IUISS

Department of Economics and Finance, Chair of Macroeconomics

AI AND ROBOTS - THE IMPACT ON THE LABOUR MARKET

SUPERVISOR

Prof. Francesco Lippi

CANDIDATE

Sara Ruggieri

ID number: 207391

Academic year 2018/2019

Index

Introduction	3
Chapter I	5
1.1 A brief history of technological change – Industrial revolution 4.0	6
1.2 Technological Unemployment	8
1.3 Current situation and polarisation of markets	11
Chapter II	14
2.1 The Luddite Fallacy	15
2.2 The productivity effect	16
2.2.1 An increase in productivity – Cobb-Douglas and Okun's Law	17
2.2.2 Is new technology Pareto-improving?	19
2.3 The displacement effect	21
2.3.1 Advantages of short-run displacement	21
2.3.2 Impairing persistent displacement	23
Chapter III	25
3.1 The Caselli-Manning Model	25
3.1.1 The benchmark model and the assumptions	
3.1.2 Findings	27
3.1.3 Deviating from the assumptions	
3.2 Comparison with the Acemoglu-Restrepo Model	31
Conclusions	
References	

Introduction

The debate concerning whether the advancement of automation and technological improvement will cause irreversible technological unemployment has puzzled economists for the last few centuries, and the advent of robots and Artificial Intelligence has made it livelier than ever. Past events cannot give us an idea of what will come next, since, according to some scholars, "this time is different" from every other time in the past when automation made workers' lives easier, and that by increasing the amount of tasks that can be carried out by machines there will be permanent job displacement, affecting both blue-collar and white-collar workers alike.

On the other hand, some other economists support a more positive standpoint on the future of automation: it is true that some jobs will be destroyed, but others will be created offsetting the negative effects on employment, and the long-term effects will result in a shift across industries. Also, according to Reinhart and Rogoff, "More money has been lost because of four words than at the point of a gun. Those words are 'This time is different."¹. The underlying idea of this sentence, which is about financial crises but which can be applied to any market disruption, is that even though every new disruption seems different to the one before, it is not. Following this line, technology will not have impact on the job market different from those it always has had: no peak in long-term unemployment, no catastrophic impact on welfare.

Despite this, the lively debate on whether AI will cause irreversible technological unemployment has not yet reached a conclusion; new features of technology, especially machine learning, are of concern because of the much wider range of tasks performable by machines. But if new technology affected the market as currently existing automation has done, it would have a positive impact due to the productivity effect and a negative impact due to the displacement of workers. Through the means of academic papers and existing literature on the topic, I will try to assess whether the positive effects will offset the negative ones or if on the contrary the level of technological unemployment that the advancement of technology will generate will be too high to be restrained.

Chapter I describes the history of automation and how it affected the job market over the centuries, starting from the invention of the steam engine during the First Industrial Revolution up to the current Industry 4.0. Despite the increase in productivity experienced after every industrial revolution, the growing concern of technological unemployment has raised questions in many economists, who tried to assess whether unemployment generated by technological advancement is a long-term problem or a transitory effect. Nevertheless, it is the general consensus that some tasks are more at risk of automation than others, especially middle-skilled routine tasks associated with middle-wage occupations. This caused workers to shift towards

¹ Reinhart, C. M., & Rogoff, K. S. (2009). *This time is different: Eight centuries of financial folly*. Princeton university press.

jobs at the extremities of the wage spectrum, and consequently the wages of workers as well as the job market as a whole to polarise, increasing income inequality.

Chapter II presents one the most iconic examples in history of the contrast between the two main forces that develop as a consequence of technological advancement, namely productivity effect and displacement effect: the Luddite riots, which saw periods of social unrest due to the displacement of textile workers as a result of the invention of mechanised looms. This social movement gave origin to the Luddite fallacy, the belief that as automation increases, workers will lose their jobs offsetting the benefits of the increase in productivity and causing a sharp drop in total welfare. Whether new technology entails a positive or negative impact on the economy depends on the combination of productivity and displacement effect. The productivity effect has an undoubtedly positive impact on the economy. Because of the better performance and the lower cost of technology compared to workers, it is possible to have lower price of goods and more goods produced for a given amount of factors of production. This productivity increase is Pareto-improving in case of perfect mobility of workers and no redistribution costs. The displacement effect is more controversial, given that, as mentioned above, it is not clear whether it is a long-term or a short-term effect. If it is only present in the short run, it means that the displaced workers are eventually reabsorbed by other industries or employed in the new tasks that new technology has generated, and may lead to higher average wages and education and better working conditions; if it is persistent in the long run however it may constitute a threat for both white-collar and blue-collar workers as the impossible reallocation of the too many displaced workers would topple the whole economy.

Chapter III describes the Caselli-Manning model presented in "*Robot Arithmetic: New Technology And Wages*" (2018), which focuses on the gainers rather than on the losers from the employment of new technology. The assumptions of the model are the following: fixed labour, constant returns to scale, constant interest rate (as technology changes), perfect competition, different levels of technology across economies and homothetic consumer preferences. It investigates the impact of technological advancement on wages and it finds very optimistic results: after assessing that at least one type of workers is made better off by new technology, it reaches the conclusion that if the price of investment goods does not rise in relative to the price of consumption goods, average real wages rise, and that if labour is perfectly elastic in supply all workers gain. Nonetheless, this model does not take into account the displacement effect, so it is compared to the Acemoglu-Restrepo model presented in "Modeling automation" (2018), in order to have a more comprehensive framework. The comparison reveals that these two models clash under several standpoints, showing that a unanimous agreement on the effects caused by new technology on the job market has not yet been reached.

Chapter I

"Technological possibilities are an uncharted sea."

Joseph Schumpeter

Undeniably, technology has the potential to significantly boost up productivity in many (if not all) sectors, but the impact that technological change has on employment is yet to be defined. The major macroeconomic outcomes, generated by the effects of new technology that historically have shaped the labour market, are twofold: technological unemployment and job market polarisation. While it is certain that the latter is caused by new technology, which is able to substitute human workers in performing tasks associated with the middle of the wage distribution, no consensus has been reached on the former: the question of whether technology causes unemployment only temporarily or whether its impacts last in the long run has been asked by some of the greatest minds of the last two centuries, who in some cases reached opposing results. In this chapter I will describe these aspects by starting with a brief history of automation (1.1), whose beginning can be identified with the invention of the steam engine during the First Industrial Revolution. The technological progress made during the first, second and third revolutions, laid the foundations to reach the current brink of Industry 4.0, which revolves around the internet and Big Data; during these ages productivity has increased and living conditions have improved, even though it raised the question of whether the technological unemployment it caused would only be a temporary effect or whether it would have permanent implications. Section 1.2 describes the views of some of the most important economists of the last centuries: some considered technological unemployment just a "temporary phase of maladjustment" (Keynes, Say) while some others were fairly concerned about its long-run implications, fearing that the number of displaced workers could not be fully reabsorbed by other industries (Marx, Ricardo, Schumpeter). Despite this, more recent analyses have found that not only technological unemployment is not a long-term issue (Feldmann, Nakamura and Zeira), but also that automation enhances the tasks complementary to it. If technological change happens too rapidly though, aggregate employment might fall. This has not happened so far, but a higher level of technology can lead to an increase in income inequality and polarisation. Section 1.3 describes the ways in which technological improvements lead to polarisation of the job market through a task-based approach. The occupations are divided by skills in three groups, high, medium and low skills, where the first employs non-routine skills, the second employs routine skills, and the last uses both routine and non-routine skills depending on the occupation: it is shown that workers performing non-routine tasks (divided in two sub-groups, abstract and non-routine manual) benefit from technological advancements, while workers performing routine tasks are at risk of displacement. This causes workers to shift towards the two ends of the skill spectrum, leading to a polarisation of the job market and consequently to a polarisation of wages, increasing inequality. But while the effect on the high-skill end of the spectrum is clearly positive and the effect in the middle clearly negative, the effect on the low-skill end is ambiguous, since it involves both routine and non-routine tasks.

1.1 A brief history of technological change – Industrial revolution 4.0

Even though the term "automation" was coined in 1948 by Delmer S. Harder, Vice President of Ford, the first mention of it can be found in Homer's Iliad. Hephaestus, the Greek god of fire and blacksmiths, used to create self-operating robot-like creatures out of metal, that would help him in his work of crafting equipment for both gods and mortals. Although this ancient tale is quite quaint (and hardly credible), it is a good example to show for how long men and machines have worked side by side: the first machines though, because of their simplicity, brought nothing but benefits to the society.

The first disruptive machine that unfolded an unprecedented technological development was the steam engine in 1765: this marked the beginning of the First Industrial Revolution. The steam engine, brought to the world by James Watt, led to an increase in productivity because steam overcame the limits of both pre-existing machines (based on coal) and human work. In fact, it allowed mass production by upgrading factories and mass transportation by upgrading railways. Factories invariably had to be located next to a river to generate (little) power through the water wheel: with the introduction of steam engine, much more power was producible, and factories could be placed next to pre-existing railroads so that the goods manufactured could be transported more easily.

The graph below represents the situation in England, as it is the cradle of these two revolutions: after a few decades of adjustment during which wages were stagnant, the first industrial revolution caused average real wages for all the groups of workers analysed to rise. With the enhancement of the production process due to mechanisation, more factories were built, and new employment opportunities arose. The population migrated from the countryside to urbanised areas for a higher pay in industries rather than a lower one in agricultural occupations, causing average income to rise and demand for consumer goods to increase.



Graph 1

Adult male average full-time earnings for selected groups of workers (constant prices)²

² Lindert, P. H., & Williamson, J. G. (1983). English workers' living standards during the industrial revolution: a new look. *The Economic History Review*, *36*(1), 1-25.

The Second Industrial Revolution, also known as Technological Revolution, started in the second half of the 19th century: the pivots of this new age of development were electricity, telephone, internal combustion engine and indoor running water. This second revolution, in addition to further boosting productivity, improved the living conditions of people: sewer systems improved public health; crop failures did not imply famine since consumer goods could easily and quickly be transported; electricity led to automation. In fact, it is during the later stages of this era that Henry Ford introduced into his factory the moving assembly line.

The Second Industrial Revolution yielded the same effects of the first one: overall, it improved living conditions and led to higher real wages (J. Mokyr, 1998). Moreover, the labour demand increased because of the new jobs created by mechanisation: *"for the economy as a whole to switch from manual techniques to a mechanized production required hundreds of inventors, thousands of innovating entrepreneurs, and tens of thousands of mechanics, technicians, and dexterous rank and file workers."* (Mokyr, 1994). Indeed, both the first two waves of industrialisation led to a change in the nature of employment, since new occupations were no longer skilled but rather routine, and to a polarisation of markets. It also shifted the pivot of technological development from England to the United States.

The Third Industrial Revolution had a completely different approach to technological improvement: in fact, this modern revolution revolved around Information and Communication Technology (ICT). The Third Industrial Revolution, which started in the second half of the 20th century, shifted the focus of development to digital automation of production thanks to electronics: the rise of telecommunications, personal computers, internet and wireless technologies. This era led to the introduction of automatons (programmable logic controllers) and robots. With this increasing industrial mechanisation, the organisation of production changed: the assembly line transformed into separate bundled processes, which no longer required routine tasks. The type new labour demanded was once again skilled.

We are now on the brink of Industry 4.0, the Fourth Industrial Revolution. The new production process revolves around the internet and Big Data to coordinate and integrate the several steps of the supply chain. This interconnections among both vertical and horizontal stages of production is deconstructing the production chain, increasingly promoting the reinforcement of global value chains. Since the whole process will be optimised, productivity will increase also thanks to the use of algorithms such as predictive maintenance and light out production (machines keep working 24/7 even when no workers are around).

Intelligent machines also affect our daily lives. Thanks to machine learning, AI technologies, such as predictive algorithms, self-driving cars, translating software and even the smartphone keyboard corrector, have made some tasks easier than before. Watching a movie, making a payment or booking a flight have never been done as effortlessly as nowadays, and many improvements are yet to come. These technologies use the huge amount of data and computing power that has exponentially increased in the last few decades.

Since machines are more productive and cost efficient than men, the fear that machines will ultimately replace workers is widespread. Klaus Schwab compared the three biggest companies of Detroit in 1990 and the biggest three companies in the Silicon Valley in 2014. What he found was that the three companies in the Silicon Valley generated the same revenues as the ones in Detroit but using ten times less workers³.

The threat of job displacement is therefore not groundless, but whilst some affirm that increasing automation will start a "creative destruction" where new jobs will replace the displaced ones, some others say that it will lead to unabsorbable technological unemployment.

1.2 Technological Unemployment

"We are being afflicted with a new disease of which some readers may not have heard the name, but of which they will hear a great deal in the years to come—namely, technological unemployment."

John Maynard Keynes, 1930

The term "technological unemployment" indicates the share of structural unemployment "*due to our discovery of means of economising the use of labour outrunning the pace at which we can find new uses for labour*" and was made popular by Keynes, who defined it as "*only a temporary phase of maladjustment*"⁴. Despite Keynes' "Economic possibilities for our grandchildren" was published in 1930, the concept of technological unemployment is incredibly older. In 350 B.C., in fact, Aristotle wrote about Hephaestus' automatons saying that in the future machines could make human work useless.

Technological unemployment did not worry Keynes, who believed that it would not be a problem in the long run and that by 2030 (a hundred years from 1930) the economy would grow by 4 to 8 times; a first analysis made by David Ricardo shares the same long-run result. In the first editions of "Principles of political economy and taxation" Ricardo affirms that, since technological progress reduces the labour requirement to produce one unit of a good, when new machinery is introduced, if the demand for that good does not increase proportionally to the increase in productivity, then workers are displaced. However, the displaced workers can be employed in the production of another good for which there is demand. This is the accomplishment of Say's Law, according to which supply determines its own demand. Therefore, Keynes agreed with Ricardo when saying that persistent technological unemployment cannot exist, since all unused resources are employed in the production of another commodity. Say as well reached the same conclusion, stating that technological improvements in the production process do displace workers in the industries using that new technology, but

³ The three biggest companies of Detroit had revenues of \$250 billion with 1.2 million employees (1990), while the three biggest companies of the Silicon Valley had revenues of \$247 billion with 137,000 employees.

⁴ Keynes, J. M. (2010). Economic possibilities for our grandchildren. In *Essays in persuasion* (pp. 321-332). Palgrave Macmillan, London.

create new jobs in the industries that produce the new machines. Moreover, due to the new cost-saving technology, prices are likely to fall, demand for goods to increase and therefore demand for labour to rise.

However, in the third and last version of his book, Ricardo rectifies himself, admitting that he had made a mistake, since he no longer believed that Say's Law could avoid permanent displacement of workers. In fact, he corrects his first analysis by saying that initially he had assumed that in a society, if the net income – defined as the sum of rents and profits, increases, then also the gross income – defined as the sum of rents, profits and wages, increases. This, however, is not to be taken as granted: in fact, capitalists are interested in increasing the net income, i.e. profits, rather than gross income, i.e. employment. Income may rise, not because of higher wages, but because labour was substituted by machines.

Karl Marx reached the same conclusion, describing the will of capitalists to shift to more capital-using technology in the short run to have a higher profit at the expenses of the proletariat, who would end up being impoverished since it would face technological unemployment in the long run due to the substitution of workers for machines. He also ruled out that the number of new jobs created could offset the number of jobs lost to machines: "*Machinery does not just act as a superior competitor to the worker, always on the point of making him superfluous. It is the most powerful weapon for suppressing strikes.*"⁵.

Even Schumpeter, who thought of innovation as the main driver of economic growth, described new technologies as a "*perennial gale of creative destruction*", realising that a rapid change in the technology employed could cause severe disruptions in the industries using it, which unavoidably implied unemployment.

Clearly, technological unemployment has always been a questionable topic, and the magnitude of the impact of new technology on employment is ambiguous. A study conducted by H. Feldmann in 2013, using a sample of 21 industrial countries, found that in the period from 1985 to 2009, when the effects of the Third Industrial Revolution were unfolded, there was a correlation between unemployment and technological change, expressed as patents available. In fact, an increase of a standard deviation in technological change caused an increase in unemployment of between 2.3 and 3.0 percentage points. Yet, this was only a temporary effect: the effect grew smaller after a year and completely faded after two years, leaving no impact in the long run. The same conclusion was drawn by Nakamura and Zeira (2018): their model empirically proved that technological unemployment diminishes over time to eventually converge to zero in the long run, even though the authors themselves treat this result with caution because of the assumption that technology being currently developed will affect the markets as in the past.

All types of automation so far, whether it be tractors in agriculture, mass-production assembly lines in industries or tabulation and calculation in the tertiary sector, have aimed at overcoming human limits by increasing productivity, reducing time per unit of output and minimising risk of errors and, overall, they lay

⁵ Fisher and Taub (2017)

foundations for the development of even more labour-saving technologies. Yet, automation has not decreased aggregate employment. The explanation is simple: while it is true that some routine tasks can be easily carried out by machines, those that cannot are supported and augmented by automation.

An example of this phenomenon can be found in the banking sector, as observed by James E. Bessen. In his analysis, he considers the relationship between ATMs and bank clerks: one would think that, with the introduction of Automated Teller Machines, there would be less demand for human tellers to hand cash to customers. What happened was quite the opposite. Between their introduction and 2010, the number of ATMs in the US more than quadrupled, and the number of bank clerks also slightly increased, going from 500,000 in the 1980s to 550,000 in 2010. Bessen observed two forces at work. First, an indirect increase in demand for clerks due to the increase in the number of ATMs; since ATMs were cost-saving for each branch, fewer human workers were needed per branch, but the reduction in costs allowed the number of urban branches to increase by 40%. Second, since the routine tasks of bank clerks were carried out by machines, bank personnel became more aware of the importance of building a relationship with each customer, which not only ensured loyalty to the bank, but also gave the opportunity to introduce them to new services and products offered.

Nevertheless, there is no consensus on whether technological unemployment is just a temporary deviation or entails a permanent effect. When the introduction of new technology in an industry erases the need for specific types of skills, those workers have to develop the new skills required by the new technology to find an occupation. This process, though, lasts for an unspecified period of time. If the effects are only in the short run, unemployed workers will develop the new skills and the economy will reach a new equilibrium where all unemployment is absorbed by the new businesses and occupations created by the innovations; but if the process of job finding lasts a long time and technology evolves fast, by the time unemployed workers have acquired the new skills, technology might have changed, making those skills once again obsolete.



Graph 2 - Employment to population ratio

However, not all skills have become obsolete. As can be seen from the graph, employment to population ratio had an increasing trend (with some fluctuations) for the last 70 years at least, with no apparent long-run rise in unemployment. Nonetheless, "technophobes" point out that the new technology currently being developed as an upgrade of automation, namely Artificial Intelligence and robotics, will disrupt previous patterns and trends.

In fact, while the "*creative destruction*" described by Schumpeter takes place, many, if not all, routine tasks will be undertaken by machines, and not all involve a human relationship side as in the example of the bank tellers. This distinction may exacerbate income inequality in addition to creating technological unemployment, creating financial instability in particular for low-skilled workers.

While we cannot predict the effect on aggregate employment of technologies such as Artificial Intelligence and robots, we can forecast that the increase in income inequality will cause a further polarisation within the job market, causing major sectoral shifts.

1.3 Current situation and polarisation of markets

As new technology and information systems improve and fall in price, the demand for skills grows biased towards certain types: because of this, technological change is said to be skill-biased, meaning that new technology increases demand for specific types of skills and lowers it for others, causing the employments of different types of occupations to grow differently and therefore changing occupational and wage distribution. For this reason, it is believed that technology, in addition to offshoring, is a major factor that caused the job market polarisation which the modern economy is experiencing. To describe the phenomenon of job market polarisation, it is convenient to have a task-based approach to the matter.

Firstly, it is important to make a distinction between skills and tasks. A task is a work activity that yields a product; a skill is an allocation of capabilities to perform various tasks. This difference is important since workers of a given level of skills may be able to change the set of tasks to perform as the level of technology changes.

As technology advances and its price falls, it creates substantial advantages for more skilled workers, whose productivity increases. A simple example can be found in computer programming. Even though the capabilities of computers nowadays are undoubtedly astonishing, they are not unlimited and, more precisely, not natural. In fact, computers, like any machine, are not equipped with flexibility or judgement, but they simply follow instructions. Therefore, their capability of performing a task depends on how clearly and comprehensively a (human) computer programmer has written the instructions and the rules the machine must follow. Because machines only follow instructions, they can carry out any task that can easily be expressed and formulated as such: these tasks are defined as routine.

Routine tasks are usually linked to middle-skill occupations such as bookkeeping, clerical work and administrative occupations in general. The graph below shows a spectrum of occupations ordered by skills (high to low). It is clear that the occupations in the middle of the spectrum experienced a fall in demand in the last period analysed, confirming that routine middle-skill tasks are the ones which suffer the most due to technological advancements, being more easily automatable.





The fall in employment for administrative, clerical and production occupations in 2007-2009 is a good example showing the likely effect of a decrease in price of machinery able to carry out such tasks. Moreover, this fall in price caused a great ease of offshoring information-related work to foreign and cheaper worksites. These two factors combined cause the demand for skilled workers that can perform complementary non-routine tasks to increase.

Non-routine tasks, as described by Acemoglu and Autor (2011), can be divided in two major groups: abstract and manual. Abstract tasks include those tasks which need capabilities such as problem solving and other soft skills, intuition and creativity, and involve occupations such as doctors, scientists, engineers, designers and managers, which are complementary to routine tasks-performing new technology (and benefit from its fall in price) since they greatly rely on information: all the occupations mentioned imply a high level of education. Non-routine manual tasks, on the other hand, do not require a high level of education but rather low-medium, and are those which require flexibility and adaptability, social interaction, physical suitability,

⁶ "May/ORG CPS files for earnings yeat 1979-2009. The data include all persons ages 16-64 who reported having worked in 2009, excluding those employed by the military and in agricultural occupations. Occupations are first converted from their respective scheme into 326 occupation groups consistent over the given period of time. All non-military, non-agriculture occupations are assigned to one of ten broad occupations presented in the figure." Acemoglu and Autor (2011)

experience and fluent spoken communication, in occupations regarding security and personal care services, maintenance, food preparation. These tasks cannot be replicated by machines because they are based on reactions to unforeseeable external inputs which, at the moment, cannot be translated into instructions.

Both these categories of occupations are hard to offshore, since they must be performed in person. Moreover, occupations which are intensive in abstract tasks and non-routine manual tasks respectively are quite the opposite as far as skills are concerned. As can be seen in Graph 3, the two groups of jobs are at the opposite extremities of the skill range. For the reasons described, middle-skilled occupations are the ones that are the most at risk of suffering a fall in demand because they are the easiest to automate, and a shrinkage of the middle-skilled group of jobs would cause a polarisation of the job market.

A clarification must be made concerning low-skill jobs. This category encompasses both non-routine and routine manual tasks, but the respective effects of new technology on these two sub-groups are the opposite of one another. As already mentioned, non-routine manual tasks are hard to automate because of the need for adaptability that they imply; differently, routine manual tasks are perfectly convertible into commanding instructions because of their repetitiveness and are at risk of automation together with middle-skilled jobs. Because of the two opposing forces at work inside the category of low-skill occupations, the effect on this side of the wage spectrum is ambiguous, but still the level of low-skilled workers demanded does not fall as much as middle-skilled ones.

Proof that the polarisation of the job market is actually caused (at least to some extent) by new technologies was given by Michaels, Natraj and Van Reenen (2014): their analysis considered three education groups (low, middle and high) and related them to investments in technology, particularly Information and Communication Technology, in 11 countries with data covering 25 years. The result was that the industries where new technology grew more rapidly were also the ones where the demand for highly educated workers grew more rapidly and the demand for medium-educated workers fell more rapidly.

If the job market polarises, wages may too, given that middle-skilled workers are also in the middle of the wage distribution: as assessed by Autor and Dorn (2013), job polarisation is a necessary condition for this to happen, but not sufficient. The additional situation that must occur is that wages of workers who work in non-routine manual tasks-intensive occupations grow at least as rapidly as abstract tasks-intensive occupations. If this is the case, job polarisation will not only have caused many middle-skilled workers to lose their job, but also it will have increased income inequality.

Chapter II

In the previous chapter two major underlying effects were described, pulling in opposite directions, which are generated by an increase in the level of technology: a positive one due to the improvement in the production process and a negative one due to the decrease in the level of employment, whether temporary or permanent. As described by Acemoglu and Restrepo, these are the productivity effect and the displacement effect. The overall impact of new technology on employment and wages is given by the sum of the two forces, and therefore it is positive if the productivity effect offsets the displacement effect.

Regardless of the prevalence of the productivity effect or displacement effect, improvements in labour-saving technology have often found a strong resistance on the workers' side: the most iconic example can be found in 19th century England with the riots of the Luddites, which is described in section 2.1 together with the explanation of why the underlying reasoning, that total welfare decreases as a consequence of substitution of workers for machines, is considered a fallacy. In fact, data shows that in the 10 most automated countries since 1991, whilst unemployment did not increase overall, GDP grew steadily, showing that technology did not harm aggregate welfare.

Section 2.2 covers the increase in productivity due to new technology: as technology becomes more available and widespread, its price falls. As a consequence, new technology enables output to be produced at a lower price than labour: the cost of production decreases, yielding higher productivity - because it overcomes human limits, and an increase in wages for workers performing those tasks which are complementary to the new technology. To explain this, section 2.2.1 exploits the Cobb-Douglas production function and finds that an increase in the level of technology leads to a positive change in wages for at least one group of workers. Moreover, because the overall welfare increases, due to Okun's Law, the level of unemployment does not increase, since new technology does not replace all workers but rather it enhances the efficiency of workers performing complementary tasks. Section 2.2.2 studies under what conditions the increase in productivity caused by new technology constitutes a Pareto improvement: it considers the utilities of two groups, employers - who gain directly by acquiring new technology in order to increase productivity, and workers, who may face job displacement or a decrease in wages. If in the economy there is perfect mobility and no redistribution costs, with adequate redistribution, technological improvement is desirable for both groups as it is Pareto-improving; if on the other hand there is imperfect mobility or redistribution is costly, the utility curve faces constriction at the expense of workers, making them worse off and causing an increase in income inequality, which worsens welfare.

The last section (2.3) describes the displacement effect, which is caused by the fact that technology is cheaper than workers and reconnects to the aforementioned question of whether technological unemployment is only a temporary matter. If the displaced workers are reallocated to other non-automated tasks after the short term (2.3.1), displacement might have a positive outcome due to the fact that, because of more automation, humans

would not be needed anymore to perform the more fatiguing and risky jobs and the average working conditions and efficiency would improve. Moreover, since the occupations that are less at risk of automation are high-skill, it is likely that the average level of education across workers would rise. On the other hand, if displaced workers are too many to be absorbed by new or existing non-automated sectors, technological improvement has unemployment long-run implications: in fact, if productivity rises simultaneously to a fall in employment, a jobless growth occurs, as happened in India in the second half of the 20th century. In addition, this section describes the feature of Machine Learning, and the probability that, because of it, new technology might affect white-collar jobs the same way it affects blue-collar ones.

2.1 The Luddite Fallacy

The Luddites were a group of English textile workers who, during the Industrial Revolution, violently protested the introduction of new mechanised textile looms by destroying them, because these new highly productive tools were taking away their jobs. The unrest was so violent and long-lasting that "The Destruction of Stocking Frames" entailed a death sentence. Fear grew during those years, as people were speculating about automation growing in the future, causing workers to lose their jobs and reducing income and total welfare: if this were true people would stop consuming and those very companies that were cutting costs by using automation would be at risk. If the rate of automation were higher than the rate of job-finding, the production-consumption cycle would collapse.

This argument, however, is considered a fallacy: *"If the Luddite Fallacy were true, we would all be out of work because productivity has been increasing for two centuries"* (A. Tabarrok, 2003). So far, mechanisation and automation have spread and developed, without significant increase in unemployment relative to total welfare. This can be proven in the figures below, which represent unemployment and GDP historical data of the most automated countries in the world⁷.



Graph 4, Unemployment, total (% of total labor force) (modelled ILO estimate)

⁷ International Federation of Robotics data show that the 10 most automated countries in the world are: South Korea, China, Germany, Japan, Sweden, Denmark, USA, Italy, Belgium and Taiwan. No data from World Bank was available for Taiwan, so it was omitted from the analysis.



Graph 5, GDP per capita, PPP (constant 2011 international \$)

Comparing Graph 4 and Graph 5^8 , it is clear that, while GDP per capita steadily grew during the past years, unemployment has not had significant changes in the countries observed. If automation had a negative impact on the economy, we would have seen by now a rise in unemployment and/or a fall in GDP, which would have been of great magnitude considering the high use of automation that these countries have. Clearly, this is not the case: unemployment does not follow a trend in any of the countries analysed and GDP is steadily increasing. Therefore, assuming theoretically that automation is the only variable factor affecting these parameters, under the ceteris paribus condition, it is implied by these graphs that it has no impact on unemployment but that it steadily raises economic welfare and, treating employment as constant, productivity as well.

Despite these data, many scholars believe that this time is different and that past events do not enable us to forecast what new technology will lead to in the future. The introduction of Artificial Intelligence and Machine Learning is daunting for the fact that, for the first time, machines may be able to perfectly substitute workers in every sector and for both blue-collar and white-collar jobs, up to the point that they will be able to learn, evolve, and even think without the need for human intervention.

But this can still be considered a science fiction scenario, since so far machines are limited to adhering to the instructions they were programmed to follow. Machines can only perfectly substitute workers in some tasks and, when they can, they lead to a significant increase in overall productivity, even if the negative side effect is temporary displacement of some workers.

2.2 The productivity effect

Robots are increasingly being considered more efficient than human labour and are consequently being employed on an incremental basis. Thanks to the growth in Big Data⁹, artificial intelligence is extremely helpful to complete complex tasks that would require much more human labour, since robots are able to detect

⁸ Data in Graph 1 and Graph 2 retrieved on World Bank.

⁹ Big Data has seen a CAGR of over 50% since 2010, found an Accenture study.

trends in large pools of data. After all, "*Data is to AI what food is to humans*" (B. Smyth, University College Dublin). Moreover, artificial intelligence eliminates some ethical biases present in humans that hinder their decision-making capabilities, which is an advantage in fields such as courtrooms and healthcare. For these reasons, advanced machines are being used for an increasingly wider range of tasks, causing a decline in the costs: in fact, robot prices have fallen by roughly 10% in the last few decades, and the sales have reached a peak of a 40% increase in year 2011, with 166,000 units sold (Frey & Osborne, 2013). Because of its fall in price, new technology allows the production of output at a lower price than labour: the cost of production decreases, yielding a higher productivity. It has in fact been observed that, between 1993 and 2007, industrial robotics alone has caused a rise in annual labour productivity growth by 0.36 pp (Graetz & Michaels, 2015). Because of the lower cost, prices fall as well, increasing real income and demand for goods and services. Moreover, demand for complementary non-automated tasks increases.

It is safe to assume that this will lead to various positive consequences. One of the most glaring examples can be found in the annual median income of workers in the Silicon Valley, the cradle of the Tech Boom, which reached an impressive level of \$94,000, almost twice as much the national annual median income (\$53,000).

2.2.1 An increase in productivity - Cobb-Douglas and Okun's Law

The increase in wages in the sector that develops technologies as such, as technology advances, is a consequence of the increase in productivity. Such phenomena can be explained mathematically by the Cobb-Douglas production function.

$$F(K,L) = A \times K^{\alpha} \times L^{1-\alpha}$$

"A" is the total factor productivity (TFP), which is regarded as the main driver of GDP growth rate. Its most important sub-sections are considered to be technology growth and efficiency: therefore, a better-performing technology will increase the productivity of the economy and consequently the rate of growth of Gross Domestic Product. The production function can therefore be rewritten as follows:

$$F(L, K, \theta) = K^{\alpha} \times (\theta \times L)^{1-\alpha}$$

Where θ is the level of labour-augmenting technology and α is a number between 0 and 1. In perfectly competitive markets, each factor receives its marginal product: this implies that the wage received by labour will be

$$w = \frac{\partial F}{\partial L} = (1 - \alpha) \left(\frac{K}{L}\right)^{\alpha} \theta^{1 - \alpha} > 0$$

This shows that, because all terms are positive, labour increases productivity, and an increase in technology increases wages. On the other hand, one might be concerned that technology could decrease the marginal product of labour: this is not true because, once again, all terms are positive:

$$\frac{\partial^2 F}{\partial L \partial \theta} = (1 - \alpha)^2 \left(\frac{K}{L}\right)^{\alpha} \theta^{-\alpha} > 0$$

The marginal product of labour increases as the level of technology increases, so new technology causes wages to rise for at least a group of workers¹⁰, the ones whose skills are complementary to the new technology.

A study made by PwC reports that the contribution of AI to the labour force is expected to have a positive impact of \$15.7 trillion on the global GDP by 2030 and, during the same period, an average increase of 26% in local economies' GDP.

On the other hand, AI cannot be treated as a mere factor of production: it is in fact a non-traditional hybrid able to behave both as labour, by efficiently replacing humans in some tasks, and as capital, with the difference that it has the ability of learning from experience (differently from a plant or a piece of equipment). This fact makes so that AI will not completely replace human labour, causing a higher unemployment rate, or capital, but instead it will increase both factors' efficiency.

To prove this, we can consider Okun's Law.

$$\frac{y-y^*}{y^*} = -\beta(u-u^*)$$

As shown in the formula, there is a positive relationship between changes in employment rate and growth rate of real Gross Domestic Product (negative relationship between real GDP and unemployment). Hence, the increase in the technology level, i.e. a greater use of artificial intelligence in the workplace, which would lead to an increase in economic welfare on both a global and local level due to higher productivity, would not lead to a decrease of employed people, according to this analysis.

Okun's Law has often been considered imprecise due to differences between observed phenomena and forecasts: the discrepancies between Okun's Law predictions and data registered during recessions were explained by some studies examining the relationship between employment and economic growth. The findings show that economic growth does have a positive impact on employment growth, but the effects are not immediate: on the contrary, the effects can be observed with a lag of a few quarters (Seyfried, 2011).

An explanation to the fact that the introduction of artificial intelligence into the workplace will not displace workers in the long-run can be found in the degree of efficiency it brings to the entity employing it: AI was in fact found to be used mainly for back-office functions that can be performed by an algorithm, for example

¹⁰ The opposite situation will be considered in the third chapter by the Caselli-Manning model.

gathering data or responding to emails (Davenport & Ronanki, 2018), which in any case boost up an employee's efficiency by cutting off the time spent on bureaucratic or routine tasks and therefore allowing a more efficient and more productive allocation of workers' time.

An example of this improvement in the quality of labour thanks to artificial intelligence can be found in the pharmaceutical industry. Traditionally, new drugs where discovered and developed thanks to educated guesses on which molecules could bind together to yield the desired therapeutic purpose, which obviously was a costly and timewasting process; this changed when a company named Atomwise applied AI to the research procedure: the software had the capabilities of identifying poor-quality candidate molecules by predicting the outcome of different experiments without running them. This new feature increased efficiency, significantly decreased costs and increased returns for the downstream lab tasks, without harming employment: in fact, the downstream tasks are run by humans, and the higher returns are likely to lead to an increase in the demand for labour for those occupations to meet the increase in productivity. Moreover, in addition to increasing the demand for labour in pre-existing tasks, new technology also has the potential to generate innovations that will create new ones.

2.2.2 Is new technology Pareto-improving?

An increase in productivity due to new technology will allow the economy to produce a higher level of output with the same combination of factors of production. Moreover, the substitution of humans with machines for certain tasks allows workers to specialise in non-automated tasks, which yields a more efficient division of labour. For these reasons, the production possibilities frontier shifts outwards, which theoretically has the effect of making everyone better off. If the PPF shifts outwards, consequently the utility possibilities curve does as well. Consider two categories, employers (who are the owners of capital and introduce new technology in the production process) and workers.



Graph 6

 E_0 represents the initial equilibrium and E_1 is the equilibrium after technological improvement. Clearly, while employers experience a rise in utility, workers are worse off because of an increase in the level of labour-saving technology. In an economy with no redistribution costs, perfect mobility or adequate redistribution, new technology can be desirable for both groups. The area in bold of the utility possibilities frontier represents the different redistributions that lead to a Pareto improvement, so that both workers and employers benefit from the surplus generated by new technology or so that neither group is worse off.

But this is not to be taken for granted since, in reality, technological improvements might change the income distribution and exacerbate income inequality causing a fall in total income. In this case, the losers are considerably worse off. In the real world, in fact, perfect mobility and absence of redistribution costs are hard to find. If there are redistribution costs, the utility possibilities frontier after the technological change is constrained, at the expense of workers.



In this case, part of the curve after the technological change lies below the initial curve, so technological change clearly does not lead to a Pareto improvement. This means that, even with redistribution, new technology will unavoidably cause losses for the group of workers. Therefore, it is comprehensible that workers fight against technological improvements, like in the case of the Luddites.

Moreover, this situation might have negative welfare implication. If the total welfare function does not consider inequality, meaning that an additional unit of income to worse-off workers has the same weight of an additional unit of income to better-off employers, technological change would still be desirable if the extra utility gained by one group more than offsets the loss in utility of the other group: in a more inequality-averse function though, total welfare falls.

For example, with the large increase in productivity in the US agricultural sector at the beginning of the 20th century, income in the sector fell because of labour-saving technology: the agricultural industry in the US in 1900 employed as much as 41% of the American labour force but by 2000, after strong mechanisation in the sector, the agricultural employment accounted for only 2% of total employment. If there had been perfect mobility, the displaced agricultural workers could have moved to cities and learnt the skills needed to

contribute to production in other sectors; but because of existing cost, falling wages and value of rural assets, these workers could not afford to shift from one sector to another in the short term. However, in the long term redistribution might be reached naturally, so that the benefits of technological change are fairly spread: the huge number of workers who would have been dedicated to agriculture, but who experienced the fall in labour demand in that sector, provided human capital for other occupations: this allowed other sectors to develop exponentially, especially the tertiary sector, and generate the wide range of diverse occupations that exist in modern society. Therefore, even though it causes serious negative impacts in inequality and welfare in the short run, new technology has the potential to lead to a Pareto improvement in the long run, even with redistribution costs or no imperfect mobility.

For this reallocation of workers to happen, they have to be displaced in the first place: the second effect this paper will analyse is, in fact, the displacement effect caused by technological change.

2.3 The displacement effect

It is undeniable that the intervention of robots in the production process will lead to lower costs of production, higher productivity and higher GDP: yet, this quest for better performing factors of production, which can be identified as one of the drivers of technological change, could be the very reason for a major disruption in the labour market.

If wages are sufficiently high, output producers will find it convenient to substitute workers with machines since, as discussed before, the price of new technology lowers as the level of technological change increases. Because of this, workers of automated tasks face the risk of being unemployed. The labour demanded by new occupations might absorb the newly unemployed, causing only a temporary impact which could even lead to long-run benefits; though if the new additional labour demand does not have a sufficient magnitude to utilize all unemployed workers, the impact persists in the long term causing serious hindrances.

2.3.1 Advantages of short-run displacement

In the short run, since the job market would not have the time to adjust to new technology, the situation may be different and artificial intelligence might cause temporary job displacement. In this case, AI would cause costs of production to decrease and productivity - also due to the fact that a robot cannot get ill, fatigued or injured, and (potential) quality of the product to increase, benefiting the economy as a whole despite the rise in income inequality. In fact, if the use of new technology causes firms to produce at lower costs, consumers will have a higher level of disposable income to spend, increasing demand for labour in sector other than the ones at risks. Displaced workers would then, in the long run, specialise in existing or newly created non-

automatable jobs (if any), with the possibilities of shifting to less heavy and/or dangerous jobs that can be performed by machines.

This phenomenon has been occurring for a very long time: before the 20th century, in the room where needlewomen were sewing fabric, there used to be a man sitting on a tall chair and reading the newspaper to them to keep them entertained. This job has evolved, for example, into the radio host, and changes and evolvements as such are very likely to keep arising. Another example, this time concerning a seriously heavy and alienating job, can be found back in 1913 with the introduction in Ford factories of the assembly line, which reduced the average time of assembly of a car body from 12.5 hours to 5 hours: certainly some factory workers lost their job in favour of machines, but new occupations were created, since specialised technicians were needed for the maintenance of the automated parts of the production line. The car assembly line process evolved, in 2018, with Mercedes, BMW and Audi discarding heavy industrial robots to replace them with light, flexible robots. This major change was implemented to meet a change in consumers' luxury goods preferences: a growing demand for customised or tailor-made products. In fact, these new, mobile robots are used to forge a hybrid process between mass production, carried out by industrial robots, and human customisation. Since human dexterity and creativity cannot be imitated by machines yet, robots collaborate with human workers, required in this new customised mass production: this need was not present with the previous robots since all tasks carried out were routine, so this evolvement of technology actually led to job creation.

The examples presented above involve blue-collar routine and repetitive occupations, while creative whitecollar tasks displacement situations can be found with much less ease. In fact, a further analysis can be done distinguishing between blue-collar and white-collar jobs: according to data from the Bureau of Labour Statistics, jobs related to a low wage have an extremely higher probability of being automated than occupations which yield a high hourly wage, and the same reasoning can be applied to low-education jobs, where the probability of automation is high and decreasing as years of education increase up to the point of graduate degree-level occupations, which were found to be completely non-automatable.







This could imply that low-paid occupations will be fully automated, making possible a shift of human workers to higher-education, higher-wage jobs, which will in the long run benefit both personal incomes and the economy as a whole, increasing the efficiency for both machine labour and human labour. Moreover, temporary displacement in these kinds of jobs due to technological change are often followed by more than offsetting gains in employment, as happened in the 20th century with the mechanisation of agriculture.

2.3.2 Impairing persistent displacement

New technology is more likely to displace low-skilled and middle-skilled workers: if the new tasks created do not require the same or similar skills to the ones needed to perform the tasks automated, workers have to acquire the new capabilities demanded. This may take some time and may not always be possible. Because of this fall in employment, the increase in productivity will not be matched with a proportional increase in demand for labour: output per worker will decrease – depending on the magnitude of the productivity effect, tapering off the wages of those who were not displaced.

If the productivity effect does not compensate for the unemployed, this may be an example of a jobless growth. An excessively high substitution of human workers in favour of machines, may cause a job displacement so high that industries are not able to absorb all the ones who remained jobless, causing unemployment to rise dramatically even in the long term, with a first disruptive impact in labour-intensive countries that would shift towards a more capital-intensive economy due to more cost-efficient solutions.

A clear example can be found in the Indian manufacturing industry. After the 1980s, India's output grew at a rate of over 7%, while employment only had a growth rate of 0.53%. The explanation lies in the fact that capital intensity (K/L) grew during those years at a rate of 4.3%: this implies that at least some of the increase in productivity was due to technological upgradation which substituted labour for machines. Moreover, the share of wages to workers fell by a half but share wages to supervisory and management only fell by a quarter, showing that low-paid blue-collar workers were those who suffered the most from this differential in growth rates.¹¹

Yet white-collar jobs, despite being intensive in non-automatable tasks, are not out of risk. Nowadays machines are able to recognise patterns which are learnt from training data sets and rearranged into algorithms. The higher the number of training datasets, the more accurate the predictive algorithm, the better the performance. Since we are currently in the age of Big Data, AI-powered technologies are able to "upgrade" the original instructions with which they were programmed: potentially in the future they will

¹¹ Kannan, K. P., & Raveendran, G. (2009). Growth sans employment: A quarter century of jobless growth in India's organised manufacturing. *Economic and Political weekly*, 80-91.

have the ability to create and programme new machines without human aid, outperforming human workers at some, if not all, types of occupations. So far, in fact, an important characteristic of technological improvement was not considered in the analysis, since it only began to rise in the second half of the 20th century, which is Machine Learning. This new characteristic of technology has the potential to hinder high-skilled, high-educated workers just as much as low and medium-skilled ones, by gaining the ability of expanding the instructions to follow outside the boundaries of the initial programming algorithm.

As a result of this, even though new technology is likely to create new jobs (as analysed in the previous section), the demand for labour is bound to decrease for tasks that can be substituted my machines, that is, which can be automated. Machines are cheaper, better performing and more productive than human workers, so profit-maximising employers will substitute machines for people above a certain wage threshold. This will reshape the whole economy: productivity will boost, prices will fall, and labour demand will shift, causing a decrease in the equilibrium wage rate. If this happened, the wages of all workers, not only those whose tasks were automated, would suffer a fall: this hypothesis will be dismissed in the Caselli-Manning model in the next chapter.

Chapter III

In this chapter, two contrasting models will be explained for comparative purposes. The first model was developed by Francesco Caselli and Alan Manning in 2018 in the paper "*Robot Arithmetic: New Technology And Wages*"¹² (3.1), and provides an optimistic point of view of technological advancement such as AI and a model which supports the idea that technology will not harm the labour market but on the contrary it will enhance it under the assumptions of fixed labour, constant returns to scale, constant interest rate (as technology changes), perfect competition, different levels of technology across economies and homothetic consumer preferences (3.1.1). The results the model reaches studying the impacts of technological improvements are positive (3.1.2): the first result is that technology makes at least one type of worker better off; the second result is that if the relative price of investment goods to consumer goods does not increase, the average real wage of workers will rise; the last result is that if supply of different types of labour is perfectly elastic, then all workers will become better off. The case in which these assumptions do not hold will be analysed as well (3.1.3).

The comparison will be made with the model developed by Daron Acemoglu and Pascual Restrepo¹³ (3.2), as this will help fill the gaps that the first model might have in certain situations. One of the discrepancies is that the Caselli-Manning model observes a rise in real wages, but it does not consider the impact of new technology on unemployment. Moreover, while the first model treats all types of new technology as the same, the Acemoglu-Restrepo model makes a distinction between labour-augmenting and capital-augmenting technology, both of which, however, predict an increase in labour demand which is inconsistent with evidence. A further contradiction between the two models is that the second one finds that technological improvements always lead to an increase in demand for capital and consequently a rise in interest rate, which is a violation of the assumptions of the first model. All these contradictions show that the impacts of new technology on the job market are hard (if not impossible) to predict.

3.1 The Caselli-Manning Model

As explained by the authors, the aim of this model is to provide the tools for a general analysis of the potential aggregate impact that AI will have on jobs, but in a different way from those who already tried to assess the effects of new technology on labour. While many analyses only focus on specific cases of partial equilibria and "first-round effects" such as displaced workers, and which could therefore bear the underlying risk of omitting major factors and assumptions, this model focuses on identifying the gainers of an increase in

¹² Caselli, F., & Manning, A. (2018). Robot arithmetic: new technology and wages. *American Economic Review: Insights*.

¹³ Acemoglu, D., & Restrepo, P. (2018). Modeling automation. In AEA Papers and Proceedings (Vol. 108, pp. 48-53).

technology instead of those who suffer job displacement. This broader investigation reaches the conclusion that, under certain assumptions (fixed labour, constant returns to scale, constant interest rate, perfect competition, different levels of technology across economies and homothetic consumer preferences), an increase in the level of technology will have a positive impact on workers, regardless of its nature, whether it substitutes or complements labour and whether it is labour-augmenting or capital-augmenting; it also identifies in what cases new forms of technology could actually have a negative impact on the labour market.

3.1.1 The benchmark model and the assumptions

The benchmark model used to develop this analysis considers a production function where output is produced by labour *L*, capital *K* and technology θ , $F(L,K,\theta)$. The assumptions are constant returns to scale, perfect competition and one type of labour and one type of capital goods.

The Caselli-Manning model expands these assumptions in the following way. There is an arbitrary number of types of labour, an arbitrary number of consumption goods, intermediate goods and capital goods and an arbitrary number of types of new technology. This model is used for Acemoglu's model as well as this one.

The price of consumer and intermediate goods is a vector denoted by p and the price of capital goods is denoted by p^{K} .

The assumptions of the model are the following.

- 1. Labour is the only fixed factor. As mentioned above, there is no fixed number of types of labour, and this factor is denoted by a row vector *L*. What is fixed, on the other hand, is the supply of the single types of labour. The price of the various types of labour is wages, denoted by a column vector *w*. The supply of labour is considered inelastic, so that workers will work for a wage equal to zero.
- 2. Production has constant returns to scale. The total cost function of producing all types of goods depends on the variables w, p, p^K and the level of technology, which is denoted by θ . As has been discussed previously, technology lowers the cost of production of goods, therefore the total cost function is non-increasing under the ceteris paribus condition and will have a negative first derivative. Naming *c* the consumption and intermediate goods cost function and c^i the investment goods cost function, we have that

$$c_{\theta} = \frac{\partial c(w, p, p^{K}, \theta)}{\partial \theta} \le 0$$
 and $c_{\theta}^{i} = \frac{\partial c^{i}(w, p, p^{K}, \theta)}{\partial \theta} \le 0$,

where c_{θ} and c_{θ}^{i} are the marginal cost of technology for consumption and intermediate goods and capital goods respectively. In addition, technology improvements are assumed to lower the price of

consumption goods at a given level of wages, both directly because of the perfect competition assumption (point 4) or indirectly, due to the lower costs of production of the other two types of goods.

3. New technology does not influence interest rate (r) of financial assets. This assumption causes the long-run supply of capital to present perfect elasticity. Moreover, if capital is assumed to have a constant depreciation rate δ and pⁱ represents price of investment goods, then

$$p^{K} = (r + \delta)p^{i}$$

4. *The markets are perfectly competitive.* Because of perfectly competitive input and output markets, price of each good is equal to its marginal cost:

$$p = c(p, (r + \delta)p^{i}, \theta)$$
$$p^{i} = c^{i}(w, p, (r + \delta)p^{i}, \theta)$$

5. *The economies compared present diverse degrees of technology in steady state.* As discussed before, these different levels of technology increase productivity, so an increase in the level of technology affects the production function so that output is increased:

$$\frac{\partial F(L,K,\theta)}{\partial \theta} > 0$$

In addition, it is possible that an increase in the level of technology reduces the marginal product of labour (differently from what was considered in section 2.2.1) such that:

$$\frac{\partial^2 F(L,K,\theta)}{\partial L \partial \theta} < 0$$

6. Consumers' preferences are homothetic. If the utility function is homothetic, for any income level, the relative demand depends only on the relative price of the goods. Because of this, there is a single CPI which is denoted by e(p) and the utilities of the several types of workers are represented by a column vector denoted by u^w .

3.1.2 Findings

The results of this model are based on the long-run comparison of steady-state wages of economies with different levels of technology advancement. Firstly, since markets are perfectly competitive, each factor of production earns its own marginal product, so wages will be given by the equation:

$$w = \frac{\partial F(L, K, \theta)}{\partial L}$$

Result 1: Improvements in technology cannot harm all types of labour.

To reach this result, consider the maximum price and the maximum wage given an increase in technology¹⁴. Given prices and wages, technology decreases $costs^{15}$, therefore for a given wage level, prices must decrease as technology increases. Hence, the largest increase in wage will always be greater than the largest increase in price, since an improvement of technology, regardless of its nature, leads to higher real wages in at least one case: in the situation of the benchmark model, where there is only one type of labour, all workers will be better off because of higher real wages; in this model, since there is an arbitrary number of types of labour, at least one type of workers will be better off – the one associated to the highest increase in wages. Therefore,

$$\frac{\partial \mathrm{log}\rho}{\partial \theta} \leq \frac{\partial \mathrm{log}w}{\partial \theta}$$

Result 2: Improvements in technology raise the average real wage of workers if the price of investment goods does not increase relative to consumption goods.

Result 1 alone is not enough to discard the chance that almost all types of labour will be worse off. Result 2 reaches the conclusion that, under the conditions described, average wage increases. Given assumption 6, the total expenditure of each type of workers can be expressed by the function $e(p)u^w$, and in equilibrium it must equal total income (wages for each type of workers). Real wages can be thought of as the total utility of workers, Lu^w . In the paper, it is proven that:

$$Ldu^{w} = pX^{k} \left[\frac{d\tilde{p}}{\tilde{p}} - \frac{d\tilde{p}^{i}}{\tilde{p}^{i}} \right] - \left[Xc_{\theta} + Ic_{\theta}^{i} \right] d\theta q$$

where \tilde{p} and \tilde{p}^{i} are relatively consumer price index and investment goods price index¹⁶. The first term in square brackets, which represents the difference in rates of inflation, only leads to a positive change in total utility of workers if positive: this happens when prices of investment goods, used to offset the loss of value of depreciation on capital goods, fall at a faster pace than price of consumer goods, for example because consumer goods are more labour-intensive.

The intuition behind this result is that the gain generated by the extra output that technology allows to produce can go to capital or labour. Because of perfect competition, new capital receives its marginal

¹⁴ Prices of consumption and investment goods are combined into a vector ρ . The maximum change in price is given by $\frac{\partial \log \rho^{max}}{\partial \theta}$ and the maximum change in wage is given by $\frac{\partial \log \varphi^{max}}{\partial \theta}$.

¹⁵ See assumption 2.

¹⁶ X denotes the vector of consumption and intermediate goods and X^k is the consumption of capitalists.

product so the extra gain must go to labour or existing capital. But if improvements in technology cause the price of investment goods to fall relative to price of consumption goods, return to capital falls and consequently return to labour rises.



As can be seen from the graph, the relative price of investment goods has been steadily decreasing, therefore the predictions of the model find a consistent basis in the real-world data.

Result 2 proves that the average wage rises, but it does not consider the variance among wages of different types of worker: even though the mean rises, there might be a huge difference between the highest and the lowest wage and the number of gainers might be minuscule, meaning that improvements in technology might lead to severe distributional effects. This problem is solved by the last result.

Result 3: If labour of different types is perfectly elastic in supply, then workers of all types must gain from technological progress.

To explain this result, a change in the assumption related to labour supply must be made. Up to this point, labour was considered fixed, but now this assumption is relaxed. In fact, in the long run a fixed labour supply is not possible, since all the new workers that enter the labour market have the choice of which type of labour to opt for, depending on many both quantitative and qualitative factors such as wages and nature of the job. It is now assumed that labour supply is perfectly elastic, which is drastic but which significantly approximates reality, since there have often been very large changes in level of employment in different types of labour with moderate relative wage changes. Perfectly elastic supply implies fixed relative wages across different types of labour.

That being said, if the supply of labour is perfectly elastic, all wages change together, so relative wages of different occupations are fixed, and they can be treated as a single type of labour. As said in result

1, in the case of the benchmark model where there is only one type of labour, wages of all workers increase as technology improves: therefore, in case of perfectly elastic labour supply, an increase in the level of technology makes all workers better off.

3.1.3 Deviating from the assumptions

This model holds only if the benchmark model assumptions of perfect competition, constant return to scale, and constant interest rate hold. The analysis proceeds by outlining the potential effect of an improvement in technology in case one of these assumptions is violated.

✤ Decreasing returns to scale.

Decreasing returns to scale are often thought to exist because of an omitted fixed factor. In this case, it can be said that an improvement in technology leads to an increase in returns for fixed factors as a whole, but labour is not the only fixed factor. In this case, it is sensible to assume that the benefits that technology leads to go to the owner of the omitted factor instead of workers.

✤ Imperfect competition.

In the goods market, imperfect competition would cause the prices to increase due to the mark-up above marginal cost. The total income of labour would then be:

$$WL = (1 - \mu)F(L, K, \theta) - P^{K}(r + \delta)K$$

Where μ is the mark-up. This does not affect Result 1 if mark-ups across different goods are constant, but this is not likely to occur. If a higher technology level leads to an increase in the mark-ups for some goods, real wages are likely to fall because of grater price differentials.

Rising interest rate.

As discussed before, new technology leads to a higher productivity and therefore to a higher rate of economic growth, which is expected to rise interest rate and therefore decrease real wages. However, this model focuses its analysis on steady states so, even though it is not specified in the assumptions, it is necessary to mention examples of what could happen in case of a dynamic economy.

✤ Non-steady states.

This model does not consider the transition from one steady state to another. In case of a technological singularity, for example, labour would not be a fixed factor anymore and under perfect competition wages and prices would fall to zero. Since there would be no limits for production imposed by fixed labour, the economy would gradually reach the state of total abundance, so a transition model would be necessary.

3.2 Comparison with the Acemoglu-Restrepo Model

The model explained above works perfectly for existing-like technology but, as implied by the last point, the economy could develop in a way which is completely different from what has occurred in the past, since a type of technology such as AI might be disruptive as it may evolve in such a way that it will become as good or even better than humans at carrying out tasks. Therefore, this model might be flawed. Because intelligent robots are still at their outset, there is no real evidence that can help us predict what is going to happen in the next few decades, so we have no choice but to rely on models such as this. Moreover, this model proves that average real wages will rise in the long run, but it does not consider the impact of new technology on unemployment. In fact, even if real wages increase on average, many workers are likely to be substituted by machines which are able to imitate human work. In other words, while the Caselli-Manning model takes into consideration the productivity effect, it does not consider the displacement effect, its negative counterpart that is as likely to arise as technology advances.

This comparison, on the other hand, is widely discussed by the models developed by Daron Acemoglu and Pascual Restrepo, who affirm that, if the productivity effect is not large enough to offset the displacement effect, labour share of income inevitably falls, causing a decrease in average real wages. Moreover, while the previous model treated all types of technological change the same regardless of its nature, this analysis makes a distinction between labour- and capital-augmenting new technology: it proves that if the technological change is capital-augmenting, equilibrium wages, labour demand and consequently labour income would rise, but this conceptualisation is neither logical nor supported by evidence; in case of labour-augmenting new technology, which is more likely to be in case of robots and AI, according to the model there would be an increase in labour demand and equilibrium wages but a decrease in labour share of income, which again is inconsistent with real evidence of the effect of higher level of automation on labour demand, since substitution of human workers in favour of robots has always lead to a less labour-intensive production process.

Another contradiction between the two models is that, according to the Acemoglu-Restrepo analysis, automation always increases demand for capital and therefore increases equilibrium rental rate: this represents a violation of the assumptions on which the Caselli-Manning model is built upon. Since $r(t) = R(t) - \delta$, where r(t) is the interest rate and R(t) is the rental rate of capital, during the transition between one equilibrium to another, interest rate would be increasing, contradicting assumption 3.

This comparation is enough to show that, despite the many attempts to develop a general and comprehensive framework, it is hard to assess the effects that new technology will have on the labour market, and it is probably too early to make accurate predictions based on existing data and evidence because of the peculiar characteristics of technological advancements that involve Artificial Intelligence.

Conclusions

There is no unanimous opinion on whether a technology as different as the ones developed so far as AI will actually cause unpredictable disruptions in the economy or whether instead it will interact with the job market as automation has done since the First Industrial Revolution – improving living and working conditions and boosting up productivity with no major consequence on long-term unemployment. The Artificial Intelligence used so far does not seem to preannounce any greater capability of completely substituting human workers in performing all tasks. Most certainly, as has been assessed with the analysis of the polarisation of wages and occupations, some tasks – middle-skill in particular, are the most automatable because of their routineness and therefore workers performing middle-wage occupations are the most likely to be displaced. If AI impacts the economy as all technology so far, technological unemployment will only be a transitory matter, and might actually lead to improvements in the level of education and wages of workers, in addition to better working conditions because of machines performing the riskiest and most fatiguing tasks. If on the other hand AI will not behave as pre-existing technology, the implications on unemployment and total welfare might be more serious, causing long-lasting technological unemployment and threatening all-skill workers to face job displacement because of the number of unemployed so high that the market cannot reallocate them all.

Regardless of the persistence of technological unemployment, all new technology, and therefore Artificial Intelligence as well, entails gains due to the increase in productivity, which more or less offset the losses due to the displacement effect. The productivity effect, in addition to allowing for the production of more output at a lower cost, also substantially benefits workers performing tasks complementary to the ones being automated, due to higher wages and higher efficiency – because of the absence of routine tasks that take time and that are instead performed by machines. Despite this positive effect, an increase in productivity due to technological improvement is not Pareto-improving, since this would require perfect mobility of workers and no costs of redistribution in the economy: these two requirements are not met in real world, so some types of workers will always be made worse off by new technology.

This proposition however is not consistent with the Caselli-Manning model, who found that if labour supply is perfectly elastic and if the prices of investment goods relative to consumption goods (fact that is consistent with evidence as shown in Graph 10) all wages experience a rise, so that the average real wage increases. However, this result is reached because this model only considers the effect of new technology on wages, whilst not considering displacement. Moreover, some assumptions of the model might be too stern compared to the real-world economy. For these reasons, some inconsistencies were found between this model and the second one analysed, the Acemoglu-Restrepo model. First of all, while CM consider all types of technology the same, AR make a distinction between labour- and capital-augmenting, where both should lead to an increase in labour demand but which does not happen in reality, consistently with the consideration that new technology is not Pareto-improving. Furthermore, according to AR new technology causes the interest rate to

rise, which is a violation of assumption 3 of the CM model. Lastly, CM treats Artificial Intelligence as type of technology that will affect the economy as any other technology has done in the past: the comparison between the two models shows that this may not be true, and that the aggregate effects on total welfare of technologies currently being developed cannot be predicted easily to form a reliable framework.

References

Acemoglu, D. (2012). Introduction to economic growth. Journal of economic theory, 147(2), 545-550.

Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics* (Vol. 4, pp. 1043-1171). Elsevier.

Acemoglu, D., & Restrepo, P. (2017). Robots and jobs: Evidence from US labor markets.

Acemoglu, D., & Restrepo, P. (2018). Modeling automation. In *AEA Papers and Proceedings* (Vol. 108, pp. 48-53).

Agrawal, A., Gans, J. S., & Goldfarb, A. (2019). *Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction* (No. w25619). National Bureau of Economic Research.

Autor, D. (2014). *Polanyi's paradox and the shape of employment growth* (Vol. 20485). Cambridge, MA: National Bureau of Economic Research.

Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.

Brynjolfsson, E., & Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science*, *358*(6370), 1530-1534.

Bahrin, M. A. K., Othman, M. F., Azli, N. N., & Talib, M. F. (2016). Industry 4.0: A review on industrial automation and robotic. *Jurnal Teknologi*, 78(6-13), 137-143.

Bessen, J. E. (2016). How computer automation affects occupations: Technology, jobs, and skills. *Boston Univ. school of law, law and economics research paper*, (15-49).

Caballero, R. J., & Hammour, M. L. (1998, June). Jobless growth: appropriability, factor substitution, and unemployment. In *Carnegie-Rochester conference series on public policy* (Vol. 48, pp. 51-94). North-Holland.

Campa, R. (2014). Technological growth and unemployment: A global scenario analysis. *Journal of evolution and technology*, 24(1).

Caselli, F., & Manning, A. (2018). Robot arithmetic: new technology and wages. *American Economic Review: Insights.*

Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard business* review, 96(1), 108-116.

David, H., & Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review*, *103*(5), 1553-97.

David, H. J. J. O. E. P. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of economic perspectives*, 29(3), 3-30.

Feldmann, H. (2013). Technological unemployment in industrial countries. *Journal of Evolutionary Economics*, 23(5), 1099-1126.

Frey, C. B., & Osborne, M. (2013). The future of employment.

Graetz, G., & Michaels, G. (2018). Robots at work. Review of Economics and Statistics, 100(5), 753-768.

Greenwood, J. (1997). *The third industrial revolution: Technology, productivity, and income inequality* (No. 435). American Enterprise Institute.

House, W. (2016). Artificial intelligence, automation, and the economy. Executive office of the President.

Kannan, K. P., & Raveendran, G. (2009). Growth sans employment: A quarter century of jobless growth in India's organised manufacturing. *Economic and Political weekly*, 80-91.

Keynes, J. M. (2010). Economic possibilities for our grandchildren. In *Essays in persuasion* (pp. 321-332). Palgrave Macmillan, London.

Kim, Y. J., Kim, K., & Lee, S. (2017). The rise of technological unemployment and its implications on the future macroeconomic landscape. *Futures*, 87, 1-9.

Korinek, A., & Stiglitz, J. E. (2017). *Artificial intelligence and its implications for income distribution and unemployment* (No. w24174). National Bureau of Economic Research.

Lindert, P. H., & Williamson, J. G. (1983). English workers' living standards during the industrial revolution: a new look. *The Economic History Review*, *36*(1), 1-25.

McGaughey, E. (2018). Will robots automate your job away? Full employment, basic income, and economic democracy. *Centre for Business Research, University of Cambridge, Working Paper*, (496).

Michaels, G., Natraj, A., & Van Reenen, J. (2014). Has ICT polarized skill demand? Evidence from eleven countries over twenty-five years. *Review of Economics and Statistics*, *96*(1), 60-77.

Mitchell, M. F. (2005). Specialization and the skill premium in the 20th century. *International Economic Review*, *46*(3), 935-955.

Mokyr, J. (1998). The second industrial revolution, 1870-1914. Storia dell'economia Mondiale, 219-45.

Nübler, I. (2016). New technologies: A jobless future or golden age of job creation. *International Labour Office Research Department Working Paper*, *13*, 22-23.

Ricardo, D. (1891). Principles of political economy and taxation. G. Bell.

Robinson, J., & Acemoglu, R. (2012). Why nations fail. Crown Publishing Group.

Seyfried, W. (2011). Examining the relationship between employment and economic growth in the ten largest states. *Southwestern Economic Review*, *32*, 13-24

Wagener, H. J., Drukker, J. W., & Marx, K. (1986). *The economic law of motion of modern society: a Marx-Keynes-Schumpeter centennial*. CUP Archive.

Wilson, J. M., & McKinlay, A. (2010). Rethinking the assembly line: Organisation, performance and productivity in Ford Motor Company, c. 1908–27. *Business History*, *52*(5), 760-778.

Wood, J. C. (Ed.). (2004). Karl Marx's economics: critical assessments (Vol. 1). Routledge.