

	(1)	(2)	(3)	(4)
	rtfpna	rtfpna	rtfpna	rtfpna
numjournalcitationspc	-55.22*	-49.62	-19.19	-87.59***
	(30.99)	(31.31)	(29.99)	(26.61)
numconfcitationspc	36.75	20.85	2.652	36.54
	(40.72)	(42.30)	(30.35)	(34.33)
enrollshare	-0.0394	0.156	-0.188	-0.292
	(0.161)	(0.216)	(0.276)	(0.187)
aiprivateinvestpc	-0.00163**	-0.00143**	0.000252	-0.00161***
	(0.000641)	(0.000651)	(0.000574)	(0.000503)
numstartupspc	14523.4***	13927.6**	1922.3	23388.0***
	(5465.6)	(5484.5)	(7053.9)	(4828.5)
Year FE	No	Yes	Yes	Yes
Country FE	No	No	Yes	No
Region FE	No	No	No	Yes
Observations	104	104	104	104
Mean of Dep. Variable	1.010	1.010	1.010	1.010

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Because the unit of observation is a country in a given year, fixed effects encompass factors related to total factor productivity in each country that do not vary over time. It is important to note that using country fixed effects means I am identifying the effect of artificial intelligence implementation on TFP by comparing within countries and identifying whether they have higher TFP in years with higher levels of artificial intelligence implementation. I am no longer comparing to see whether countries with more artificial intelligence have higher TFP. Losing the latter source of identification removes the possible omitted variable bias that occurs in my model but also means that I have less ability to detect the effects of artificial intelligence implementation on TFP. Once I control for the fixed effects in the model of specification three, the effects of both private investment in artificial intelligence and the number of startups focused on artificial intelligence become insignificant. This indicates that neither investment in artificial intelligence nor increasing the number of startups has a positive effect on total factor productivity.

The fourth specification, like the third, includes two-way fixed effects. However, the fixed effects are for both year and region instead of year and country. In this specification, the coefficients on private investment in artificial intelligence and the coefficient on number of startups focused on artificial intelligence are both significant. A one dollar increase in private investment in artificial intelligence per capita decreases total factor productivity by 0.00161. The increase of one startup per million people is correlated with a total factor productivity increase of 0.023388. Also, the coefficient on the number of journal citations per capita is significant for the first time in any specification. A one unit increase in the number of journal citations per million people would decrease total factor productivity by 0.000088.

The two-way fixed effects model in the third and fourth specifications are the most important because the year and location effects help account for potential endogeneity. In the third specification, when both year and country fixed effects are included, the significance on all independent variables disappears at both 5% and 10% level. This lack of significance could be a result of the relatively small number of years in my dataset. With country fixed effects, we are looking at a given country in one year compared to the same country in the previous or following year. This affects my significance levels because artificial intelligence investment and TFP do not vary much from year to year and vary more from country to country.

Another reason there may not be any significant variables in the third specification could be that I am seeing the effects of the productivity paradox. Although I expect there to be an increase in productivity when artificial intelligence implementation increases, the lack of significance on any independent variable indicates that there may not actually be an effect. This is consistent with the findings presented in previous literature as many economists and researchers have failed to find the influence of artificial intelligence on overall productivity. Again, this could be because of the small number of years in my dataset. It is common to use very large datasets that span over large time periods when performing analysis with TFP.

The fourth specification which includes fixed effects for region and year has three significant variables: number of journal citations, the number of startups, and total private investment. The significant negative coefficient on the number of journal citations per capita is unexpected. This could be because producing high quality research has no direct economic effects. Also, money put into university research is not available for other uses which might have more of an immediate impact and the benefits of university research will have a time lag. Therefore, the positive effects of high-quality research are largely spillover effects, meaning the research done in a given country produces an opportunity for businesses to implement findings into their markets. However, published research also provides the opportunity for technology firms in other countries to take advantage of findings that their own country has not produced. The spillover effects would largely be positive, because countries not doing research on artificial intelligence would benefit from the country that did. Also, and perhaps the more likely answer, academic research that is published in academic journals does not produce any products or investment. Therefore, the high-quality research papers that are published and cited are important but do not actually contribute to the productivity of a given country as a whole.

The negative significant coefficient on artificial intelligence private investment per capita is also relatively surprising. One would think that increasing investment in technology ought to yield higher levels of productivity. However, this effect is not immediate and may take time for the investment to pay off. With panel data of four years, I am unable to see the effects of an investment over the long term. The initial effect may be negative because the money that is spent on artificial intelligence could have been spent in another industry that pays off more quickly. Using this dataset, we are only able to see the negative effects of the initial investment instead of the potential positive effects of the future investment.

The coefficient on the number of startups focused on artificial intelligence per capita is positive and significant. Unlike private investment, a startup typically has a much faster impact on the economy because of its role in product development. This variable measures the total number of startups focused on artificial intelligence that were funded in a given year. Because of its positive significance, it may have overcome the expected delay in productivity that we have seen throughout literature due to the productivity paradox.

5.2 Logarithmic Regressions and Results

Next, I discuss the results from the logarithmic functional form. I use a logarithmic construction for my specifications in Table 4. A log-log model determines elasticity, with a percentage increase in “x” causing a percentage increase/decrease in “y”. In the specifications in Table 4, I use a log-log structure to estimate regression equations with no fixed effects, one-way fixed effects for time, two-way fixed effects for both time and country, and two-way fixed effects for both time and region that a country is located.

Table 4

Results of Logarithmic Specification

	(1)	(2)	(3)	(4)
	logrtfpna	logrtfpna	logrtfpna	logrtfpna
lognumjournalcitations	-0.0268**	-0.0191	-0.0223	-0.0279**
	(0.0124)	(0.0130)	(0.0224)	(0.0125)
lognumconfcitations	0.0207*	0.00544	-0.0195	0.00342
	(0.0107)	(0.0134)	(0.0145)	(0.0125)
logenrollshare	-0.000163	0.0127	0.00336	-0.00287
	(0.00783)	(0.0112)	(0.0195)	(0.0108)
logaiprivateinvest	-0.00924*	-0.00664	-0.00372	-0.00133
	(0.00515)	(0.00527)	(0.00433)	(0.00448)
lognumstartups	0.0239***	0.0282***	0.0299**	0.0247***
	(0.00833)	(0.00870)	(0.0124)	(0.00797)
Year FE	No	Yes	Yes	Yes
Country FE	No	No	Yes	No
Region FE	No	No	No	Yes
Observations	104	104	104	104
Mean of Dep. Variable	0.00848	0.00848	0.00848	0.00848

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

In the first specification I regress the natural log of $RTFP^{NA}$ on the independent variables. Similar to the previous model with per capita estimations, the coefficients on both the number of journal citations and the number of artificial intelligence focused startups are significant at the 5% level. A one percent increase in the number of journal citations is correlated with a 0.0268% decrease in total factor productivity. For startups focused on artificial intelligence, a one percent increase in the number of startups leads to a 0.0239% increase in total factor productivity. The estimate for startups is consistent with the results obtained from the per capita model, suggesting that there is a positive correlation between the number of startups in a given country and their total factor productivity. The coefficient on the number of journal citations, however, is newly significant without fixed effects and shows that an increase in the number of journal citations actually decreases overall total factor productivity. Unlike the first specification in the per capita model, the estimate for private investment in artificial intelligence is insignificant at the 5% level.

After including the time fixed effects in the second specification, only the coefficient on the number of startups remained significant and became slightly larger. A one percent increase in the number of startups is associated with a 0.0282% increase in total factor productivity. Once again, the effect of this variable on total factor productivity is consistent with the first specification of my logarithmic model as well as the effects obtained from the per capita model.

Like the per capita model presented earlier, I use two-way fixed effects for year and country in the third specification. The two-way fixed effects model eliminates the significance of the number of journal citations, but the number of startups variable continues to be significant. The positive significance on the number of startups variable is robust and maintained throughout all of the logarithmic specifications. A one percent increase in the number of startups in a given country is associated with a 0.0299% increase in total factor productivity. Since this model has two-way fixed effects, this coefficient indicates that when more startups are funded in a given country grows from one year to the next, there is growth in overall productivity.

The coefficient on the number of startups is both robust and defies the productivity paradox. In each logarithmic specification, and in the fourth specification of the per capita model, the coefficient on the number of startups remains significant. This could have broader implications for governments and different countries around the world. Depending on the system of government in a given country it may be easier or more difficult to create and launch a startup. However, if increasing the number of startups produces higher productivity, it may be advantageous to shift policy so that startup implementation is easier. The fast-acting effects that occur in the number of startups variable implies that this may be the best area to focus for countries looking to rapidly boost productivity through artificial intelligence implementation.

The fourth specification is less important than the third specification but nonetheless contains meaningful information. Using fixed effects for both region and time is useful but does not carry as much weight as the fixed effects for both country and time used in the third specification. Similar to the fourth specification in the per capita model, the fourth specification in the logarithmic model has both the number of journal citations and the number of startups significant. The negative coefficient on the number of journal citations is most likely due to the lack of production that occurs with research alone. Unlike opening new startups, research does not contribute to production of or investment in artificial intelligence. The positive coefficient on the number of startups variable is robust and unsurprising given my previous results.

Although I have found the number of startups variable to be significant, it is interesting to consider why the other variables may not be. Research, increased investment, and higher enrollment share ought to produce higher levels of productivity. The productivity paradox that occurs with artificial intelligence may be in action with each one of these variables. One of the primary concerns with the productivity paradox is a time concern. That is, increased artificial intelligence implementation does not increase productivity because we have not given enough time for it to develop to the point that we see its effects in productivity measurements. For each one of these variables, the time lag explanation of the productivity paradox is an appropriate justification for their insignificance.

Both the number of journal citations and the number of conference citations measure the quality of research that is produced by a given country. The impact of research on a given country is not immediate. It takes time for research spillovers to occur and companies to implement new ideas into their products. Overall research and development are critical in future improvement of technology that will eventually increase productivity, but their delayed effects are not seen in my regression. With a larger dataset that includes many years' worth of data, the effects of research may be more apparent.

Private investment in artificial intelligence follows the same trends. Private investment in artificial intelligence technology assists in its development but the effects of the investment may not be seen for years. Money that is invested today helps companies with production, research and development, etc., in the future and it is possible there is a lag in productivity as a result. Unfortunately, with a short time span in my data, I am unable to see the longer-term effects of artificial intelligence investment on overall productivity. In some specifications, the coefficient on artificial intelligence private investment may be negative because the money that is spent on artificial intelligence implementation could be spent in other areas that have a faster impact on productivity.

The potential delayed effects of my enrollment share variable are easily explained. The enrollment share variable is defined as, "the fraction of a country's enrollments that are in courses teaching artificial intelligence and related skills to measure the relative interest in artificial intelligence content." Since this independent variable measures interest in artificial intelligence for a given country through education, it is safe to infer that the skills acquired and learned in the courses take time to be implemented into the workplace. Once the educated individuals reach the workforce and begin to make an impact there should be an increase in overall productivity.

The lack of significance presented above suggests that artificial intelligence implementation, for the most part, does not affect TFP in the short term. To answer the question of does artificial intelligence implementation affect TFP in the long term, I would need to use a different dataset. Research analyzing TFP growth typically has more than fifty years' worth of data and my dataset only contains panel data from four years. Unfortunately, there are not many years' worth of data available on artificial intelligence implementation because it is a new and developing industry. Although previous research on artificial intelligence uses a less comprehensive dataset, my research still exhibits aspects of the productivity paradox that is consistent with earlier findings.

When discussing the true cause of the productivity paradox in artificial intelligence, the positive significance on the number of artificial intelligence startups when using the logarithmic construction may offer support for the time-lag explanation. The significance defies the expected productivity paradox and suggests that increasing the number of startups focused on artificial intelligence will increase overall productivity of a country. As discussed earlier, I believe this is due to the fast impact a startup is able to have on an economy. It is important to note, however, that startups are often using technology that was developed a few years prior to the launch of the startup. Following economic beliefs that technology does increase overall productivity, the positive significance on the number of startups variable supports these beliefs and rejects the notion that artificial intelligence does not actually affect overall productivity. These findings lead me to conclude that that the productivity paradox is largely a result of the lack of time given for the implementation of artificial intelligence to appear in the productivity statistics. Since a delay of a few months would have been captured by my variables, my results suggest that it may be years before there is a major effect on TFP due to artificial intelligence.

CHAPTER 6

CONCLUSION

My research suggests that artificial intelligence implementation only contributes to the productivity statistics when it is tangible. In other words, we do not see productivity dividends from artificial intelligence when it is in the research phase and instead must wait until it is implemented by firms. More specifically, my analysis shows that the number of startups implemented by a given country has a positive effect on total factor productivity. I look at the effect of artificial intelligence implementation on a nation's overall productivity by performing a regression analysis. I have regressed total factor productivity on various measures of investment in artificial intelligence.

Previous literature does not find that artificial intelligence implementation has any effect on productivity and suggests there is a productivity paradox. A productivity paradox occurs when technology is rapidly improving but its improvement is not reflected in the productivity statistics. A productivity paradox occurred during the 1980s and 1990s and some economists believe it is occurring again today. Justifications for this paradox are largely focused on time delays, measurement issues, or that there may not actually be a positive impact of technology on productivity. My research offers support for the time delay justification and does not find any effect of artificial intelligence on TFP in the short run. Although my research does not suggest an exact period of time before productivity results may appear, it does suggest that it is quite a few years rather than a few months. This explanation is consistent with economic theory presented by Brynjolffson, Rock, and Syverson (2017).

I develop OLS models and run regressions with both one and two-way fixed effects. Because of the nature of my model, two-way fixed effects are necessary to account for potential endogeneity issues. Also, I utilize per capita estimates for my independent variables in order to account for the large difference in size and wealth of countries in my dataset. In addition to using per capita estimates, I use a logarithmic construction following the methodology set forth by Letta and Tol (2018).

The regression results from my per capita estimates for the number of startups coefficient was both positive and significant until the two-way fixed effects for country and time were included in my model. These results are consistent with the productivity paradox that has been found in previous research and were expected. However, the coefficient on the number of startups variable is significant in three out of the four specifications, including the model with two-way fixed effects for both year and region. Since the parameter estimates do not change beyond two standard deviations in the significant specifications, I can make the assumption that there is an effect on TFP due to the number of startups but significance in the third specification (two-way fixed effects for time and country) is lost because of the country dummy variables. This could occur because of the relatively small size of my dataset. There may not be enough data points to simultaneously estimate both country fixed effects and the effect of the number of startups and as a result enough multicollinearity is created that we do not see significant results in the third specification.

The regression results from the logarithmic estimation yielded slightly different results. While most of my independent variables appear to have no effect on productivity, the coefficient on the number of startups focused on artificial intelligence is both positive and significant. This result is surprising and is inconsistent with the productivity paradox found previously. Its positive significance suggests that increasing the number of startups focused on artificial intelligence in a given country will increase that country's overall productivity. More broadly, this offers some empirical evidence supporting the economic theory that improving technology should have a positive effect on productivity.

The positive and significant coefficient on my logarithmic estimation offers support for one explanation of the productivity paradox. The primary explanation for the productivity paradox is that the expected positive effects are not captured by the statistics because there has not been enough time since their implementation. This justification is consistent with economic theory and was an explanation for the Solow Productivity paradox of the 1980s and 1990s as well. The significance on my number of startups variable reinforces this explanation because of the rapid impact opening a startup has on the economy. Startups focused on artificial intelligence contribute to the development and implementation of new technology that immediately has an effect on the production of an economy. As a result, I expect the coefficient on the number of startups to be unaffected by the potential time lag after implementation. This hypothesis is supported by my results.

The lack of significance on my other independent variables also supports the above justification for the productivity paradox. My other variables measure aspects of artificial intelligence implementation that are performed today in order to create future benefits. For these variables, it is expected that a time lag may prevent their effects from appearing in the productivity statistics. The positive significance of the number of startups variable combined with the lack of significance of my other dependent variables ultimately support the time lag justification of the productivity paradox.

There are three clear potential implications of my findings. The first is regarding artificial intelligence investment. My research supports the time lag explanation of the productivity paradox and therefore countries should not be discouraged from investing in artificial intelligence even if the results do not appear immediately in the productivity statistics. My research does not suggest that artificial intelligence will pay dividends in future years, and due to my short time panel, I cannot address this issue, but it does offer support for the time lag justification of the productivity paradox.

The second potential implication of my research is largely focused on policy. Since the number of startups focused on artificial intelligence has a positive impact on productivity, it is advantageous for countries to create policies that encourage startup development. Policy adjustments based on this research would include easing restrictions on starting a business and possibly even supplementing the new companies with government aid. My research indicates that increasing the number of startups focused on artificial intelligence in a given country should increase productivity and ultimately benefit the country as a whole. This implication is consistent with my earlier conclusion that the use and adoption of technology by firms is what is most important.

The third potential implication of my research, which also centers around the time lag justification of the productivity paradox, is that poor countries may not be able to participate in the same level of artificial intelligence implementation as wealthy countries. Since it takes time for the investment in artificial intelligence to pay off, the reward from that investment may not come for many years into the future. While this may be a reasonable expectation of wealthy countries, poor countries may not have the time and resources to develop a technology that will not positively impact their economy for a number of years. It may be advantageous for those countries to invest in other forms of technology that will increase their productivity immediately instead of artificial intelligence which appears to have a delayed impact. It is important to note here that positive spillover effects could occur in poor countries. If the technologically advanced countries develop this technology, countries that did not develop it can benefit from it once it becomes publicly available.

Although I have taken all of the necessary steps to ensure the reliability of my results, it is important to note that there are still some concerns regarding my conclusions. Confounding variables are one potential cause of concern. Confounding variables occur when there is an unmeasured variable that effects both my independent and dependent variable. Since artificial intelligence is a relatively new field, it is possible that a key variable that effects both artificial intelligence implementation and productivity has been left out of the equation and skewing my results. However, I am not aware of a way to test this for my data. The second concern has to do with the construction of TFP. TFP contains a wide range of factors and it is calculated instead of observed. Economists expect changes in TFP to be a result of technological change, but that may not always be the case. Letta and Tol (2018) state, “changes in TFP can also be due to managerial or behavioral change, changes in the structure of the economy or company entry and exit within sectors, changes in regulation or taxation, changes in the provision of public goods, changes in market power, or changes in international trade.” Ultimately, the increase in TFP I see from my results may not be exclusively from artificial investment and may involve a number of other factors. My data does not allow me to identify exactly what is affecting TFP growth.

The primary concern with my results and interpretations comes from my small dataset. When analyzing trends in TFP, economists typically use large datasets in order to capture the trends accurately. Unfortunately, since artificial intelligence is a relatively new field, data is not available for many years. As a result of a small dataset, the number of independent variables I can use is restricted and the potential for measurement issues is increased. I chose the Stanford HAI data because I felt it was the best dataset to capture holistic artificial intelligence implementation rather than focusing on one particular aspect. If this data had been available for more than four years, I believe I would have been able to look at the long-term effects of artificial intelligence implementation on total factor productivity instead of only focusing on the short term.

In the future I would suggest returning to this construction with more years of data. Creating this unique dataset by combining Stanford HAI with the Penn World Tables provides insightful information about the impacts of artificial intelligence. Some economists have developed alternative variables to TFP to measure the impact technology has on an economy (Antonelli and Quartero, 2017, Bar Shira, et al, 2003, Carlaw and Lipsey, 2003). Using one of these alternative variables may assist in determining the specific effect of artificial intelligence implementation on productivity as a whole. More research will be necessary in the future to determine which justification for the productivity paradox is correct and when the impact of artificial intelligence can be seen in productivity statistics.

BIBLIOGRAPHY

Aghion, Phillippe, Benjamin F. Jones and Charles I. Jones. 2017. "Artificial Intelligence and Economic Growth," NBER Working Paper # 23928.

Aghion, Philippe, Celine Antonin and Simon Bunel. 2019. "Artificial Intelligence, Growth and Employment: The Role of Policy." *Economics and Statistics*: 149-164.

Agrawal, Ajay, Gans, Joshua and Avi Goldfarb. 2019. "Economic Policy for Artificial Intelligence." *National Bureau for Economic Research.*

Albarrán, Daniel Gallardo, Robert Inklaar and Pieter Woltjer. 2019. "The Composition of Capital and Cross-Country Productivity Comparisons." *International Productivity Monitor* (36): 34-52.

Antonelli, Cristiano and Francesco Quatraro. 2010. "The Effects of Biased Technological Change on Total Factor Productivity: Empirical Evidence from a Sample of OECD Countries." *Journal of Technology Transfer* (35): 361-383.

Atrostic, K. and Sang V. Nguyen. 2005. "IT and Productivity in U.S. Manufacturing: Do Computer Networks Matter." *Economic Inquiry,* (vol. 5, 3): 493-506.

Bailey, Micheal A. 2017. *Real Econometrics.* Oxford University Press.

Bar-Shira, Ziv, Israel Finkelshtain, and Avi Simhon. 2003. "Cross-Country Productivity Comparisons: The "Revealed Superiority" Approach." *Journal of Economic Growth* (8, no. 3): 301-323.

Bonsu, Kwadwo Osei and Jie Song. 2020. "Turbulence on the Global Economy Influenced by Artificial Intelligence and Foreign Policy Inefficiencies," *Journal of Liberty and International Affairs* (6, no. 2): 113-122.

Brynjolfsson, Erik, Daniel Rock and Chad Syverson. 2017. "Artificial Intelligence and the Modern Productivity Paradox," NBER Working Paper #24001.

Byrne, David M., John G. Fernald and Marshall B. Reinsdorf. 2016. "Does the United States have a productivity slowdown or a measurement problem?" *Brookings Papers on Economic Activity.*

Cameron, Euan, Jon Andrews and Jonathan Gillham. 2017. "The economic impact of artificial intelligence on the UK economy." *PricewaterhouseCoopers LLP*.

Carlaw, Kenneth I. and Richard G. Lipsey. 2003. "Productivity, Technology and Economic Growth: What is the Relationship?" *Journal of Economic Surveys*: 457-495.

Cetindamar, Dilek, Thorsten Lammers and Yi Zhang. 2020. "Exploring the knowledge spillovers of a technology in an entrepreneurial ecosystem – The case of artificial intelligence in Sydney." *Thunderbird International Business Review* (62): 457-474.

Crafts, Nicholas. 2010. "The Contribution of New Technology to Economic Growth: Lessons from Economic History." *Journal of Iberian and Latin American Economic History* (28, no. 3): 409-440.

Eberhardt, Markus and Francis Teal. 2010. "Productivity Analysis in Global Manufacturing Production," University of Oxford Discussion Paper Series #515.

Feenstra, Robert C., Robert Inklaar and Marcel P. Timmer. 2015. "The Next Generation of the Penn World Table." *American Economic Review* (105, no. 10): 3150-3182.

Fox, Kevin J. 2012. "Problems with (Dis)Aggregating Productivity, and Another Productivity Paradox." *Journal of Productivity Analysis* (37): 249-259.

Goldfarb, Avi and Catherine Tucker. 2011. "Online Display Advertising: Targeting and Obstrusiveness." *Marketing Science* (vol. 30, no. 3): 389-404.

Goldfarb, Avi and Daniel Trefler. 2018. "AI and International Trade," NBER Working Paper # 24524.

Horowitz, Michael C. 2018. “Artificial Intelligence, International Competition, and the Balance of Power.” *Texas National Security Review* (vol. 1, 3).

Hogendorn, Christiaan and Brett Frishmann. 2017. “Infrastructure and General Purpose Technologies: A Technology Flow Framework.” Wesleyan Economic Working Paper #2017-001.

Hulten, Charles R. 2000. “Total Factor Productivity: A Short Biography,” NBER Working Paper #7471.

Hulten, Charles R. 2001. “Total Factor Productivity: A Short Biography”. *The University of Chicago Press.*

Jones, Charles I. and Peter J. Klenow. 2016. “Beyond GDP? Welfare Across Countries and Time.” *American Economic Review* (106, no. 9): 2426-2457.

Lafuente, Esteban, Zoltan J. Acs, Mark Sanders and László Szerb. (2019). “The Global Technology Frontier: Productivity Growth and the Relevance of Kirznerian and Schumpterian Entreprenuership. *Small Business Economics*: 1-40.

Letta, Marco and Richard S.J. Tol. 2018. “Weather, Climate, and Total Factor Productivity.” *Environmental and Resource Economics* (73): 283-305.

McKinsey & Company. 2017. “A Future that Works: Automation, Employment, and Productivity. *McKinsey Global Institute.*

Mims, Christopher. 2018. “Inside the New Industrial Revolution.” *The Wall Street Journal.*

Mou, Xiaomin. 2019. “Artificial Intelligence: Investment Trends and Selected Industry Use.” *International Finance Corporation.*

Naqvi, Al. 2017. “Competitive Dynamics in Artificial Intelligence Economy: The Wicked Problem of Cognitive Competition,” *Journal of Economics Library* (4, no. 2): 187-193.

O’Neil, Cathy. 2016. *Weapons of Math Destruction.* Crown Publishing.

Panch, Trishan, Szolovits, Peter and Rifat Atun. 2018. “Artificial Intelligence, Machine Learning and Health Systems.” *Journal of Global Health* (vol. 8, 3).

Qiulin, Chen, Xu Duo and Zhou Yi. 2019. “AI’s Effects on Economic Growth in Aging Society: Induced Innovation and Labor Supplemental Substitution.” *China Economist* (14): 54-66.

Rossana, Robert J. 2011. *Macroeconomics,* Taylor & Francis Group. *Proquest Ebook Central.*

Santos, Roberto S. and Lingling Qin. 2019. “Risk Capital and Emerging Technologies: Innovation and Investment Patterns Based on Artificial Intelligence Patent Data Analysis.” *Journal of Risk and Finanical Management* (vol. 12, 189).

Seamans, Robert and Manav Raj. 2018. “AI, Labor, Productivity and the Need for Firm Level Data,” NBER Working Paper #24239

Syverson, Chad. 2017. “Challenges to Mismeasurement Explanations for the US Productivity Slowdown.” *Journal of Economic Perspectives* (vol. 31, 2): 165-186.

Triplett, Jack E. 1998. “The Solow Productivity Paradox: What Do Computers Do to Productivity?” *The Brookings Institute.*

van Ark, Bart. 2016. “The Productivity Paradox of the New Digital Economy.” *International Productivity Monitor* (31): 3-18.

Watanabe, Chihiro, Naveed Kashif, Yuji Tou and Pekka Neittaanmäki. 2018. “Measuring GDP in the Digital Economy: Increasing Dependence on Uncaptured GDP,” *Technological Forecasting and Social Change:* 2-15.