



HAL
open science

Robotics and artificial intelligence: what impacts on employment and inequalities ?

Florent Bordot

► **To cite this version:**

Florent Bordot. Robotics and artificial intelligence: what impacts on employment and inequalities ?. Economics and Finance. Université de Strasbourg; Scuola superiore Sant'Anna di studi universitari e di perfezionamento (Pise, Italie), 2022. English. NNT : 2022STRAB011 . tel-04359797

HAL Id: tel-04359797

<https://theses.hal.science/tel-04359797v1>

Submitted on 21 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

**UNIVERSITÉ DE STRASBOURG
SCUOLA SUPERIORE SANT'ANNA**

ÉCOLE DOCTORALE ED221 AUGUSTIN COURNOT

[UMR 7522]

THÈSE présentée par :

[**Florent Bordot**]

soutenue le : **20 juin 2022**

pour obtenir le grade de : **Docteur de l'université de Strasbourg**

Discipline/ Spécialité : Sciences économiques

**Robotics and artificial intelligence: what
impacts on employment and inequalities?**

THÈSE dirigée par :

M. LLERENA Patrick

M. MINA Andrea

Professeur, Université de Strasbourg (France)

Professeur, Scuola Superiore Sant'Anna (Italie)

RAPPORTEURS :

M. PIANTA Mario

Mme SAVONA Maria

Professeur, Scuola Normale Superiore (Italie)

Professeur, University of Sussex (Royaume-Uni)

AUTRES MEMBRES DU JURY :

M. LAMPERTI Francesco

M. LANG Dany

M. LORENTZ André

MCF, Scuola Superiore Sant'Anna (Italie)

MCF, Université Paris 13 (France)

MCF, Université de Strasbourg (France)



Sant'Anna
Scuola Universitaria Superiore Pisa

Université

de Strasbourg

ROBOTICS AND ARTIFICIAL INTELLIGENCE: WHAT IMPACTS ON EMPLOYMENT AND INEQUALITIES?

Author

Florent Bordot

Supervisors

Andrea Mina

Patrick Llerena

Acknowledgements

I would like to thank my thesis supervisors, Patrick Llerena and Andrea Mina, who guided me along the stages of this doctoral thesis.

I would like to give special thanks to André Lorentz, who has advised me to attend the PhD program in Pisa, accompanied and encouraged me during these four years. This thesis is the outcome of a long process of exchange of ideas initiated five years ago, when I was still a Master student.

I also thank Mauro Napoletano, who agreed to be a member of my thesis monitoring committee, and Mario Pianta, Maria Savona, Dany Lang and Francesco Lamperti who accepted to be members of the jury.

Thanks to Pierre, Louis, Kenza, Michele, Luca and Marie with whom I shared the daily life of a PhD student. I will keep fond memories of the time we spent together.

Last but not least, I'm thankful to my family for always believing in me. A special thanks to Julie for her unconditional love and support.

Contents

Introduction	7
1 AI, robots and unemployment: evidence from OECD countries	22
1.1 Introduction	22
1.2 Literature review	24
1.3 Research hypotheses	25
1.4 Methodology and data	26
1.5 Results	28
1.6 Concluding remarks and issues for further research	35
2 Automation and labor market polarization in an evolutionary model with heterogeneous workers	38
2.1 Introduction	38
2.2 The Model	44
2.2.1 Production	45
2.2.2 Demand, Market Dynamics and Pricing Behavior	46
2.2.3 Employment, wages and the labor market	48
2.2.4 Productivity gains, Individual learning and Innovation	51
2.2.5 Entry/exit mechanism	54
2.2.6 Timeline of the model	55
2.3 Simulation Results	56
2.3.1 Initialization and simulation protocol	56
2.3.2 Polarization measurement	56
2.3.3 Empirical validation	57
2.3.4 Baseline scenario	59
2.3.5 Automation scenario	62
2.3.6 Slow automation scenario	65
2.3.7 Summary of the results	66
2.4 Conclusions	68
3 Is labor market polarization inevitable? Some answers from a policy simulation exercise.	71
3.1 Introduction	71

3.2	The core of the model	76
3.3	Extensions and policy experiments	78
3.3.1	Labor market flexibility	78
3.3.2	Redistributive policy	79
3.3.3	Pay-settings institutions policy	80
3.3.4	Training policy	81
3.4	Results	82
3.4.1	Initialization and baseline scenario	82
3.4.2	Firing and resignation regulation (FRR)	82
3.4.3	Minimum wage (MW)	84
3.4.4	Unemployment benefits (UB)	86
3.4.5	Training policy (TP)	87
3.5	Results summary	88
3.6	Conclusions	90
	General conclusion	92

List of Figures

1	Average hours worked and real GDP per capita in OECD countries, 1970 – 2019	11
1.1	Evolution of the average stock of robots (panel (a)) and AI patents stock (panel (b))	27
2.1	The structure of the model	44
2.2	Evolution of firms’ productivity (panel (a)) and Real GDP (panel (b)) - Base 100 in step 0 - 10-periods moving average - Baseline scenario	58
2.3	The unemployment-inequality nexus	58
2.4	Automation spikes	59
2.5	Polarization indexes (equations 3.12 and 3.11) - Baseline scenario	60
2.6	Wage distribution Step 1 (panel (a)) vs. Step 250 (panel (b) - Baseline Scenario	60
2.7	Skill index distribution Step 1 (panel (a)) vs. Step 250 (panel (b) - Baseline Scenario	61
2.8	Median skill, median and average real wage evolution - Baseline - Base 100 on step 0	61
2.9	Polarization indexes - Automation scenario	62
2.10	Dynamics of the Gini indices	63
2.11	Wage distribution (panel (a)) and Skill index distribution (panel (b) - Automation scenario - Step 250	63
2.12	Median skill and median real wage evolution - Automation scenario	64
2.13	Share of automated tasks - Automation and slow automation scenarios .	65
2.14	Dynamics of specialization indexes	67
3.1	Prices indices (panel (a)) and vacancy rates (panel (b)) - Baseline and FRR scenarios	83
3.2	Real minimum wage and labor productivity (panel (a)) Wage polarization indices (WPI) dynamics(panel (b) - Baseline and MW scenarios	85
3.3	Debt ratio - UB and HUB scenarios	86
3.4	Dynamics of polarization indexes and keys macroeconomic variables - PC vs Baseline scenarios	91
3.5	Real GDP - Baseline scenario - 1000 steps and 100 replications	99
3.6	Labor productivity - Baseline scenario - 1000 steps and 100 replications	99
3.7	Structure du modèle du Chapitre 2	127

List of Tables

1	Firm-level studies	18
2	Industry-level studies	19
3	Regional-level studies	20
1.1	<i>Panel estimates</i> - $U_{i=0}^{k=25-64}$ and $U_{i=1}^{k=25-64}$	30
1.2	<i>Panel estimates</i> - $U_{i=2}^{k=25-64}$ and $U_{i=3}^{k=25-64}$	31
1.3	$U_{i=1}$ differentiated by age group - GMM-SYS estimates	32
1.4	$U_{i=2}$ differentiated by age group - GMM-SYS estimates	33
1.5	$U_{i=3}$ differentiated by age group - GMM-SYS estimates	34
2.1	<i>Automation and baseline scenarios</i>	64
2.2	<i>Slow automation scenario versus automation and baseline scenarios</i>	66
3.1	<i>Automation and baseline scenarios</i>	82
3.2	<i>FRR and baseline scenarios</i>	84
3.3	<i>Automation and baseline scenarios</i>	85
3.4	<i>LMW = Low Minimum Wage scenario</i>	85
3.5	<i>UB and baseline scenarios</i>	87
3.6	<i>HUB, UB and baseline scenarios</i>	87
3.7	<i>TP and baseline scenarios</i>	88
3.8	<i>Results summary</i> - "/" = No significant effect ($p\text{-value} > 0.05$)	88
3.9	<i>PC and baseline scenarios</i>	89
3.10	$U_{i=1}$ differentiated by age group - FE estimates	96
3.11	$U_{i=2}$ differentiated by age group - FE estimates	97
3.12	$U_{i=3}$ differentiated by age group - FE estimates	98
3.13	Main equations of the model	100
3.14	Parameters setting.	101
3.15	Data used to initialize the model.	102
3.16	Parameters setting. - = same value than in the baseline scenario / = absence of the variable in this scenario	103
3.17	<i>FRR and TP scenarios</i>	104
3.18	<i>FRR and MW scenarios</i>	104
3.19	<i>FRR and UB scenarios</i>	104
3.20	<i>MW and UB scenarios</i>	105

3.21	<i>MW and TP scenarios</i>	105
3.22	<i>UB and TP scenarios</i>	105
3.23	<i>Synthèse des résultats du Chapitre 3 - "/" = Pas d'effet statistiquement significatif</i>	128

Introduction

The link between process innovation and employment is old and complex; as it is a vector of development, but also a source of instability. This antagonism makes it desired and feared, both on the market of goods and services for well-established producers using old production technology, and on the labor market for workers at risk of losing their jobs. This fear of the effects of technical progress on employment is not new: in Ancient Rome, the emperor Vespasian is said to have graciously paid an engineer who invented a machine for transporting heavy columns to the Capitol building site, while refusing to use this invention to “enable my poor carriers to earn their bread.” (Suetonius (2019)). More contemporarily, the 19th century was marked by revolts of textile artisans, demanding better wages and castigating the unfair competition created by the appearance of new machines that mechanized work and reduced costs. The most famous of these, the Luddite revolt, began in the town of Arnold in 1811 and spread throughout England in the following two years. In France, in 1831, the appearance and spread of looms that lowered production costs, and therefore prices, was one of the triggers for the revolt of the Canuts, who were craftsman who worked silk on old manual looms giving much lower yields.

The social consequences generated by technological unemployment are materialized in political concerns. In 1961, President Kennedy declared: “The major challenge of the sixties is to maintain full employment at a time when automation is replacing men” (MacBride (1967)). More recently, the subject has come back to the forefront: an economic report for the US President (Lee (2016)) was dedicated to “Artificial intelligence, automation and the economy”, and highlights the positive effects of AI on productivity, but also its destabilizing impact on the labor market. The authors of the report recommend three areas of intervention: investing in artificial intelligence because it is a major source of competitiveness and economic development, strengthening the education and training system to best adapt the skills of Americans to future jobs, and supporting workers in professional transition. Benoit Hamon, candidate for the French presidential election in 2017, focused his campaign on the effect of automation on employment, proposing a “universal income” that would provide an income for the entire population. Presidential candidate in 2020, Andrew Yang campaigned on a similar idea, called the “freedom dividend”, that would provide one thousand dollars per month to every American adult over the age of 18.

At first sight, fears of a massive increase in unemployment resulting from technological progress do not seem justified. Indeed, historically, process innovations have been

beneficial to humanity. The British agricultural revolution, which began in the 17th century, has increase considerably agricultural yields, allowing the country to escape the Malthusian trap and to rapidly increase its population which tripled between 1700 and 1850 (Pretty (1991)). The labor surplus generated by process innovations in agriculture was one of the major factors in triggering the Industrial Revolution: without significant productivity gains in the primary sector, the expansion of the secondary sector would have been constrained by labor shortages.

This link between technical progress, structural change, and economic development is summarized in the Petty-Clark law, which states that “the movement of the labor force from agriculture to manufacturing, and from manufacturing to trade and services” was the “most important concomitant of economic progress” (Clark (1940), p.76). This claim is confirmed by Nuvolari and Russo (2019) who point out that, with the exception of resource-rich and small open countries, most developed countries have gone through these stages in their development. A recent and striking example of the Petty-Clark law is China, which has managed to significantly reduce its development gap with OECD countries in record time. According to Khan and Fatima (2016), the agricultural sector accounted for 54.3% of Chinese employment in 1994, compared to only 34.8% in 2011. Over the same period, the shares of industry and services increased by 6.8 and 12.7 percentage points respectively, and according to the World Bank GDP per capita increased by more than 5 times over the same period.

One explanation for this shift of labor from the agricultural sector to industry, and then from industry to services, is provided by the spill theory developed by Alfred Sauvy (1980). The concept is that in the short term and on a microeconomic scale, technical progress destroys jobs; but in the long-run and on a macroeconomic scale, the productivity gains generated by technical progress make it possible to lower prices and thus increase purchasing power, opening the way to the consumption of new goods and services and thus to the expansion of new sectors of activity into which the surplus labor force will flow, ultimately leading to a reduction in unemployment. The case of the agricultural revolution cited above illustrates this mechanism well, with an agricultural population that represented more than a third (36.8%) of the British working population in 1759 and that only accounted for less than a quarter (23.5%) in 1851 (Broadberry et al. (2015)). Between 1876 and 1911, more than 3 million French people left agriculture (Smith (1999), Duby and Wallon (1976)).

These changes in the sectoral composition of employment were accompanied by changes in the occupation of the territory: this was the rural exodus. Contrary to what the spillover theory might predict, the transition from the countryside to the cities, from fields to factories, did not happen without a struggle. In France, vagrancy and mendicity were widespread at the end of the nineteenth century, rising from between 75,000 and 200,000 vagrants in the first half of that century to 400,000 to 500,000 in the 1890s (Smith (1999)). Without a social protection system guaranteeing everyone a minimum income to survive, the neediest had to turn to charity, which was very heterogeneous from one region of the country to another (Wagniard (1998)).

The spillover process did not stop with the secondary sector but continued in the

middle of the 20th century with the expansion of the tertiary sector. The globalization of trade, on the one hand, leading to an increasing relocation of industrial activities to developing countries, and productivity gains on the other hand, have led to a significant reduction in the share of industrial jobs in developed economies. Data from the OECD STAN database confirm this phenomenon: in France, between 1971 and 2017, the share of manufacturing jobs in total employment fell from 22.58% to 9.25%, while the share of service sector jobs rose from 53.78% to 80.73%. A similar trend can be observed in the United Kingdom, with the share of manufacturing jobs falling by 21.25 percentage points and the share of service sector jobs rising by 25.7 percentage points over the same period. In line with the spillover theory, the question arises whether we will see a spillover of service jobs into a new quaternary sector, the contours that has yet to be defined, or whether this surplus of labor will not rather lead to an increase in unemployment.

Technological unemployment and the classics

The question of the existence or not of technological unemployment is not new, and it is striking to note that founding authors of political economy have led reflections on this subject. In his famous book “The Wealth of Nations”, Adam Smith (1776) mentions that the introduction of better machines makes it possible to reduce the number of workers needed to achieve a certain level of production, but he also stresses that the productivity gains will lead to an economic expansion that will ultimately benefit real wages. Smith also emphasizes the deflationary nature of technical progress, which reinforces the purchasing power of workers, and uses a concrete example to illustrate this point: “A better movement of a watch, that about the middle of the last century could have been bought for twenty pounds, may now perhaps be had for twenty shillings” (Smith (1776) p.307).

In chapter 6 of his famous book “A treatise on political economy”, Jean-Baptiste Say (1803) recognizes that, in the short term, the introduction of machines can generate an increase in unemployment. In the long-run, however, Say believes that the net effect on employment is positive, thanks to two mechanisms: an increase in the demand for labor to produce the machines, and an increase in consumption linked to a drop in the price due to productivity gains. To illustrate this argument, Say takes the example of the introduction of the spinning-jennies into Normandy in 1789, which allowed the French cotton spinners to stay competitive with the rest of the world. Even if some of them had to reconvert, the situation in terms of employment would have been much worse, according to Say, if they had kept the old production methods, more labor-intensive and so more costly. Another example taken by Say is the invention of the printing press, which put almost all copyists out of work. The decrease in the cost of books and the increase in their distribution allowed the printing industry to develop, so that the need for labor, according to Say, exceeded one hundred times the number of copyists existing before the invention of the printing press.

In the third edition of his book “On the Principles of Political Economy and Taxa-

tion”, David Ricardo (1821) devotes an entire chapter, entitled “On Machinery”, to this subject. He concludes that mechanization does not lead to long-term unemployment if the new machines are financed by profit. On the other hand, if the machines are financed out of the “wage fund”, in other words, the budget originally intended for the wage bill, then unemployment will increase and the living conditions of the working class will deteriorate: “I am convinced, that the substitution of machinery for human labor, is often very injurious to the interests of the class of laborers” (Ricardo (1821) p.283). This last point is in contrast to the position he initially took in earlier versions of the same book, in which he adopted a vision similar to that defended by Say (1803). In December 1819, during a parliamentary speech on Owen’s plan “for ameliorating the conditions of the lower classes”, Ricardo declared: “It could not be denied, on the whole view of the subject, that machinery did not lessen the demand for labor” (Kurz (1984)). This shift in opinion reflects the emergence of doubts about the effectiveness of market mechanisms in performing their self-regulatory function.

John Stuart Mill (1848) recognizes that machines can be detrimental to the interests of workers, but for him this effect is only temporary and corresponds to periods of adjustment in the labor market. To illustrate this, he uses the example given by Say (1803) of the copyist monks who saw themselves competing with the invention of the printing press and notes that the compositors and pressmen outnumber the copyists who have been thrown out of employment. As highlighted by Leontief (1979), John Stuart Mill later reversed this position and admitted that the introduction of ever more efficient machines in the production process could, in the long-run, reduce the aggregate demand for labor.

Marx (1867) dedicated an entire chapter to the question of machinery in volume 1 of his book “Capital”. In a section of this chapter, he analyzes the compensation mechanisms and criticizes them. He uses the example of the “direct” compensation effect, according to which job losses due to the introduction of a new machine are compensated by the hiring of workers for the construction of the same machine. Marx shows that it is unlikely that the compensation will be total since, for the capitalist of the first firm to agree to buy the machine of the second firm, the acquisition price must be lower than the wage bill saved by the introduction of this machine into the production process. Knowing that the price of the machine includes the cost of the raw materials, machines, and labor necessary for its production, increased by the surplus value, the wage bill required for the construction of the machine will necessarily be lower than the wage bill saved by the acquiring firm. Marx recognizes the existence of an indirect mechanism of job creation, that he clearly distinguishes from the traditional compensatory mechanisms defended by the “bourgeois political economists”: by increasing productivity, machines allow for a tenfold increase in production and thus in the purchase of inputs, thereby boosting production in the industries selling these inputs and thus increasing the demand for labor in these sectors. To illustrate his argument, Marx uses the example of the expansion of coal and metal mining that coincided with the Industrial Revolution in British factories. Despite these partial compensations, this instability in the labor market contributes, according to Marx, to the misery of the working class.

Debates about technological unemployment also continued throughout the 20th century. In an essay entitled “Economic Possibilities for our Grandchildren”, Keynes (1930) popularized the term “technological unemployment” by calling it a “new disease”, and predicted a future characterized by a strong reduction in working hours and an increase in non-economic activities. Even if, in the short term, the adjustments could be painful, notably through an increase in the unemployment rate, Keynes remained optimistic and declared that thanks to productivity gains, humanity was on the right track to solving its “economic problem”, i.e. to achieve a sufficient production capacity to satisfy the needs of the entire population. While Keynes was a bit too optimistic by predicting that technological progress would lead to a 15-hours-workweek, he however remarkably well captured the overall downtrend in working time and uptrend in wealth creation. Using the OECD database, we plot the evolution of the average annual working hours and real GDP per capita in Figure 1; and we can observe that the average working time has decreased by 11% between 1970 and 2019 while real GDP per capita has increased by 145% over the same period; two evolutions consistent with the trends predicted by Keynes in 1930.

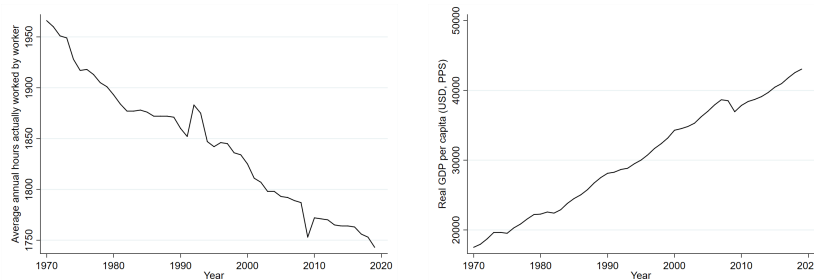


Figure 1: Average hours worked and real GDP per capita in OECD countries, 1970 – 2019

Kaldor (1932) writes a critique of Emil Lederer (1931)’s book in which the latter establishes a direct link between technical progress and the increase in unemployment. While Kaldor does not deny that technical progress may cause economic dislocation in the short-run, he argues that the real cause of a temporary rise in unemployment is not directly technical progress, but rather a distortion of prices linked to the emergence of market power from competitive advantages gained through technical progress. Apart from the question of the quantitative effect, Kaldor reflects on the qualitative impact of technical progress on employment and the existence of feedback loops. In particular, Kaldor makes the connection between the increase in technical knowledge and productivity gains, emphasizing that the level of qualification of the workforce determines the capacity of workers to fully exploit the productive potential of machines (Lorentz (2016)). This concept, in rupture with the neoclassical approach, will be used in Chapter 2 of this thesis dedicated to the polarization of the labor market, and will constitute a key mechanism for the endogenization of productivity gains.

Schumpeter (1954) explicitly rejected the conclusions of Ricardo and Marx. For

him, innovation is a process of destruction-creation: technical progress destroys jobs in some sectors but creates more jobs in others. Innovation cannot be reduced to process innovations alone, but also includes product innovations that open up new markets and therefore create jobs. He also notes that machines do not always replace human labor, but are sometimes complementary to it by making it possible to carry out tasks previously unfeasible by workers alone. For Schumpeter, the debate on technological unemployment is no longer relevant, stating that “the controversy that went on throughout the nineteenth century and beyond, mainly in the form of argument pro and con ‘compensation’ is dead and buried: as stated above, it vanished from the scene as a better technique filtered into general use which left nothing to disagree about.” (Schumpeter (1954), p. 652).

Leontief (1979) notes that rapid productivity gains are not something new, but that their potentially adverse effects on employment were cushioned by an unprecedented decline in working time in the United States, from 67 hours in 1870 to 42 hours per week in the 1940s. Figure 1 shows that the decline in working time has continued in developed countries, but at a slower pace than in the first half of the twentieth century, with an average annual working time reduction of 8.5 percent between 1978 and 2019. Leontief advocates wage increases to enable workers to reduce their working time. Leontief et al. (1986) perform an input-output model of the U.S. economy to estimate the impact of automation on labor. The technical coefficients are derived from dynamic tables and reflect the intensity of each input required to produce a unit of output in a given sector. The authors test four scenarios, differentiated by the expected level of investment in computers and robots, and make projections to the year 2000. The first scenario, characterized by relatively low levels of investment in computers and robotics, results in a near doubling of the total number of jobs between 1978 and 2000, while the increase is only 76% over the same period in the scenario where investment in these two technologies is the highest. It is interesting to compare these predictions with the empirical data: according to the U.S. Census Bureau data (Clark et al. (2003)), the total number of jobs in the United States was 129.7 million, an increase of only 46% since 1978, much lower than that the expectations of Leontief and Duchin in their most pessimistic scenario. This gap between predictions and reality is due in particular to the fact that the authors underestimated the pace of technical progress and productivity gains in certain sectors. For example, their projections result in a total number of between 5.3 and 5.4 million workers in the agricultural sector, regardless of the scenario studied; whereas according to the Bureau of Labor Statistics there were only 2.4 million jobs in agriculture in July 2000.

Reading the writings of Smith, Say, Ricardo, Stuart Mill, Marx, Schumpeter, Keynes, Kaldor and Leontief on the question of unemployment teaches us that, behind the plurality of opinions on the existence or not of technological unemployment, there are radically different conceptual approaches.

The Keynesian approach treats technological unemployment as any other form of involuntary unemployment, deriving its origin in a lack of aggregate demand. The solution is to stimulate it through expansionary public policy, thus raising expected demand

and so production and finally the employment level. For the structuralists, technological unemployment is structural unemployment that emerges from the mismatch between skill demand and skill supply. By impacting the qualitative aspect of labor demand, technological progress generates dislocations in the labor market. One way to address this is through education and training policies, to have a workforce with skills that can adapt quickly to these changes. Marxist analysis is very critical of compensation mechanisms and sees technological unemployment as any other kind of unemployment: a way to increase the reserve army to strengthen the power of the capitalists over the working class. The neoclassical theory differs from the three others mentioned above because it refutes the very existence of technological unemployment. In an economy where prices are perfectly flexible and factors of production are mobile, the surplus of labor generated by productivity gains will lead to a decrease in the equilibrium wage and, ultimately, labor demand increase will offset the initial decrease.

The process innovation - employment nexus

Unlike the economists of the 18th, 19th, and first half of the 20th century, contemporary economists have the opportunity to take advantage of the the progress of econometric techniques, the increase in the quantity and quality of data to statistically test the effectiveness of these compensation mechanisms. The dominant approach is the firm-level analysis, in which researchers use micro-data to determine whether innovation is detrimental or beneficial to employment. Some studies use R&D expenditures or patents as indicators of technological progress. Stam and Wennberg (2009) use data on Dutch start-ups and find that, on average, there is no effect between R&D and employment growth; however, they highlight a positive effect for high-tech firms and the 10% of firms with the fastest growth rate. Bogliacino et al. (2012) analyze European firms and find that R&D has a positive impact on employment growth. Coad and Rao (2011) use both patents and R&D spending and find that innovation is positively associated with employment growth. Van Roy, Vertesy and Vivarelli (2015) study a panel of European firms and find that patenting activity is positively correlated with job creation, but only in the high-tech manufacturing sector. The limitation of these studies is that by using patents and/or R&D expenditures, it is not possible to know which part is related to product innovation and which part is due to process innovation. Therefore, the observed effect is a net effect, and it is impossible to isolate the effect specific to process innovations, whose are particularly interesting because their effects on employment are more ambiguous than for product innovations. Indeed, there is a consensus in the literature regarding the positive impact of product innovations on employment (Antonucci and Pianta (2002), Benavente and Lauterbach (2008), Bogliacino and Pianta (2010), Crespi and Tacsir (2012), Harrison et al. (2014), Vivarelli (2014)), which open up new markets, stimulate demand, production and therefore labor demand.

Some studies make a clear distinction between product and process innovation. A first set of research highlights the adverse effect of the latter on employment. Harrison et al. (2014) find that, assuming constant output and price, process innovations tend

to reduce employment. However, they point out that when this assumption is removed, productivity gains are associated with a reduction in prices and thus an expansion of output, supporting the price compensation mechanism. Peters et al. (2014) find limited evidence that process innovation hurts employment growth, with effects confined to the manufacturing sector. Crespi et al. (2019) find that process innovation harms employment in Argentinian, Chilean, Costa Rican, and Uruguayan firms; but the effect is small. Aboal et al. (2015) find that process innovation reduces employment growth for unskilled workers and has no effect on skilled workers. Using CIS data in six European countries, Evangelista and Vezzani (2012) find that process innovation does not affect employment, except for firms in the manufacturing sector for which the impact is negative.

In contrast to these results, Zimmermann (2009) points out that process innovation has a positive impact on employment in German small and medium-sized firms, with a stronger effect than the one of product innovation. A similar conclusion is shared by Lachenmaier and Rottmann (2011) for German manufacturing firms, and by Giuliodori and Stucchi (2012) for Spanish manufacturing firms. Triguero et al. (2014) use Spanish microeconomic data covering the period 1990 and 2008, and find that persistent process innovation and employment growth are positively correlated. Peluffo (2020) finds that process innovation has a positive impact on employment in Uruguayan firms. Similar results are found by Laguna and Bianchi (2020), Okumu et al. (2019) for African firms and Zhu et al. (2021) in China.

Finally, some studies come up with mixed results. Harrison et al. (2014) find that process innovation has no impact on French and Spanish firms and have no significant effect on German and British firms. Hall et al. (2009) use Italian micro-data covering the period 1995 - 2003 and find no significant effect between process innovation and employment. A similar conclusion is shared by Benavente and Lauterbach (2008) who use data on Chilean firms between 1998 and 2001. Bianchini and Pellegrino (2019) use a dataset on Spanish manufacturing firms and conclude that there is no significant relationship between process innovation persistence and employment. Other authors (Baensch et al. (2019), Cirera and Sabetti (2019), Mitra (2020), Castillo et al. (2014)) also conclude that process innovation has a neutral effect on employment. Using a panel of 265 Italian firms over the period 1998-2010, Barbieri et al. (2019) find no significant statistical relationship between embodied technical change, a proxy for process innovation, and employment. A similar result is found by Lim and Lee (2019) for Korean firms, De Elejalde et al. (2015) for Argentine firms, and Benavente and Lauterbach (2008) in Chile. Hou et al. (2019) analyze microdata from France, Germany, the Netherlands, and China, and find that overall process innovation does not affect employment, except for German firms in the manufacturing sector where a negative effect is found.

The first lesson we can draw from this microeconomic literature is that the results are not clear, with a majority of studies finding no statistically robust link between process innovations and employment. A second observation is that this level of analysis does not allow us to draw conclusions about the net effect of process innovations on overall employment. In the case of studies concluding that process innovations have a

positive effect on employment, there is no way of knowing whether these job creations are not at the expense of job losses among their less innovative competitors. Conversely, the studies that conclude that there is a negative effect do not allow us to observe the effects on, for example, the companies that produce these process innovations. The interpretation of the results is therefore limited to the sample used and does not allow for extrapolation to estimate the effectiveness of the compensation mechanisms.

To address this issue, some authors have taken their analytical framework to the sectoral level. Antonucci (2007) uses data on 22 manufacturing sectors in 10 European countries covering the period 1994 - 2001 and finds that process innovation has a negative impact on employment, but this effect is offset by an increase in demand generated by productivity gains. This result is not shared by Bogliacino and Pianta (2010), who analyze a sample of 38 industries in 8 European countries (Germany, Spain, France, Italy, the Netherlands, Portugal, the United Kingdom and Norway) covering the period 1996 - 2004, and highlight a negative correlation between spending on innovative machinery per employee and hours worked. Using Estonian data covering the period 2001 - 2006, Meriküll (2010) finds a negative correlation between process innovation and employment. Lucchese and Pianta (2012) analyze data from 21 manufacturing sectors in Germany, France, Italy, the United Kingdom, the Netherlands, and Spain for the period 1995-2007, and show that process innovation impacts negatively employment during economic downturns. Dosi, Piva, Virgillito and Vivarelli (2021) use sectoral data covering 19 European countries over the period 1998 - 2016 and find that the replacement of means of production (machines, tools, software, etc.) exerts a negative impact on labor in sectors that are consumers of technology but do not themselves perform R&D activity to develop their own process innovation.

Sectoral analyses provide interesting insights, but if all sectors of the economy are not covered by the sample, it is impossible to determine the net effect on aggregate employment. An alternative solution to this problem is to use a macro-level approach. Feldmann (2013) analyzes a sample of 21 industrial countries over the period 1985 to 2009 and uses triadic patents (i.e., patents filed with the three main patent offices: the European Patent Office, the U.S. Patent and Trademark Office, and the Japanese Patent Office) as a proxy for technological change. Although patents are not a perfect proxy for process innovation, as product innovation can also be patented and patent strategies can artificially inflate the number of patents, Feldmann (2013) finds a clear and significant positive correlation between triadic patents and unemployment rate. This result is not shared by Matuzeviciute et al. (2017) who also use triadic patents as a proxy, but find no significant relationship with the unemployment rate in 25 European countries for the period 2000 - 2012. Evangelista et al. (2014) use the ICT indicator and find a positive correlation with employment; and a similar result is found by Marcolin et al. (2016) for 28 OECD countries over the period 2000-2011, reaching the same conclusion with the use of patents instead of ICT intensity. A major limitation of this macroeconomic literature is that, due to a lack of data and the use of proxy variables, process innovations are not clearly defined. It is therefore difficult to know whether the results obtained allow us to conclude on the positive or negative effect of process

innovations on employment.

AI, robots and employment

Among process innovations, two technologies are catching a particular interest: robotics and artificial intelligence. Although the origins of industrial robotics (Zamalloa et al. (2017), Gasparetto and Scalera (2019)) and AI as a discipline of study (Warwick (2013)) date back to the mid-20th century, the rise of these two technologies is relatively recent and coincides with the increase in computing power and the decrease in the price of hardware components.

Robotic machines are distinguished from other traditional machines by their high degree of autonomy, resulting in a considerable reduction of the manpower required in the production process. While not the only factor, the deployment of industrial robots in the manufacturing sector in many developed countries is a good example of this phenomenon, illustrated by a decrease in the number of jobs despite an increase in production. According to data from the OECD STAN database and the International Federation of Robotics, the value added produced by the French manufacturing sector has increased by +48%, employment has decreased by 28% and the stock of industrial robots has increased threefold between 1993 and 2017. Over the same period, the value added produced by the German manufacturing sector increased by 54%, employment decreased by 12% and the stock of industrial robots increased by 2.9 times. Similar trends can be observed in Italy, Spain, the UK and other OECD countries.

Artificial intelligence diverges from robotics by making the automation of cognitive tasks possible. It is similar to standard algorithms, but differs in that machine learning algorithms are evolutionary and therefore adaptive. They make it possible to automate tacit knowledge-intensive tasks, such as driving a car, because it is not necessary for programmers to code all the knowledge required to perform the task. This paradigm shift, which pushes back the frontier of automatable tasks, raises questions about a possible acceleration of automation in the service sector which has been relatively spared until now.

The literature aimed at quantifying the impacts of AI and robots on employment is relatively recent. The seminal study by Frey and Osborne (2017), who use expert assessment and determine that 47% of US jobs are at risk of automation, paved the way for studies with similar methodology; giving a percentage of jobs threatened by automation of 25% in Singapore (Fuei (2017)), 35% Finland and 33% in Norway (Pajarinen et al. (2015)), 42% in Germany and 35% in the UK (Frey and Osborne (2014)). Arntz et al. (2016) mitigate the result found by Frey and Osborne (2017) by analyzing the risk of automation no longer at the occupation level, but at the task level; finding that only 9% of jobs at risk of automation in the US. Dengler and Matthes (2018) follow a similar methodology and find that 15% of German employees are at risk. All these studies have the merit of quantifying, with a relatively wide range of results, the proportion of jobs that could be automated, but they suffer from two major drawbacks. The first one is that they are based on the opinions of experts, who are vulnerable to

a whole set of cognitive biases that are well known in behavioral economics (Thaler and Ganser (2015)), and the second one is that they treat automation in a binary and homogeneous way, without any distinction between the different sources of automation.

The use of data on robots and AI can address both limitations. Tables 1, 2 and 3 list the firm-level, sectoral and regional-level literature respectively. The analyses at the firm level (Table 1) lead to a consensus regarding the positive effect of robots on total employment (Koch et al. (2021)), Acemoglu and Restrepo (2020), Bessen et al. (2020), Domini et al. (2021), Dixon et al. (2021), Aghion et al. (2020)), although some authors raise negative effects on certain groups of workers (Humlum (2019), Bonfiglioli et al. (2020), Dixon et al. (2021)). The effect on wages is more mixed: while Humlum (2019) and Acemoglu and Restrepo (2020) find a positive effect, Aghion et al. (2020), Bessen et al. (2020), and Koch et al. (2021) conclude that it has no statistically significant effect on wages.

At the sectoral level (Table 2), the effects on employment are less clear than at the firm level. Klenert et al. (2020) and Aghion et al. (2020) find a positive effect of industrial robots on total employment, Compagnucci et al. (2019) and Carbonero et al. (2020) observe a negative impact; while Graetz and Michaels (2018) and De Vries et al. (2020) find no statistically significant effect. Regarding wages, Graetz and Michaels (2018) and Aghion et al. (2020) find no statistically significant impact, while Blanas et al. (2019) highlight a positive effect on the wages of highly skilled workers.

Finally, region-level studies (Table 3) also struggle to reach a consensus. Regarding employment, Mann and Püttmann (2018) find that automation patents are positively correlated with total employment. Chiacchio et al. (2018) and Acemoglu and Restrepo (2020) use data on the stock of industrial robots and observe a negative impact on employment. Dauth et al. (2021) provide a more nuanced conclusion, showing that robots have a positive impact on sector and service employment and negative on manufacturing employment; while Caselli et al. (2021) and Dottori (2021) find no statistically significant effect on total employment. On the wage side, Acemoglu and Restrepo (2020) find a negative impact, while Chiacchio et al. (2018) observe that there is no effect.

We note that at the firm level, while studies on process innovations lead to contradictory results, the conclusions regarding the effect of robots are unanimously positive. Some authors, however, point out that even if the effect on total employment tends to be positive, negative effects are present on low-skilled workers (Humlum (2019), Bonfiglioli et al. (2020)) or medium-skilled workers (Dixon et al. (2021)). At the industry level, the literature mostly highlights a negative impact on employment, a conclusion shared by 5 of the 8 studies presented in Table 2. Finally, at the macro level, the results obtained in both types of literature (process innovations and robots) do not allow us to conclude whether the overall effect is positive or negative.

Firm-level analysis	Variable of interest	Dataset	Period	Results
Koch et al. (2021)	Robots adoption, robot density	Spanish manufacturing firms	1990 - 2016	(+) total employment (/) wages
Acemoglu, Lelarge and Restrepo (2020)	Robots adoption	French manufacturing firms	2010 - 2015	(+) total employment (+) wages
Humphum (2019)	Industrial robots	Danish employer-employee data	1995 - 2015	(-) employment production workers (+) employment tech workers (+) wages
Bessen et al. (2020)	Automation event	Dutch private sector firms	2000 - 2016	(+) total employment (/) wages
Domini et al. (2021)	Automation-intensive goods	French manufacturing employers	2002 - 2015	(+) total employment
Bonfiglioli et al. (2020)	Imports of industrial robots	French firms	1994 - 2013	(+) total employment HS workers (-) employment LS workers
Dixon et al. (2021)	Imports of robots	Canadian firms	1996 - 2017	(+) total employment (-) employment MS workers (+) employment of LS or HS workers
Aghion et al. (2020)	Electric motive force Import of industrial machines	French manufacturing firms	1994 - 2015	(+) total employment (/) wages

Table 1: Firm-level studies

Industry-level analysis	Variable of interest	Dataset	Period	Results
Graetz and Michaels (2018)	Industrial robot density	14 industries in 17 countries	1993 - 2007	(/) total employment (/) wages (-) employment LS workers
Carbonero et al. (2020)	Industrial robot density	14 industries in 41 countries	2005 - 2014	(-) total employment
Compagnucci et al. (2019)	Industrial robots	9 industries in 16 OECD countries	2011 - 2016	(-) total employment
Blanas et al. (2019)	Industrial robots	30 industries in 10 high-income countries	1982 - 2005	(-) employment of LS or MS workers (+) employment of HS workers (+) wage share of HS workers
Klenert et al. (2020)	Industrial robot density	14 industries in european countries	1995 - 2015	(+) total employment
De Vries et al. (2020)	Industrial robot density	14 industries in 37 countries	2005 - 2015	(/) total employment (-) employment share of routine jobs
Acemoglu, Lelarge and Restrepo (2020)	Industrial robot density	21 industries in 6 OECD countries	2000 - 2015	Heterogeneous effects among sectors
Aghion et al. (2020)	Industrial robots Import of industrial machines	French manufacturing sector	1994 - 2015	(+) total employment (/) wages

Table 2: Industry-level studies

Region-level analysis	Variable of interest	Dataset	Period	Results
Acemoglu and Restrepo (2020)	Exposure to robots	Local labor markets in the U.S.	1990 - 2017	(-) total employment (-) wages
Chiacchio et al. (2018)	Industrial robot density	116 regions in 6 EU countries	1995 - 2017	(-) total employment (/) wage
Dottori (2021)	Industrial robot density	784 locals labor markets in Italy	1993 - 2016	(/) total employment
Caselli et al. (2021)	Exposure to robots	377 locals labor markets in Italy	2011 - 2018	(/) total employment
Mann and Pittmann (2018)	Automation patents	722 locals labor markets in the U.S.	1976 - 2014	(+) total employment
Dauth et al. (2021)	Exposure to robots	402 locals labor market in Germany	1994 - 2014	(-) employment in manufacturing sector (+) employment in service sector

Table 3: Regional-level studies

Limits of the literature

A major limitation of the literature is that macroeconomic studies are scarce and the existing ones do not lead to a consensus, yet it is this level of analysis that matters to policymakers. Knowing whether robots destroy jobs at the level of certain firms and sectors is a useful information, but it will not have the same importance for economic policy if the net effect on aggregate employment is positive and negative. The literature also lacks macroeconomic studies on the impact of AI on employment, mainly due to limitations in data availability. This is why, in the first chapter of this thesis, we have chosen to conduct a cross-country analysis that both the effects of robots and AI on the unemployment rate of 33 OECD countries. We will use unemployment rates differentiated by education level and age groups to capture the heterogeneity of the effects.

In the second chapter, we develop a macroeconomic model where the dynamics are micro-founded by the agents' behaviors. The objective here is to explain the mechanisms behind the polarization of the labor market induced by automation, characterized by an increase in the share of employment and wages of low- and high-skilled workers, and a decrease in the employment and wages shares of medium-skilled workers induced by automation. The explanation given by the literature is based on the routine-biased technical change hypothesis developed by Autor et al. (2003) and refined by Acemoglu and Autor (2011), according to which automation generates polarization by targeting routine tasks performed by medium-skilled workers. The problem with this explanation is that it tends to be tautological: by making the first assumption that routine tasks are mainly performed by medium-skilled workers, and then making the second assumption that automation targets routine tasks, one implicitly assumes that automation targets tasks performed by medium-skilled workers. Once these two assumptions are used in a theoretical model, it is straightforward that the dynamic will result in a polarized labor market. The aim of chapter 2 is to explore whether a theoretical model can generate labor market polarization without using the routine-biased technical change hypothesis. We also discard the limiting low-skilled/medium-skilled/high-skilled dichotomy and make no assumptions about the type of workers threatened by automation.

Finally, there is a body of research analyzing the effect of different policies on employment and wages, but the literature specifically devoted to policies seeking to reduce the effects of robotics and AI on employment and inequality is scarce. The objective of Chapter 3 is therefore to test, using the agent-based model developed in Chapter 2, four policies aimed at combating labor market polarization: a measure restricting layoffs and resignations, a minimum wage, an unemployment insurance system, and a training policy. In a final scenario, we test all the policies and show that there are synergies between them.

Chapter 1

AI, robots and unemployment: evidence from OECD countries

In this first chapter, we investigate the relation between artificial intelligence, robots and unemployment on a panel of 33 OECD countries covering the 2005 – 2017 period. We find that a 10% increase in the stock of industrial robots is associated with a 0.42 point increase in the unemployment rate. For artificial intelligence (AI), we use patents as a proxy of AI-related technological capabilities and find a positive correlation with the aggregated unemployment rate, albeit statistically weaker than the one found for robots. We then run the regressions on unemployment rates differentiated by education and age, and observe highly heterogeneous effects between groups. For example, the effect of robots is 2.5 times greater for 25-34 year-olds below upper secondary education levels than for the 55-64 year-olds with a tertiary degree. Lastly, the effect of robots is strongest on the unemployment rate of people with a medium level of education, providing some evidence that robots could contribute to the polarization of the labor market. A similar effect is found with AI, but the results are less robust than for robots.

1.1 Introduction

The resurgence of the concept of technological unemployment both in the labor economics literature and in the media highlights fears that technical progress could be job-destructive. This anxiety is rooted in the view that exponential technical progress will lead to a relentlessly growing automation of production.

Two relatively recent branches of computer sciences have come to fuel this vision: robotics and AI. The first is a continuation of the previous technological revolutions that have taken place in history, and aims to lighten the drudgery of work, to push back the frontier of technology while improving the efficiency of production. From the ancient Roman cranes that built aqueducts to the modern cranes that build ever taller skyscrapers, from the invention of printing (attributed to Gutenberg in the mid-15th century, despite evidence that Tang Dynasty China (618 - 907) had already mastered

printing techniques by the 8th century (Gunaratne (2001)), to modern printing presses that have contributed to accelerating the diffusion of information by speeding up printing while decreasing costs, from manual looms replaced by the spinning jenny in the middle of the 18th century to modern looms that increase the quantity produced tenfold, all these inventions have made human work less arduous, more productive and increased the quality of manufactured products. Robotics is part of this continuity, allowing a drastic reduction of the amount of human work per unit of capital, by automating the operation of machines.

Artificial intelligence is also in line with the search for efficiency, but differs from the previous innovations in that it does not automate manual tasks, but cognitive tasks. In this sense, it is part of the continuity of the digital revolution with the rise of computer science, making it possible to perform computational tasks easily and quickly. What is new with AI lies in its adaptive characteristics, compared to traditional algorithms, which are static. The difference could be formulated as follows: a traditional algorithm is programmed to perform a task; a machine learning algorithm is programmed to learn to perform a task. When the degree of complexity of a task is low, the first approach is much more efficient and consumes less time and resources. However, when the degree of complexity increases, the second method is much more powerful and sometimes leads to a result that was not anticipated by the developers. This type of algorithm is particularly effective for problems requiring tacit knowledge, which is, by definition, difficult to codify. For example, machine learning techniques have greatly accelerated the development of autonomous cars. Driving is a simple task, but it requires a great capacity of adaptation that was very difficult to implement with traditional algorithms that would have required anticipating all possible scenarios that the car could have faced.

The objective of this paper is to answer a very simple question: is there any statistical evidence that robotics and AI technologies could increase unemployment? We want to reason at the macro level because what matters to us here is the net effect of job destruction and creation, i.e., testing the effectiveness of compensating mechanisms. Calvino and Virgillito (2018) distinguished between two types of compensation mechanisms: classical and Keynesian-Schumpeterian. The former encompasses four different mechanisms: an increase in labor demand in the firms producing the process innovations, a drop in prices due to productivity gains that then leads to an increase in aggregate demand, a decrease in wages that in turn increases labor demand, and an increase in investment financed by the decrease in the unitary cost of production generated by process innovation, thus triggering new hiring to expand production. The second category includes two mechanisms: an increase in wages due to productivity gains for the workers remaining in the firm, thus boosting consumption; and the creation of new products that open new markets and thus increase labor demand.

The paper is structured as follows: we will present in Section 1.2 the literature review on the link between robotization, AI and unemployment, we will then develop our research hypotheses in Section 1.3, then in Section 1.4 we will present the methodology and data we used to answer them, Section 1.5 will be devoted to the analysis of the

results and we will conclude in Section 1.6.

1.2 Literature review

Opinions on technological unemployment are divided between pessimistic and optimistic views. This last position was shared by Keynes (1930), who called for a decrease in working time to avoid technological unemployment: “Three-hour shifts or a fifteen-hour week may put off the problem for a great while” (Keynes (1930) p.5). Rifkin (1995) went even further by predicting the end of work, depicting an almost fully automated economy that would produce an abundant quantity of goods and services at a marginal cost of close to zero (Rifkin (2014)).

The evidence on technological unemployment in the growing body of research on the effects of technical progress is mixed. When looking at how many jobs could be potentially automated in the next two decades, the literature provides estimates as high as 47% (Frey and Osborne (2017)) to as low as 9% for the US (Arntz et al. (2016)). Nedelkoska and Quintini (2018) used data from the PIAAC database covering 32 OECD countries and found that about 14% of jobs are highly automatable. Using the automation risk measure developed by these authors, Georgieff and Milanez (2021) analyze whether or not jobs considered at a high risk of automation in 2012 experienced a net loss of workers seven years later. They found that these occupations saw a weaker employment growth than other jobs, and even sometimes a modest decline in employment levels. However, the authors point out that, at the aggregate level, automation does not result in net job destruction. Felten et al. (2019) proposed an “AI Occupational Impact” measure and found that, while artificial intelligence has a small positive effect on wages, it has no impact on employment. Acemoglu, Autor, Hazell and Restrepo (2020) used data on online vacancies and found no correlation between AI exposure and employment and wages. This conclusion is not shared by Webb (2019), who found that jobs with high exposure to automation technologies experienced a decline in employment and wages. Building on Frey and Osborne (2017) and Felten et al. (2018), Fossen and Sorgner (2019) found that a high computerization risk is associated with an increasing likelihood of changing occupation or being unemployed.

At the firm level, Bessen et al. (2019) investigated the relation between automation and jobs and concluded that automation increases workers’ likelihood of leaving the firm. However, they note that this effect is gradual and small in magnitude. Koch et al. (2021) conducted an analysis on a panel of Spanish manufacturing firms and pointed out that robot adoption led to a net increase in jobs of 10%, while firms that did not invest in robots suffered job losses over the period. Aghion et al. (2020) used French data and found that automation has a positive impact on employment, both at the firm and industry level. Montobbio et al. (2022) matched labor-saving robotics patents with firm-level data, and highlighted a sector-specific effect on manual and cognitive tasks. In addition to studies on robotics and artificial intelligence, an entire body of literature addresses the effects of process innovations on employment. Van Roy, Vivarelli and Vertesy (2015) found that process innovation tends to increase employment,

whereas other studies observed no evidence of a meaningful relationship between the two (Harrison et al. (2014); Calvino (2019); Barbieri et al. (2019)).

While analysis at the firm level has the advantage of offering a large amount of data, often of better quality than macroeconomic data, it has the disadvantage of not providing a practical framework for answering the question of the net effect of technological progress on employment at the aggregate level. Indeed, some job losses in one firm or sector can potentially be offset by the creation of new jobs in other firms or sectors.

Meso-level and macro-level analyses allow for bypassing the limitations of firm-level studies. Dauth et al. (2018) looked at the effects of robot adoption on the German labor market, and found that it has changed the job distribution across industries without decreasing the aggregate level of employment. Graetz and Michaels (2018) used sectoral data on a panel of seventeen countries and found no significant reduction in total employment. They did, however, point out that increased robotization is associated with a decrease in the employment share of low-skilled workers. Acemoglu and Restrepo (2020) used a task-based theoretical framework to analyze the effect of robot density on local labor markets in the US, and found that one additional robot per thousand workers is associated with a drop of 0.2 percent in the employment-to-population ratio and a decrease in wages of 0.37 percent. They also emphasized the heterogeneity of the effects by showing that the lower the workers' level of education, the stronger the negative impact on jobs and hourly wages. Chiacchio et al. (2018) employed a similar methodology across six European Union countries and estimated a negative impact of 0.16 to 0.20 points on the employment-to-population ratio. They found a particularly strong displacement effect for young males with a medium level of education. Autor and Salomons (2018) turned their attention to the effects of robotization on the labor share and concluded that automation has contributed to its decline since the 1980s.

1.3 Research hypotheses

Based on the empirical and theoretical literature, we formulate four research hypotheses. First, we expect to observe a positive correlation between the aggregate unemployment rate, robots and AI (H1). While some authors (Graetz and Michaels (2018), Chiacchio et al. (2018), Acemoglu and Restrepo (2020)) have highlighted a negative (positive) impact on employment (unemployment), the quantitative effect of AI on jobs at the macro level is currently not well covered. Second, we expect a positive correlation between robots and the unemployment rate of the least educated workers, and a negative correlation between robots and the unemployment rate of those with a high level of education (H2). This hypothesis is consistent with the skill-biased technical change (SBTC) theory: unskilled workers are negatively impacted by technological progress while skilled workers benefit from it. Conversely, we test the hypothesis that AI and robots mainly have an impact on the unemployment rate of medium-skilled workers (H3). This is in reference to the literature on labor market polarization (Autor et al. (2006)) for the US, Salvatori (2018) for the UK, Furukawa and Toyoda (2018) for Japan, Goos et al. (2014) for Europe, according to which AI algorithms mainly substitute for

routine tasks performed by medium-skilled workers.

Finally, we expect heterogeneous effects within a same education group and between age classes. Two opposing effects could come into play: an “experience” effect and a “knowledge obsolescence” effect. The first effect is due to the fact that older workers tend to have more work experience and are, therefore, at a similar education level, are more qualified. As a result, the tasks performed by older workers should be more complex than those of younger workers, and thus they should be less exposed to technological unemployment. The second effect postulates that young people tend to be more technologically savvy than seniors. For example, 25-34 year-olds tend to be more comfortable with computer and mobile technologies than their 55-64 year-old counterparts with a similar level of education. This effect particularly makes sense for the most educated workers, who are most likely to be employed as data scientists, software developers, robot engineers, etc.

While we expect the former effect to prevail for the less educated, we expect the latter to prevail for the more educated. As a result, we expect a positive correlation between robots and the unemployment rate of young people with a low education level (and a negative or not significant effect on the unemployment rate of older people with the same education), and a negative correlation between AI and the unemployment rate of young people with a high education level (and a positive or not significant effect on the unemployment rate of seniors with the same level of education) (H4).

1.4 Methodology and data


To test these four hypotheses, we study the relationship between AI, robots and unemployment rates. This approach allows us to observe the net effects of compensation mechanisms at the macro level, and not just the impacts in specific firms or industries. Our dataset consists of a panel of 33 OECD countries and covers the 2005 – 2017 period, except for Iceland (2009 - 2016), Ireland (2005 - 2014) and Israel (2007 - 2017). We use unemployment rates that we regress on GDP growth which, as has been known since Okun (1963), explains a significant part of the variation in the unemployment rate (U); on the inactivity rate (IR), which can have a direct impact on unemployment by increasing or decreasing the labor force, and on the growth rate of the robots (IR) and AI patent stocks (AI). The choice of taking the first difference instead of the levels is motivated by the non-stationarity of the data, as illustrated in Figure 1.1. To account for potential lags between the variation of our robots and AI variables and the effect on unemployment, we also include the growth rate of t-1 and t-2 for these two variables. Given that unemployment can sometimes keep increasing even as its causes have faded (known as the “hysteresis effect”), we also include a lag of the dependent variable. Formally, we estimate the following specification:

$$U_{ij}^k(t) = \alpha + \beta_1 U_{ij}^k(t-1) + \beta_{2,3} Y_j(t-m) + \beta_{4,5} IR_{ij}^k(t-m) + \beta_{6,7,8} Robots_j(t-n) + \beta_{9,10,11} AI_j(t-n) + \epsilon_{ij}^k(t) \quad (1.1)$$

Where $U_{ij}^k(t)$ is the unemployment rate of education group i of age class k in country j at time t , Y is the real GDP per capita expressed in dollar purchasing power parity, IR is the inactivity rate, $Robots$ the growth rate of the stock of industrial robots and AI the growth rate of the stock of patents related to artificial intelligence. $m \in 0, 1$ and $n \in 0, 1, 2$ define the lags for Y , $Robots$ and AI variables respectively.

Data on real GDP per capita and unemployment rates are provided by the OECD. Unemployment and inactivity rates are collected from the labor force surveys conducted by the national statistical offices and centralized by the OECD. The different levels of educational attainment correspond to the 2011 International Standard Classification of Education (ISCED) elaborated by UNESCO. The OECD (2017) uses the ISCED levels 0 to 2 for the category “below upper secondary education”, 3 to 4 for the category “upper secondary or post-secondary non-tertiary education” and 5 to 8 for the category “Tertiary education”. For each education category, we have four age groups: 25-34, 35-44, 45-54, and 55-64 year-olds.

Data on the operational stock of robots were obtained from the International Federation of Robotics (IFR). The IFR provides macro data starting from 1993 to 2017 for most countries. However, changes in methodology in 2004 - 2005 led to an underestimation of the operational stock for some countries. To avoid any major bias and to ensure methodological consistency in our sample, we chose to begin our analysis in 2005. An important limitation of the IFR database is that for the period covered, data on service robots are not available. Therefore, when we use the term “robots”, we refer to industrial robots defined by the IFR (2016) as “an automatically controlled, reprogrammable, multipurpose manipulator programmable in three or more axes, which can be either fixed in place or mobile for use in industrial automation applications” (IFR (2016) p. 25). Figure 1.1 (a) shows the evolution of the average robot density in our sample. During the 2005-2017 period, the series follows an upward trend.



(a): Average stock of robots

(b): AI patents stock

Figure 1.1: Evolution of the average stock of robots (panel (a)) and AI patents stock (panel (b))

Lastly, we identified patents related to AI using a methodology developed by the World International Patent Office (WIPO (2019)). The approach is based on a combi-

nation of keywords and IPC / CPC codes. We then computed the patent stock by using the perpetual inventory method. We start in 1956, the date of Dartmouth conference organized by John McCarthy, considered as the event that formalized AI as a new field of study and research (Warwick (2013)). Figure 1.1 (b) retraces the evolution of the average AI patent stock, which follows a similar trend to that of robots. Conceptually, we use patents as a proxy for technological capabilities, that is, the ability to select, acquire, generate and apply technologies. What we are trying to estimate is the relationship between the growth of AI-related technological capabilities and the unemployment rate. The faster these capabilities evolve, the more the organizational structures and production processes of firms could be disrupted, which would impact the unemployment rates through quantitative and qualitative human resources adjustments.

Some concerns could be raised regarding the use of patents as a proxy for progress in AI-related technological capabilities. First, this approach does not allow us to capture unpatented innovations, and is therefore sensitive to changes in patenting behavior (Öcalan-Özel and Pénin (2016)). However, the surge in the number of patents since 2013 is consistent with the exponential number of scientific publications in the field (WIPO (2019)), and does not seem to be driven by a massive change in the patenting strategy of the firms. Second, one may ask whether these technologies have had time to spread through the economy. Some surveys indicate that AI technologies are already implemented in an increasing number of firms. Chui and Malhotra (2018) conducted a survey of over 2000 American firms in over 10 industries, and found that 47% of these companies have embedded at least one AI capability into their business processes. Garfinkel (2020) gathered data on over 3000 chief information officers in over 89 countries, and found that 37% of firms had already deployed AI technologies in their operations or were planning to do so in the near future.

1.5 Results

First, we estimate equation (1.1) with a fixed effect model. However, the estimator for dynamic panels may be biased (Nickell (1981)). To overcome this bias, we re-estimate the model with the GMM system estimator (Blundell and Bond (1998)). For each GMM regression, we compute the Hansen test for overidentification restrictions.

The results are presented in Table 1.1 and 1.2. The dependent variable $U_{i=0}^{k=25-64}$ represents the unemployment rate of the 25-64 years-olds, $U_{i=1}^{k=25-64}$ for the unemployment rate of the 25-64 years-olds with below up-per secondary education, $U_{i=2}^{k=25-64}$ of the 25-64 years-olds for unemployed with upper secondary or post-secondary non-tertiary education and $U_{i=3}^{k=25-64}$ of the 25-64 years-olds for unemployed with tertiary education. Since the fixed effect estimator could be biased, we will only comment on the results obtained with GMM. Considering the aggregate unemployment rate ($U_{i=0}^{k=25-64}$) in Table 1.1, we find Okun's law through the negative correlation between real GDP growth per capita and the unemployment rate. Robots seem to have an impact on the unemployment rate, and the second lag of the AI variable does appear to be positively correlated (a 1% increase of the AI patent stock in t-2 is associated with an increase

of 0.027 points of the unemployment rate in t). This temporal difference could be explained by the measurement methods of these two variables: while the “robots” variable reflects the stock of industrial robots used in the production process, the “AI” variable reflects the stock of patents related to artificial intelligence. While the former variable reflects productive capacities, the second is a shortcut for technological progress in the field of artificial intelligence, and there is a lag between the discovery of a new technology and its use in production

Examining the unemployment rates by education level provides some interesting insights. First, we note that the absolute value of Okun's coefficient is inversely correlated to the level of education: the unemployment rate of those with below upper secondary education ($U_{i=1}^{k=25-64}$) is 2.7 time more sensitive to the economic conjuncture than the unemployment rate of those with a tertiary degree ($U_{i=3}^{k=25-64}$). Second, we note that the effect of AI on unemployment is concentrated on $U_{i=1}^{k=25-64}$ and $U_{i=2}^{k=25-64}$, while no significant correlation is found for the most educated part of the labor force $U_{i=3}^{k=25-64}$. Third, we observe that robots are positively correlated with unemployment rates, for each education category. When we compute the net effect over a three-year period, we find that $U_{i=1}^{k=25-64}$ and $U_{i=2}^{k=25-64}$ are more affected than $U_{i=3}^{k=25-64}$, with the strongest effect on $U_{i=2}^{k=25-64}$.

To test the fourth research hypothesis, we estimate equation (1.1) with unemployment rates differentiated by both education level and age group. The results of the GMM estimates are presented in Table 1.3, 1.4 and 1.5; and results for the fixed effects estimates are presented in the appendix. A first observation is that while some effects are shared across age groups, other effects identified in Tables 1.1 and 1.2 are concentrated in only one or two age groups. For example, the negative correlation between $AI(t)$ and $U_{i=1}^{k=25-64}$ observed in Table 1.1 is found in three of the four age groups presented in Table 1.4; and we observe a similar pattern for the positive correlation between $Robots(t)$ and $U_{i=2}^{k=25-64}$ (Table 1.3). In contrast, the positive correlation between $Robots(t)$ and $U_{i=1}^{k=25-64}$ highlighted in Table 1.1 appears to come only from the 25-34 years-olds (Table 1.3). Second, the variation in the magnitude of the coefficients shows that there is a lot of heterogeneity in the effects within the same education category but across different age groups. For example, the correlation coefficient between $Robots(t)$ and $U_{i=2}^{k=45-54}$ is almost twice as high as that between the same variable and $U_{i=2}^{k=25-34}$. This gap is, however, partially filled by the presence of some compensation mechanisms that are illustrated by a negative correlation with the lagged values of the Robots variable, but this effect is present only for the 35-44 and 45-54 age groups.

Finally, we note that the most robust of the effects highlighted by Table 1.1 and 1.2 is the positive correlation between the stock of robots and the unemployment rates of people with a medium level of education ($U_{i=2}^{k=25-64}$). The correlation is indeed positive and statistically significant for every age group. On the other hand, the effect of AI on unemployment is not very robust: while we can observe an effect on $U_{i=2}^{k=25-64}$, it is no longer visible when we disentangle this unemployment rate by age groups. The negative correlation with $U_{i=1}^{k=25-64}$ is also a bit surprising, and would imply that there is a complementarity between low-skilled workers and AI technologies. While this can

make sense for some specific tasks, as labeling some data that will then be used to train a supervised learning algorithm, it is harder to find an explanation for this relation at the macro level.

	$U_{i=0}^{k=25-64}$		$U_{i=1}^{k=25-64}$	
	FE	GMM-SYS	FE	GMM-SYS
$U_i^k(t-1)$	0.216*** (0.064)	0.102 (0.095)	0.177* (0.088)	0.066 (0.078)
$Y(t)$	-0.312*** (0.039)	-0.341*** (0.073)	-0.469*** (0.075)	-0.496** (0.066)
$Y(t-1)$	-0.096** (0.039)	-0.147*** (0.035)	-0.192*** (0.050)	-0.270*** (0.053)
$IR_i^k(t)$	-0.059 (0.117)	-0.072 (0.143)	0.112 (0.091)	0.017 (0.170)
$IR_i^k(t-1)$	-0.186 (0.114)	-0.255* (0.135)	0.104 (0.089)	0.139 (0.126)
$AI(t)$	-0.012 (0.016)	-0.019 (0.017)	-0.062 (0.050)	-0.186*** (0.054)
$AI(t-1)$	-0.007 (0.015)	-0.014 (0.012)	0.028 (0.028)	0.050 (0.039)
$AI(t-2)$	0.024** (0.011)	0.027*** (0.009)	0.020 (0.024)	0.049** (0.023)
$Robots(t)$	0.010 (0.003)	0.042*** (0.009)	0.015** (0.006)	0.035** (0.017)
$Robots(t-1)$	-0.003 (0.003)	-0.007 (0.010)	0.007 (0.010)	-0.000 (0.017)
$Robots(t-2)$	-0.001 (0.006)	-0.009 (0.006)	0.012* (0.006)	0.012 (0.012)
<i>Intercept</i>	0.127 (0.408)	0.185 (0.424)	0.178 (0.725)	1.283 (0.999)
Time dummies	Yes	Yes	Yes	Yes
Observations	319	319	319	319
R^2	0.7228		0.6435	
AR(1)		0.000		0.000
AR(2)		0.690		0.222
Hansen		0.505		0.203

Notes: Robust standard errors in parenthesis, clustered at country-level. * p < 0.10
 ** p < 0.05 *** p < 0.01

Table 1.1: *Panel estimates - $U_{i=0}^{k=25-64}$ and $U_{i=1}^{k=25-64}$*

	$U_{i=2}^{k=25-64}$		$U_{i=3}^{k=25-64}$	
	FE	GMM-SYS	FE	GMM-SYS
$U_i^k(t-1)$	0.136* (0.071)	0.041 (0.090)	0.102 (0.092)	0.145* (0.081)
$Y(t)$	-0.374*** (0.042)	-0.418*** (0.081)	-0.174*** (0.032)	-0.182*** (0.056)
$Y(t-1)$	-0.134*** (0.047)	-0.201*** (0.043)	-0.094*** (0.026)	-0.112*** (0.024)
$IR_i^k(t)$	-0.056 (0.101)	-0.101 (0.140)	0.059 (0.092)	0.145* (0.081)
$IR_i^k(t-1)$	-0.244* (0.131)	-0.322** (0.142)	-0.001 (0.052)	-0.009 (0.069)
$AI(t)$	-0.018 (0.022)	-0.042 (0.029)	0.004 (0.010)	0.015 (0.021)
$AI(t-1)$	-0.020 (0.018)	-0.033 (0.020)	-0.004 (0.011)	-0.014 (0.015)
$AI(t-2)$	0.028* (0.014)	0.028* (0.015)	-0.002 (0.010)	-0.005 (0.015)
$Robots(t)$	0.014*** (0.004)	0.056** (0.011)	0.006 (0.003)	0.026*** (0.009)
$Robots(t-1)$	-0.008** (0.004)	-0.017 (0.012)	0.004 (0.004)	0.004 (0.008)
$Robots(t-2)$	-0.004 (0.008)	-0.015** (0.007)	-0.000 (0.003)	-0.004 (0.006)
<i>Intercept</i>	0.535 (0.477)	0.951 (0.685)	0.111 (0.251)	0.118 (0.419)
Time dummies	Yes	Yes	Yes	Yes
Observations	319	319	319	319
R^2	0.6688		0.5729	
AR(1)		0.000		0.003
AR(2)		0.763		0.841
Hansen		0.247		0.135

Notes: Robust standard errors in parenthesis, clustered at country-level. * $p < 0.10$
** $p < 0.05$ *** $p < 0.01$

Table 1.2: *Panel estimates - $U_{i=2}^{k=25-64}$ and $U_{i=3}^{k=25-64}$*

	$U_{i=1}^{k=25-34}$	$U_{i=1}^{k=35-44}$	$U_{i=1}^{k=45-54}$	$U_{i=1}^{k=55-64}$
$U_i^k(t-1)$	-0.313*** (0.098)	-0.344*** (0.063)	-0.088 (0.142)	-0.231** (0.115)
$Y(t)$	-0.499*** (0.158)	-0.520*** (0.138)	-0.665*** (0.142)	-0.313** (0.115)
$Y(t-1)$	-0.639*** (0.115)	-0.402*** (0.043)	-0.253*** (0.080)	-0.252*** (0.046)
$IR_i^k(t)$	0.018 (0.108)	-0.032 (0.140)	0.131 (0.081)	0.077 (0.099)
$IR_i^k(t-1)$	-0.006 (0.150)	0.029 (0.073)	-0.004 (0.098)	0.083 (0.093)
$AI(t)$	-0.245*** (0.081)	-0.127** (0.059)	-0.073 (0.124)	-0.210** (0.099)
$AI(t-1)$	0.065 (0.090)	-0.001 (0.020)	0.038 (0.034)	0.003 (0.036)
$AI(t-2)$	0.097 (0.066)	0.002 (0.048)	0.052 (0.036)	0.047** (0.023)
$Robots(t)$	0.095*** (0.029)	-0.021 (0.036)	-0.018 (0.031)	0.058 (0.037)
$Robots(t-1)$	0.007 (0.024)	0.034 (0.027)	-0.001 (0.015)	-0.002 (0.012)
$Robots(t-2)$	0.001 (0.017)	0.039 (0.032)	0.007 (0.009)	0.019 (0.013)
<i>Intercept</i>	0.927 (2.495)	1.949* (1.138)	0.860 (1.174)	2.002 (1.284)
Time dummies	Yes	Yes	Yes	Yes
Observations	319	319	319	319
AR(1)	0.000	0.001	0.005	0.000
AR(2)	0.455	0.954	0.601	0.233
Hansen	0.394	0.305	0.119	0.101

Notes: Robust standard errors in parenthesis, clustered at country-level. * p < 0.10
** p < 0.05 *** p < 0.01

Table 1.3: $U_{i=1}$ differentiated by age group - GMM-SYS estimates

	$U_{i=2}^{k=25-34}$	$U_{i=2}^{k=35-44}$	$U_{i=2}^{k=45-54}$	$U_{i=2}^{k=55-64}$
$U_i^k(t-1)$	0.131 (0.127)	-0.059 (0.124)	-0.029 (0.067)	-0.186** (0.091)
$Y(t)$	-0.567*** (0.162)	-0.466*** (0.079)	-0.379*** (0.100)	-0.289** (0.057)
$Y(t-1)$	-0.108 (0.068)	-0.227*** (0.069)	-0.211*** (0.062)	-0.305*** (0.058)
$IR_i^k(t)$	-0.090 (0.422)	-0.099 (0.121)	0.118 (0.108)	-0.081 (0.055)
$IR_i^k(t-1)$	-0.190 (0.142)	0.008 (0.163)	-0.278** (0.128)	-0.083 (0.051)
$AI(t)$	0.001 (0.066)	-0.086** (0.033)	-0.056 (0.042)	0.009 (0.026)
$AI(t-1)$	0.033 (0.069)	-0.028 (0.032)	-0.025 (0.021)	-0.067** (0.026)
$AI(t-2)$	0.032 (0.037)	0.023 (0.024)	0.017 (0.017)	0.000 (0.022)
$Robots(t)$	0.037*** (0.011)	0.059*** (0.021)	0.072*** (0.018)	0.020* (0.010)
$Robots(t-1)$	0.016 (0.017)	-0.034** (0.013)	-0.009 (0.015)	0.012 (0.012)
$Robots(t-2)$	-0.007 (0.012)	-0.017** (0.007)	-0.015* (0.008)	-0.004 (0.010)
<i>Intercept</i>	-0.709 (1.592)	1.827** (0.840)	0.815 (0.782)	0.942 (0.798)
Time dummies	Yes	Yes	Yes	Yes
Observations	319	319	319	319
AR(1)	0.001	0.000	0.000	0.002
AR(2)	0.624	0.120	0.764	0.491
Hansen	0.207	0.484	0.283	0.286

Notes: Robust standard errors in parenthesis, clustered at country-level. * p < 0.10
** p < 0.05 *** p < 0.01

Table 1.4: $U_{i=2}$ differentiated by age group - GMM-SYS estimates

	$U_{i=3}^{k=25-34}$	$U_{i=3}^{k=35-44}$	$U_{i=3}^{k=45-54}$	$U_{i=3}^{k=55-64}$
$U_i^k(t-1)$	-0.094 (0.099)	-0.010 (0.099)	-0.007 (0.070)	-0.283** (0.106)
$Y(t)$	-0.300*** (0.087)	-0.092 (0.057)	-0.105*** (0.057)	-0.047 (0.045)
$Y(t-1)$	-0.112** (0.046)	-0.139** (0.060)	-0.110*** (0.034)	-0.189*** (0.048)
$IR_i^k(t)$	0.545** (0.081)	0.097 (0.233)	-0.399* (0.222)	-0.011 (0.043)
$IR_i^k(t-1)$	-0.057 (0.087)	0.112 (0.103)	-0.015 (0.089)	-0.065 (0.039)
$AI(t)$	0.044 (0.034)	0.011 (0.039)	-0.013 (0.015)	0.028 (0.025)
$AI(t-1)$	-0.009 (0.021)	-0.065 (0.039)	0.066 (0.052)	-0.041** (0.019)
$AI(t-2)$	-0.039** (0.018)	0.000 (0.024)	0.026* (0.014)	0.012 (0.027)
$Robots(t)$	0.007 (0.013)	0.021* (0.010)	0.016 (0.010)	0.037*** (0.013)
$Robots(t-1)$	0.002 (0.008)	0.028* (0.014)	-0.002 (0.006)	0.018 (0.013)
$Robots(t-2)$	0.009 (0.006)	-0.010 (0.009)	-0.002 (0.004)	-0.012 (0.009)
<i>Intercept</i>	0.357 (0.637)	0.485 (0.699)	-0.945 (0.826)	-0.146 (0.492)
Time dummies	Yes	Yes	Yes	Yes
Observations	319	319	319	319
AR(1)	0.002	0.001	0.004	0.001
AR(2)	0.784	0.503	0.233	0.131
Hansen	0.224	0.221	0.268	0.332

Notes: Robust standard errors in parenthesis, clustered at country-level. * p < 0.10
** p < 0.05 *** p < 0.01

Table 1.5: $U_{i=3}$ differentiated by age group - GMM-SYS estimates

Based on the results obtained, we can now answer the four research questions formulated in section 2. First, we can validate H1 (a positive correlation between *Robots*, *AI* and *U*). Indeed, the results presented in Table 1.1 indicate that $Robots(t)$ and $AI(t-2)$ are positively correlated with the aggregate unemployment rate, and the effect of robots is also found in Table 1.1 and 1.2 when we run the regressions on unemployment rates differentiated by education level.

Secondly, we can conclude that H2 (positive correlation between *Robots* and $U_{i=1}^{k=25-64}$ and negative correlation between *Robots* and $U_{i=3}^{k=25-64}$) is not validated. Indeed, we find a positive correlation between robots and the unemployment rate for all three education groups. It should be noted, however, that the magnitude of the correlation is decreasing with the level of education, with a coefficient 1.34 times larger for the less educated in comparison to the most educated. This result is interesting because it goes against the popular opinion which holds that robotics constitutes a “skill-based technical change” that is detrimental to the less educated and beneficial to those with more skills. At the macro level and according to our estimates, it is detrimental for all types of workers, but only slightly less so for those who are skilled.

The third hypothesis (H3), postulating that robots and AI mainly impact the unemployment rate of people with a medium level of education, is only partially validated. Over a three-year period, we note that the net effect of AI is positive (thus increasing unemployment) only for $U_{i=2}^{k=25-64}$, while the effect on $U_{i=1}^{k=25-64}$ is negative and no significant effects are found for $U_{i=3}^{k=25-64}$ (Tables 1.1 and 1.2). We note, however, that this positive effect disappears when we run the regressions by age groups. Regarding robots, the evidence is more solid, with a 1% increase in the robot stock associated with a 0.041 increase of $U_{i=2}^{k=25-64}$ over a three-year period, while the effect is 0.035 and 0.026 for $U_{i=1}^{k=25-64}$ and $U_{i=3}^{k=25-64}$ respectively. When we disentangle $U_{i=2}^{k=25-64}$ by age groups, we notice that the positive effect of robots is present in all four categories.

Finally, the last hypothesis H4 (a positive correlation between robots and $U_{i=1}^{k=25-34}$ (and a negative or not significant effect on the unemployment rate of older people with the same education), and a negative correlation between AI and $U_{i=3}^{k=25-34}$ (and a positive or not significant effect on the unemployment rate of seniors with the same level of education) is only partially validated. While we indeed find that robots are positively correlated with $U_{i=1}^{k=25-34}$ and not with $U_{i=1}^{k=55-64}$ (Table 1.3), we also find that AI is negatively correlated with both $U_{i=3}^{k=25-34}$ and $U_{i=3}^{k=55-64}$ (Table 1.5). This result could imply that the “experience” effect is stronger than the “knowledge obsolescence” effect.

1.6 Concluding remarks and issues for further research

Using panel data on 33 OECD countries, we have investigated the link between AI, robots and unemployment. We have found that both robots and AI tend to increase unemployment (Table 1.1), providing additional evidence to the literature on technological unemployment. We reran the regression on unemployment rates differentiated by education level and age group, and found a strong heterogeneity in the effects of AI and robots (Table 1.1 to 1.5). Some effects highlighted at the aggregate level, such

as the impact of robots on the unemployment rate of those with the lowest levels of education, seem to be concentrated only on some age groups (in this example the 25-34 year-olds). This analysis at the intersection of education and age (as proxies for qualification and experience) bring new elements to the literature by showing that the low/medium/high skilled dichotomy is an overly simplistic framework for studying the effect of these technologies on employment. Finally, we show that the effects of robots are the strongest on the unemployment rate of people with a medium level of education, which supports findings from the literature on technological change-induced labor market polarization. We also find that, over a three-year period, AI increases the unemployment rate of people with a medium level of education (Table 1.2) while the effect is negative or not significant for the others (Table 1.1 and 1.3 respectively), but the result is not very robust as it is no longer observable when we disentangle the unemployment rates by age groups. As discussed in the previous section on data and methodology, the use of patents as a proxy for technological capabilities is questionable. A better way to assess the effect of AI on non-employment would consist in using data on the stock of AI-related software and algorithms. Unfortunately, such data do not exist at the macro level; a first approximation is to use data on the stock of software and databases published by national statistical offices but it is not yet possible to know what part of this stock includes AI-related software. For robots, data on service robots (medical robots, automated vehicles, etc.) would be needed to obtain a more complete estimate of the overall stock of robots. It would also be interesting to assess whether, as in the case of industrial robots, the relationship between service robots and unemployment is null or if some effects can be identified. Lastly, a longer period is needed to observe the effects of these two technologies on unemployment in the long-run.

Through this first chapter, we have highlighted that robotization and progress in the field of AI are associated with an increase in the unemployment rate. We then qualified this result by showing that the effects are heterogeneous according to age and education level, even exerting a decrease in the unemployment rate of certain categories. These empirical results make it possible to demonstrate and quantify the existence of a link between robotization, AI and jobs, but do not allow to highlight the underlying mechanisms explaining this link. The purpose of this second chapter is to provide a theoretical investigation of these mechanisms, with a particular focus on the one leading to technological unemployment and a polarization of the labor market.

Chapter 2

Automation and labor market polarization in an evolutionary model with heterogeneous workers¹

The purpose of this chapter is to investigate the mechanisms underlying the relationship between automation and labor market polarization. To do so, we build an agent-based model (ABM) in which workers, heterogeneous in nature and level of skills, interact endogenously on a decentralized labor market with firms producing goods requiring a specific set of skills to realize the tasks necessary for the production process. The two scenarios considered, with and without automation, confirm that automation is indeed a key factor in polarizing the structure of skill demand and increasing wage inequality. This result emerges even without reverting to the routine-based technical change (RBTC) hypothesis usually found in the literature, giving some support to the complexity-based technical change (CBTC) hypothesis. Finally, we also highlight that the impact of automation on the distribution of skill demand and wage inequality is correlated with the velocity of technical change.

2.1 Introduction

Advances in robotics and artificial intelligence are increasingly raising questions about the impact of technology on employment and the dynamics of the labor market. By pushing the frontier of the automatable set of tasks, these technologies threaten jobs previously spared by previous technological revolutions.

Since the work of Frey and Osborne (2017) concluding that 47% of US jobs were at risk, the literature investigating the link between automation, employment and wages

¹Joint work with André Lorentz.

has grown leading to a larger set of estimates. Using a similar approach but at the task level instead of the job level, Arntz et al. (2016) find that only 9% of jobs are at risk in the US. Vermeulen et al. (2018) used the Bureau of Labor Statistics employment projections and expert assessments to estimate the effect of new technologies on employment and find that, on average, job destruction is offset by job creation.

More recently, Dosi, Piva, Virgillito and Vivarelli (2021) analyze the effectiveness of these compensation mechanisms, i.e., the mechanisms by which jobs destroyed by technology are offset by the creation of new jobs, and find that while extensive investment has a clear positive effect on employment growth, there is some evidence supporting a negative effect of replacement investment on labor demand. Graetz and Michaels (2018) examine the impact of industrial robots on hours worked in several OECD countries, and conclude that there is no statistically significant effect; a result shares by Dauth et al. (2018) for Germany. Acemoglu and Restrepo (2020) find that one additional industrial robot per thousand workers leads to a 0.2 percentage point decrease in the employment rate and a 0.37 percent decrease in wages. A similar study by Chiacchio et al. (2018) for the EU labor market finds a reduction in the employment rate from 0.16% to 0.20%, but no effect on wages.

While the quantitative effects on employment are still being debated, there is a relative consensus on the emergence of labor market polarization, a process through which the wages and share of jobs of medium-skilled workers decline in favor of high-skilled and low-skilled occupations. A growing body of work has highlighted the existence of a process of this polarization in various countries. Autor et al. (2006) show evidence of wage and job polarization between 1980 and 2004 in the United States. This phenomenon also occurs in Europe (Goos et al. (2014), Náplava (2019)), Japan (Furukawa and Toyoda (2018)), the United Kingdom (Salvatori (2018)), and is induced by many factors: immigration, population aging, female labor force entry, offshoring and technical change. Furthermore, Autor and Dorn (2013) conclude that technological progress is the main driver of labor market polarization; a result confirmed by Goos et al. (2014) in 16 western countries. Burstein et al. (2019) study the evolution of income inequality in the US between 1984 and 2003, and find that computerization explains about 60 percent of the increase in the skill premium. Gardberg et al. (2020) find evidence of polarization of the labor market in Sweden from 1996 to 2013, and point out that the impact of computerization could be larger on income inequality than on employability. Finally, Scholl and Hanson (2020) find no link between automation and changes in wages and employment, but a major limitation of their study is the quality of the data used: to measure the level of automation of a job, the authors use the “Degree of automation” variable from O*net whose reliability is questionable. Based on this variable, a cashier job was less automated in 2018 than it was in 2002, with a 25% drop in the “degree of automation” over this period. These data are based on expert estimates and are, as the authors point out, sensitive to the subjectivity of the person performing the rating.

Most of these approaches are grounded in the routine-biased technical change (RBTC) hypothesis, according to which automation is primarily targeting routine tasks. Caines

et al. (2017) offer an alternative explanation by proposing the complex task-biased technical change hypothesis (CBTC), whereby it is the degree of complexity, not the degree of routinization, that is the main explanatory component of the impact of automation on the wage and employment structure. To illustrate this point, the authors took the example of jobs that have both a high routinization index and a high complexity index, and have a very low probability of being automated in the near future, such as financial managers, accountants and auditors, statistical clerks or clinical laboratory technologists and technicians. Using cross-sectional micro-data, the authors regress their routine and complexity indexes on the evolution of wages and find a positive and statistically robust correlation between the complexity index and the evolution of wages between 1980 and 2005, while the routine index is not statistically significant. Concerning changes in the structure of employment, the results are weaker, with the statistical significance of the correlation coefficient between employment and the complexity index changing according to the econometric specifications used. Regarding the routine index, it is only significant at the 5% level in the sub-sample comprising only female workers.

The dominant analytical framework used to explain a polarization induced by technological progress is the task-based approach developed by Autor et al. (2003). In this theoretical model, the production process is disentangled into two types of tasks: routine and non-routine tasks, applying a corresponding decomposition of the labor input. Workers have a heterogeneous productivity endowment and may reallocate to only one of the two tasks, or a combination of the two. The routine tasks and the stock of computer capital are assumed to be perfect substitutes, so the firm's input mix is chosen based on the relative cost of the workers and the price of computers. Over time, the price of computers falls exogenously and puts downward pressure on the wages of workers performing routine tasks, while the demand for non-routine tasks increases leading to a polarization of the labor market.

Acemoglu and Autor (2011) also use the task-based framework to analyze the effect of technical progress on the distribution of skills and wages. Their model includes 3 types of workers: low- medium- and high-skilled workers, and the production is characterized by a continuum of tasks. Each worker can perform any task, but their productivity in a task is subject to their skill. The continuum is divided into 3 segments determined by 3 endogenous thresholds, and each segment characterizes the range of tasks in which each of these 3 types of workers has a comparative advantage. Wages are set according to the marginal productivity of labor.

The model we present in this article is based on a task-based framework, but differs from the modeling choices made by Autor et al. (2003) and Acemoglu and Autor (2011). First, we do not share the assumption that any worker can perform any task, productivity being seen only as an adjustment variable proportional to the gap between the agent's skills and those required by the job. In contrast to this assumption, we model a more segmented labor market, in which an agent cannot apply for a job if he or she does not fully meet the minimum skill levels required. Below these skills requirements, the agent's productivity level would be virtually null and the employer has no interest in hiring him or her. To simplify and keep a tractable model, we use 6 different occupa-

tions in our model (data are presented in the Appendix and the methodology used to aggregate the data is discussed in Section 2.3.1). Based on the assumption previously made, we consider, at time 1, that an engineer, mathematician, or computer scientist could potentially apply for a job in the “other services” or “production occupations” category, given that he or she meets all the skill requirements. A manager could potentially apply for a job in the “sales, office and administrative” category, but not in the “production occupations” category, because his or her technical skills are too low. Note that this is only true at the very beginning of the simulation, as these skill requirements and the skills of each agent evolve endogenously, increasing the heterogeneity of skills within each occupational group and thus the opportunities to change jobs.

Acemoglu and Autor (2011) assume that there is a law of one price, so that even if workers of the same skill level perform different tasks, they receive the same wage at the equilibrium. In the model developed in this paper, we do not assume a law of one price and wages between agents of the same skill level are heterogeneous, depending on the firm and the job held by an agent. We assume that, on average, the wage is determined by the minimum level of skills required for a given job, and thus that a highly qualified worker may receive a low wage if he is underemployed. The existence of labor force mismatching is a well-documented empirical fact (Pellizzari and Fichen (2017) for evidence in developed countries, Brunello and Wruuck (2019) for an extensive literature review on this topic). By introducing a potential mismatch in labor market mechanisms, our model can reproduce these market inefficiencies.

Other theoretical frameworks have been used to study the impact of new technologies on productivity and wages. The EURACE model has been used to show the link between intangible digital investment and productivity (Bertani et al. (2020)), and the link between technological regimes, market concentration and wages inequality Dawid and Hepp (2021). A second family of models uses and expands the original K+S model developed by Dosi et al. (2010): Dosi, Pereira, Roventini, Virgillito et al. (2021) extend the Dosi et al. (2018) model, itself based on Dosi et al. (2010), to study the conditions under which compensation mechanisms works fully. They show that with low wage sensitivity to labor market conditions and a full indexation to productivity gains, the regulation of layoffs and resignations, labor creation and destruction tend to be in balance and the wage structure remains fairly stable. Distinguishing themselves from these two families of models, Fierro et al. (2021) developed their own ABM to show that automation can trigger structural changes and labor market polarization.

Mellacher and Scheuer (2020) build on Dosi et al. (2018) to extend the K+S model to study polarization. They developed a model with three types of workers, engineers, administrative and laborers that correspond to high-skilled, medium-skilled and low-skilled workers respectively. The model reuses the mechanisms of the original K+S model (Dosi et al. (2006)) and successfully generates a polarization of the labor market. However, this model suffers from some limitations. First, the authors assume that technical change is skill-biased against middle-skilled workers (administrators in the model), and thus, given this assumption, it is clear that the dynamics of the simulation generate a polarization of the labor market. Second, there is no heterogeneity within

each category: all engineers have the same skills, all administrators have the same skills, all laborers have the same skills, and they are constant over time. This modeling implies that all administrator jobs are always middle-skill jobs, and none of them can be high-skill or low-skill jobs. Therefore, all employment and wage dynamics occur between groups leading to the absence of within-mechanisms, as tasks content changes, that may explain labor market polarization. Third, there is very limited mobility across job categories: in each period, 0%, 0.5%, 1%, or 1.5% of unemployed agents increase their skill level at the end of the period. This stochastic process implies that there is no path dependency in the agents’ skill trajectory, no learning by doing, and no skill deterioration as in Dosi et al. (2018). The choice to reserve re-skilling for a fraction of unemployed agents implies that the only source of skill improvement is a training system, and thus that work experience is meaningless. De-skilling, i.e. an agent moving from a high-skilled job to a medium-skilled job, is not possible in this model.

Our model differs in all these points. We develop an agent-based model with heterogeneity between and within jobs. We define six categories of jobs derived from US data, that we use to initialize the model (more details are provided in section 2.3.1). We do not establish “highly skilled/moderately skilled/low skilled” categories because skills are multidimensional and the competence of an individual is relative, not absolute. For example, a doctor is generally considered a highly skilled worker while a hairdresser would be considered low skilled. However, if the doctor is trying to give a haircut to one of his patients, it is very likely that he or she is much less efficient than the hairdresser. For practical reasons, we keep an aggregate skill index so that we can calculate a skills polarization index, but this does not interfere with the dynamics of the model: to apply for a job, the applicant’s skill vector must be strictly greater than or equal to the skill vector required for the job. If any of the agent’s skills are below the minimum requirement, he is not be allowed to apply, even if the agent’s overall skill level is above the required skill index for the given position.

Then, we do not assume that technical progress is biased against routine tasks and/or medium-skilled workers. The model contains two kinds of technical progress: one which improves the efficiency of labor through the improvement of machines à la Kaldor (1957). The second type of technical progress aims at automating some tasks, therefore transforming the skills composition of the labor demand. Regarding the first type of technical progress, the firms target the tasks that have the lowest productivity and thus slow down the production process; and for automation, firms use a simple cost-cutting rule by trying to automate the most expensive tasks.

We also introduce a path dependence in the evolution of agents’ skills, as in Dosi et al. (2018). We extend the mechanism by including the task structure in the learning-by-doing equation: the rate of learning depends on the time spent on a given task and, corollary, skill deterioration is inversely proportional to the time spent on a task. For example, if, for the same initial skill level, agent 1 spends 70% of his working time on task A and 30% on task B, the skill required to perform task A increases more rapidly than for an agent spending an equal amount of time on the two tasks. Conversely, the skill required to perform task B improves more slowly. The fraction of working

time allocated to each task is directly affected by technical progress: by increasing the productivity of workers in certain tasks, labor-saving technical progress balances the allocation of time between tasks: workers spend proportionally less time on tasks for which the efficiency of the capital needed to perform them has increased. For labor-displacing technical progress, the effect is even more extreme, as the automation of a set of tasks removes the use of the required skills associated with them. As a result, agents become more specialized meaning at the same time more efficient on the remaining tasks as well as losing, through time, the ability to perform the now automated tasks.

The structure of the model reflects both the between and within-jobs dynamics since the task content of each job evolve endogenously: a set of tasks requiring a skill level of 5 is different, and in a sense more complex, than a set of tasks requiring the same type of skill but with a level equal to 3. These tasks evolve with technical progress and the average skill level of workers in the same job within the same firm (equation 18). Finally, we have full mobility of agents between jobs as long as their skills match those required by the firms. Therefore, as mentioned earlier, our model allows for some mismatch on the labor market.

This model is in the post-Keynesian tradition, embodying four of the five presuppositions at the core of this school of thought (Hein and Lavoie (2019)). First, effective demand drives short- and long-run dynamics, with firms setting their desired level of output (and thus the level of labor demand) based on past demand. Second, the future is uncertain, making any inter-temporal maximization calculation impossible. Third, there is path dependence in both innovation and skills dynamics: difficulties in recruiting engineers and/or poor innovation output have a persistent effect on the firm's productivity; and periods of underemployment or unemployment have a permanent effect on an agent's skill levels. Fourth, distributions matter: a broad polarization of the skills structure increase frictions in the labor market and thus tend to increase unemployment. The fifth presupposition, that money is non-neutral, is not covered in this model given the absence of a banking system and monetary policy.

This model also draws on evolutionary theory, through mechanisms of learning by doing, trial and error in wages adjustment and simple behavioral rules allowing agents to reach an objective (increasing profits for firms and increasing wages for workers). Agents adapt to a Knightian uncertain environment (Knight (1921)), where information is imperfect and an optimization program is impossible to implement (Nelson and Winter (1982)). Unlike equilibrium models (Autor et al. (2003), Acemoglu and Autor (2011)), disequilibrium situations, both in the goods market and in the labor market, are the norm and equilibrium situations are the exception. Finally, the model also used Schumpeterian dynamics: technical progress is endogenous and generates a process of both quantitative and qualitative creative destruction that transforms both the structure of production (by modifying the intensity of each job in the production function) and the structure of tasks within each occupation.

Based on the model described in the next section, we seek to answer two research questions:

- RQ 1: Can we generate a polarization of the labor market without automation?

As the literature review illustrates, many studies have focused on technology to explain labor market polarization. While we do not deny its impact, one may wonder whether this polarization phenomenon is not the “natural” consequence of a labor demand that increasingly values advanced and specialized skills. If this is the case, labor market polarization could well occur in a scenario without automation. However, we believe that technological progress can reinforce this polarization by accelerating the qualitative change in labor demand, and thus by reinforcing the “skill premium” of workers with high-demand skills, while making the skills of other workers redundant and thus increasing the wage gap.

- RQ 2: In the scenario with automation, can we still generate a polarization without using the routine-biased technical change hypothesis?

One common issue with the theoretical models that try to explain labor market polarization is that they assume that automation targets mainly so-called “routine” tasks, which are often assumed to be performed by medium-skilled workers. Thus, the argument tends to be tautological: there is polarization in the labor market because it is assumed that medium-skilled workers are negatively affected by automation. What happens if we remove this assumption? Can we generate a polarization without making any assumptions about the nature of the tasks impacted by automation? For example, can a simple behavioral rule of cost reduction through automation of certain tasks, constrained by the degree of complexity of those tasks, be sufficient to see the emergence of a polarization of the labor market?

2.2 The Model

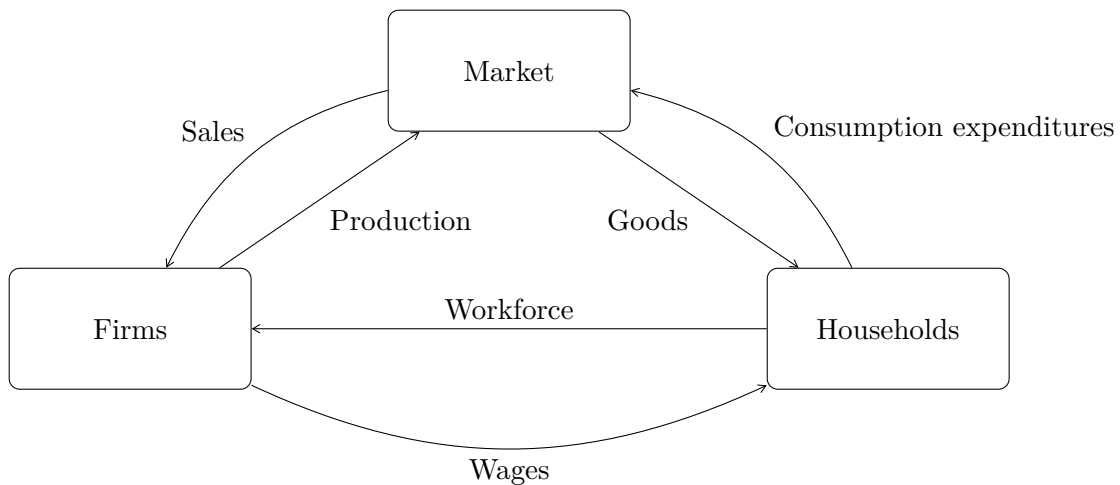


Figure 2.1: The structure of the model

The model is populated by heterogeneous agents differentiated by a vector of skills. An agent can candidate to a job, which is characterized by a vector of required skills,

if and only if they are sufficiently qualified. The model is demand-driven, with firms adapting their level of production to the level of demand realized in the previous period. To produce the desired quantity, firms hire workers in the 6 different types of jobs, and the degree of complementarity among them is given by a Leontief-type production function that represents the firm's productive structure. Wages are heterogeneous across jobs, and evolve according to the firm's productivity gains and the difficulty of recruiting. When an agent has a job, he improves some of his skills according to the time he spends performing tasks related to these skills. On the other hand, skills that are not used in the job deteriorate over time. Unemployed agents suffer a loss of skills proportional to the length of their period of unemployment.

Firms invest in R&D to improve the productivity of capital and, in the scenario with automation, to substitute capital for labor. The structure of wages and skills is directly impacted by the degree of automation: when a task is automated, workers no longer improve their skills previously used to perform this task, and the level of skills required by the firm for new candidates falls, thereby increasing the labor supply available for this job.

2.2.1 Production

The model is composed by a population of F heterogeneous firms indexed $f \in [1; F]$ producing a homogeneous consumption good. The firms use labor to produce this good. The production capacities of the firms are assumed to reflect the evolution of capital, and the productivity it embodies as well as the organizational changes at the firm level. In this respect, capital does not appear explicitly in the production function, but, in line with the idea of a technical progress function (Kaldor (1957)), affects the dynamics of labor productivity. Labor is used in a combination of M jobs indexed $m \in [1, m]$, reflecting the various production phases and activities at the firm level (i.e., production, logistics, marketing, accounting... and so on). Formally the level of production $Y_f(t)$ by each firm f can be described as follows:

$$Y_f(t) = \min \left\{ \frac{L_{1,f}(t)}{A_{1,f}(t)}; \dots; \frac{L_{m,f}(t)}{A_{m,f}(t)}; \dots; \frac{L_{M,f}(t)}{A_{M,f}(t)} \right\} \quad (2.1)$$

where $L_{m,f}(t)$ represents the labor force affected to the job m and $A_{m,f}(t)$ measures the labor force required to produce the $Y_f(t)$ units of good. In other words, $\frac{1}{A_{m,f}(t)}$ measures the level of productivity of the job m . The job level productivity is assumed to reflect both the technology used to produce, the degree of automation of the jobs and the skill level of the workers.

A job m requires a combination of tasks to be realized. We assume here that there exists a vector of I tasks indexed $i \in [1; I]$. Each task are complementary and can be realized either by workers or automatized. The units of goods produced by the $L_{m,f}$ workers hired on the job m in the firm f can then be expressed as follow:

$$\frac{L_{m,f}(t)}{A_{m,f}(t)} = \min \left\{ \frac{L_{m,f}(t)}{B_{1,m,f}(t)}; \dots; \frac{L_{m,f}(t)}{B_{i,m,f}(t)}; \dots; \frac{L_{m,f}(t)}{B_{I,m,f}(t)} \right\} \quad (2.2)$$

where $B_{i,m,f}(t)$ measures the amount of the labor force required to perform the task i per unit produced by the job m . In other words, $\frac{1}{B_{i,m,f}(t)}$ measures the level of productivity of task i on the job m .

The level of productivity of each job m can therefore be written as:

$$\frac{1}{A_{m,f}(t)} = \min \left\{ \frac{1}{B_{1,m,f}(t)}; \dots; \frac{1}{B_{i,m,f}(t)}; \dots; \frac{1}{B_{I,m,f}(t)} \right\} \quad (2.3)$$

The labor force required to perform the task i per unit produced on the job m , $B_{i,m,f}(t)$, depends on the individual worker's efficiency of every single workers $j \in [1; L_{m,f}(t)]$ employed on the job m in firm f , unless the task is fully automatized. We assume here, that the task level efficiency of each worker depends on its ability to grasp the efficiency level embodied in the technology developed by the firm through R&D ($a_{i,m,f}(t-1)$). $b_{j,i,m,f}(t) \in [0; 1]$ measures the degree at which the individual worker j benefits from the efficiency embodied in the technology developed by the firm. The more skilled the workers are on a specific task, the more it benefits from the productivity embodied in the technology (see section 2.2.4), the higher their productivity. More formally, $B_{i,m,f}(t)$ can be formalized as follows:

$$B_{i,m,f}(t) = \begin{cases} \varepsilon & \text{if the task is automatized} \\ \left(a_{i,m,f}(t-1) \sum_{j=1}^{L_{m,f}(t)} \frac{b_{j,i,m,f}(t)}{L_{m,f}(t-1)} \right)^{-1} & \text{if the task is performed by workers} \end{cases} \quad (2.4)$$

Where ε is a parameter with a value close to zero. The firms plan their production according to their expected level of demand ($Y_f^D(t)$) and their current expectations based on their past expectations ($Y_f^D(t-1)$) corrected by the demand faced in the past ($D_f(t)$) and accounting for the level of their inventories ($Y_f(t-1) - D_f(t-1)$):

$$Y_f^D(t) = \alpha Y_f^D(t-1) + (1 - \alpha) D_f(t-1) - (Y_f(t-1) - D_f(t-1)) \quad (2.5)$$

In this respect, we assume that firms have no information on the level of the current state of demand when planning their production and building adaptive expectations. The parameter $\alpha \in [0; 1]$ captures the degree of adaptation of the firm: when $\alpha = 1$, the firm is fully myopic and always plan the same quantity as the previous year adjusted by its stock; whereas when $\alpha = 0$, the firm perfectly adjust to the demand faced at the previous period.

2.2.2 Demand, Market Dynamics and Pricing Behavior

The demand addressed to each firm ($D_f(t)$) is a share $z_f(t)$ of the total expenditures of the consumers. Note that, for the sake of simplicity, we assume that firms produce a homogeneous product. The aggregate demand corresponds to the sum of the consumption level desired by the agents. Formally, the demand addressed to each firm, in

nominal terms, can be expressed as:

$$D_f(t) = z_f(t) \left[\sum_{j=1}^J C_j^D(t) + \sum_{\tau=1}^{t-1} \left(\sum_{j=1}^J C_j^D(\tau) - \sum_{f=1}^F p_f(\tau) Q_f(\tau) \right) \right] \quad (2.6)$$

Where the desired level of consumption is determined by the past consumption level and the actual income, with $\psi \in]0; 1[$:

$$C_j^D(t) = \psi * C_j(t-1) + (1 - \psi) * W_j(t) \quad (2.7)$$

We assume here a non-Walrasian market dynamics (Gaffard (2018)) on this homogeneous good market. There is no market clearing, but at each period, the total amount of expenditures is shared among firms according to a replicator dynamics mechanism representing an imperfect competition process (Metcalfe (1994)). Within such a competition framework, market shares ($z_f(t)$) are allocated according to the relative fitness levels of firms in the market: firms with a fitness level ($E_f(t)$) higher than the average ($\bar{E}(t)$) experience a growth in their market share while the firms with a fitness level lower than the average experience a decline in their market share. Formally, the market share dynamics can be expressed as follows:

$$z_f(t) = z_f(t-1) \left[1 + \phi \left(\frac{E_f(t)}{\bar{E}(t)} - 1 \right) \right] \quad (2.8)$$

where the parameter ϕ defines the strength of the competition among firms.

The fitness level, or competitiveness level, of a firm f is a function of their price $p_f(t)$ and of their past level of unsatisfied demand, normalized by their size ($\frac{D_f(t-1)}{p_f(t-1)Q_f(t-1)}$). While the first component reflects the price competitiveness of the firm, the second, in turn, reflects the firm's ability to reply to the demand expressed by consumers. These two factors are weighted by the parameter $\omega \in [0; 1]$, so that as $\omega \rightarrow 1$, the higher the influence of price competitiveness. More formally, the firm's competitiveness level at each time step can be computed as follows:

$$E_f(t) = \frac{1}{\omega p_f(t) + (1 - \omega) \left(\frac{D_f(t-1)}{p_f(t-1)Q_f(t-1)} - 1 \right)} \quad (2.9)$$

Consequently, the average fitness on the market ($\bar{E}(t)$) is computed as the weighted average fitness of the firms accounting for their market shares:

$$\bar{E}(t) = \sum_{f=1}^F z_f(t-1) E_f(t) \quad (2.10)$$

Firms set their price ($p_f(t)$) applying a mark-up ($\mu_f(t)$) to their unit production costs. The unit production cost is computed using an estimate of the current wage bill per unit produced as follows:

$$p_f(t) = (1 + \bar{\mu}) \sum_{m=1}^M w_{m,f}(t) A_{m,f}(t) \quad (2.11)$$

Where $\bar{\mu}$ is a fixed mark-up. The number of units sold by a firm ($Q_f(t)$) resulting from the market interactions is therefore the minimum between the units of goods it can supply ($Y_f(t)$) and the units of goods demanded by consumers ($\frac{D_f(t)}{p_f(t)}$):

$$Q_f(t) = \min \left\{ Y_f(t); \frac{D_f(t)}{p_f(t)} \right\} \quad (2.12)$$

The resulting profit levels of firms $\pi_f(t)$ can therefore be defined as follows:

$$\pi_f(t) = p_f(t)Q_f(t) - \sum_{m=1}^M w_{m,f}(t)L_{m,f}(t) \quad (2.13)$$

These profits are accumulated by firms to fund their R&D activity aimed at modifying their production capacities and developing machines to automate part of the production tasks.

2.2.3 Employment, wages and the labor market

According to their production plans, firms define their desired level of labor for each type of job m . In order to set the demand for labor ($L_{m,f}^D(t)$), the firms account for their desired level of production ($Y_f^D(t)$), the turnover in the labor force ($T_{m,f}(t-1)$), as well as the workers already employed by the firm for each of the jobs. We assume here that depending on the institutional framework, the labor contracts can be either short-run contracts running for a given time step only, or long-run contracts that can be broken only when either a worker leaves the job or when the firm lays off workers to reduce production. When the labor market only counts short-term contracts, the demand for labor on each job exactly equals the level required to produce the planned production. When the labor market is characterized by long-term contracts, the level of labor demand for each job corresponds to the level required to produce the planned production net from the labor force already in-house corrected from the turnover. A parameter $\lambda \in \{0;1\}$ controls the institutional frame of the labor market, so that, when $\lambda = 1$, all the contracts are short-term contracts, and conversely when $\lambda = 0$, all the contracts are long-term contracts. This parameter allows to take into account the effects of working contracts rigidity, which can have a direct impact on the quality of the matching between labor supply and demand, and so on employment, wages and GDP Dosi et al. (2017). Formally, the level of labor demand per job can be expressed as follows:

$$L_{m,f}^D(t) = \lambda A_{m,f}(t-1)Y_f^D(t) + (1-\lambda) \left(A_{m,f}(t-1)Y_f^D(t) + T_{m,f}(t-1) - L_{m,f}(t-1) \right)$$

This expression can be simplified as follows:

$$L_{m,f}^D(t) = A_{m,f}(t-1)Y_f^D(t) + (1-\lambda)(T_{m,f}(t-1) - L_{m,f}(t-1)) \quad (2.14)$$

$L_{m,f}^D(t)$ is interpreted as the number of slots available in a queue to hire new workers, as in Fagiolo et al. (2004).

For each type of job, a firm proposes a non-negotiable wage $w_{m,f}(t)$. If the firm was unable to attract enough workers to satisfy its need on a specific job, it increases its wage proposal. More formally, wages per job are set as follows:

$$w_{m,f}(t) = w_{m,f}(t-1) \left[1 + \xi_1 \left(\frac{A_f(t-1)}{A_f(t-2)} - 1 \right) + \xi_2 \max \left\{ 0; 1 - \frac{L_{m,f}^S(t-1)}{L_{m,f}^D(t-1)} \right\} \right] \quad (2.15)$$

where $\xi_1 \in [0; 1]$ measures the sensitivity of wages to the firm's productivity growth and ξ_2 reflects the weight of the wage premium resulting from an unbalance between labor supply and demand. The nominal wage is downwardly rigid, consistent with empirical evidence (Jo (2019) and Babecky et al. (2010)).

The labor supply ($L_{m,f}^S(t)$) for a job m to a firm f corresponds to the total count of job-seekers applying on the queue for that given job (see Fagiolo et al. (2004)). Each job is characterized by a vector of tasks that each requires a minimum level of skill to be performed. For a job m in a firm f , the vector $\bar{S}_{m,f}(t)$ contains the minimum skill level $\bar{s}_{i,m,f}(t)$ required to perform each of the necessary i tasks:

$$\bar{S}_{m,f}(t) = \begin{pmatrix} \bar{s}_{1,m,f}(t) \\ \vdots \\ \bar{s}_{i,m,f}(t) \\ \vdots \\ \bar{s}_{I,m,f}(t) \end{pmatrix} \quad (2.16)$$

We assume, here that the firms are conservative in the skill levels required in the job openings, so that they expect a minimum skill level corresponding to the average current skills of their workforce. Hence for a given task i , as long as the task is not automated, the minimum skill level required $\bar{s}_{i,m,f}(t)$ corresponds to the average skill level of their current employees on a job m performing the task i . If the task is automated, there is no minimum skill level required. More formally, this can be expressed as follows:

$$\bar{s}_{i,m,f}(t) = \begin{cases} 0 & \text{if the task is automated} \\ \frac{\sum_{j=1}^{L_{m,f}(t-1)} s_{j,i,m,f}(t-1)}{L_{m,f}(t-1)} & \text{if the task is performed by labor} \end{cases} \quad (2.17)$$

These skill requirements are published by the firms when advertising a job opening and are known by potential candidates.

Each agent $j \in [1; J]$ is characterized by a vector of skills $S_j(t)$ containing his/her skill levels for each possible task i :

$$S_j(t) = \begin{pmatrix} s_{j,1}(t) \\ \vdots \\ s_{j,i}(t) \\ \vdots \\ s_{j,I}(t) \end{pmatrix} \quad (2.18)$$

An agent can apply for a job if and only if his/her skills fit the minimum required for that job. Formally, each element of the vector of skills of a potential candidate have to be greater or equal to the elements of the vector of minimum requirements for that job:

$$s_{j,i}(t) \geq \bar{s}_{i,m,f}(t) \quad \forall i \in [1; I]$$

At each period unemployed agents screen the job openings. Agents that are already employed have a probability ι to look for opportunities and only apply for a new job, if and only if, at least one job opening for which they are qualified enough and that proposes a wage higher than the wage they currently have. In this case, the agent quits his/her current job, increasing the current turnover of the firm where the agent was employed ($T_{m,f}(t)$). Agents that are currently unemployed consider all the job openings for which they match the minimum requirements.

We assume here that each agent can only apply to one job at a time, which is randomly drawn from the set of jobs the agent can apply for. The probability to apply is proportional to the wage offered: the higher the wage offered for a given job in this set, the higher the probability to apply for that job. Formally, this probability corresponds to the ratio of the wage offered for that job over the sum of the wages offered in the set of jobs considered by the agent. Once the agent has picked a job, he joins the queue of applications, adding up to the labor supply for that specific job $L_{m,f}^S(t)$.

Once all the potential candidates have joined a queue, the firms proceed to the selection. The labor market is, as for the final good, a non-Walrasian market: there is no market clearing defining the price in the short-run. The labor market is fully decentralized and relies on direct interactions between the potential candidates and the firms. Each candidate has located him/herself on only one queue. For each of its job openings, the firm screens among the applicants in the queue and proceed to the selection. If the number of applicants $L_{m,f}^S(t)$ is lower than the demand for labor on this given type of job $L_{m,f}^D(t)$, then all the applicants are hired. If the number of applicants $L_{m,f}^S(t)$ exceeds the number of openings for a given job $L_{m,f}^D(t)$, then number of applicants hired is equal to $L_{m,f}^D(t)$. When there are more candidates than open positions, the firm selects according to the following decision rules: first, agents are evaluated according to their skill levels, then they are then ranked accordingly, and finally the firms hire according to the ranking until the openings for the job m are filled.

We assume that the firms only have imperfect information about the skill level of the applicants. The estimated level of a skill for a task i to be performed on a job m in the firm f by an applicant j ($s_{j,i,m,f}^E(t)$) is stochastic and drawn from a Gaussian Law, centered on the actual skill level ($s_{j,i}(t)$) but whose variance is proportional to the distance between the actual skill of the candidate and the minimum skill required by the firm for a given task. This reflects the idea that the further away from the skill set of the firm are the skills from the applicant, the more difficult it is for the firm to estimate them precisely. Formally, this estimate can be described as follows:

$$s_{j,i,m,f}^E(t) \sim N(s_{j,i}(t); (s_{j,i}(t) - \bar{s}_{i,m,f}(t))) \quad (2.19)$$

The applicants are then ranked according to a skill-score ($S_{j,m,f}^E(t)$) computed as an

average of the estimated skill levels ($s_{j,i,m,f}^E(t)$) weighted by the intensity of each tasks ($B_{i,m,f}(t-1)$) within job m :

$$S_{j,m,f}^E(t) = \sum_{i=1}^I s_{j,i,m,f}^E(t) \frac{B_{i,m,f}(t-1)}{\sum_{i=1}^I B_{i,m,f}(t-1)} \quad (2.20)$$

The applicant with the highest score is hired first. All the applicants ranked above the $L_{m,f}^D(t)$ th position are hired by the firm f on a job m , the other remains unemployed. The reverse process occurs when a firm wants to lay off employees: the employer ranks all employees from least to most qualified, and lays off the x^{th} best-ranked agent, where x is the number of employees to lay off.

For each firm f , the labor force $L_{m,f}(t)$ available to perform a given type of job m corresponds to the number of workers that remained in the firm from the previous periods augmented with the workers hired through the process above.

$$L_{m,f}(t) = (1 - \lambda)(L_{m,f}(t-1) - T_{m,f}(t)) + \min \{L_{m,f}^D(t); L_{m,f}^S(t)\} \quad (2.21)$$

The ability of a firm to compete on the labor market both to keep its employees and to hire enough workers to satisfy its needs therefore constrains the production capacities of the firm and its current resources (profits). The latter is required for the firm to increase its competitiveness in the long-run and survive the market selection mechanisms.

2.2.4 Productivity gains, Individual learning and Innovation

In the long-run, firms can experience gains in labor productivity, through two main channels: on the one hand, the labor force gains in efficiency through a learning-by-doing process, as the more a worker realized a task, the more efficient he is; on the other hand, firms can improve the efficiency of their production technology through process innovation. These technical changes result from the R&D activity of the firm and is therefore constrained by the firms' ability to fund its R&D activity. These two mechanisms are completing each other: the higher the skill of the workers on a specific task, the more the firms benefit from the productivity gains resulting from the technical changes resulting from their R&D activity.

We assume that individual skills evolve according to the employment path of the worker through a learning-by-doing process at the individual level. The agents improve their skills on the job: the more they use a specific skill, the more they improve it and, conversely, the less they use a specific skill the more it depreciates. Such mechanisms can be found in evolutionary microeconomic models of production (see among others Llerena et al. (2014)) as well as in evolutionary macroeconomic models of labor market dynamics (see among others Dosi et al. (2018)). These models however only account for the employment history, not the nature of the job. We complete these approaches in that we assume that the whole skill set of each worker evolves according to their employment path, both in the terms of their time on the job and with respect to the nature of the job (i.e., the set of the tasks to be realized and their frequency). Furthermore, the

worker's learning pace is correlated to the distance between its skill and the frontier. Say it differently, it is increasingly difficult to learn when the worker is approaching the maximum level of skill. More formally, at each period, agents suffer a loss of skill unless compensated by the individual learning by performing this task:

$$\frac{\Delta s_{j,i}(t)}{s_{j,i}(t-1)} = \delta_1 \left(\frac{B_{i,m,f}(t-1)}{\sum_{i=1}^I B_{i,m,f}(t-1)} - 1 \right) + \delta_2 \frac{B_{i,m,f}(t-1)}{\sum_{i=1}^I B_{i,m,f}(t-1)} (s_i^{MAX} - s_{j,i}(t-1)) \quad (2.22)$$

where δ_1 and δ_2 are the parameters controlling respectively the amplitude of the depreciation and learning mechanisms. s_i^{MAX} represents the highest possible level of mastery for a given task. Consequently, $s_i^{MAX} - s_{j,i}(t-1)$ captures the idea that the closer the individual skill gets to the maximum level of skill for a given task, the harder it becomes to progress further. The learning curve is therefore not linear and has a concave shape. This process is further shaped by the nature of the worker's activity, so that the more time he spends on a task, the faster he learns and the slower he forgets.

As noted above, the more skilled workers are at a specific task, the more they catalyze the potential productivity embodied in the technology into the actual efficiency of production. $b_{j,i,m,f}(t)$ measures the amplitude of this mechanism so that, first, the higher the individual skill, the higher the amplitude; second, it is only by reaching the maximum skill level s_i^{MAX} , that the firm can only fully benefit from the productivity embodied in its production technology.

$$b_{j,i,m,f}(t) = e^{\kappa(s_{j,i,m,f}(t-1) - s_i^{MAX})} \quad (2.23)$$

In this respect, the worker develops an individual capability to absorb the productivity gains from technical change.

This technical change is rooted in the firm's R&D activity. Each period, the firm devotes part of its resources to R&D expenditures ($R_f(t)$). These investments are constrained by the resources accumulated by the firm in the past. These assumptions are both in line with empirical evidence (Coad and Rao (2010)) and standard in the evolutionary literature (see among others Llerena and Lorentz (2004), Lorentz and Savona (2008), Ciarli et al. (2010), Dosi et al. (2006), Dosi et al. (2018)). These articles highlight the financial constraints imposed to the firms, which can only finance R&D with its past reserves Amendola and Gaffard (1998). More formally, we assume that the firms invest a fixed share η of their sales in R&D:

$$R_f(t) = \min \left\{ \eta p_f(t-1) * Q_f(t-1); \sum_{\tau=1}^{t-1} (\pi_f(\tau) - R_f(\tau)) \right\} \quad (2.24)$$

The firms' R&D expenditures are used to hire engineers. The agents hired as engineer are dedicated to the R&D activity solely and cannot be affected to the production process. This assumption is quite standard in modern evolutionary macroeconomic models (see among others Llerena and Lorentz (2004), Lorentz and Savona (2008), Ciarli et al. (2010), Dosi et al. (2006)). Engineers are hired for the time of the R&D project

and the number of openings $L_{E,f}^D(t)$ is deduced from the actual R&D expenditures $R_f(t)$ and of the wage paid $w_{E,f}(t)$:

$$L_{E,f}^D(t) = \frac{R_f(t)}{w_{E,f}(t)} \quad (2.25)$$

Symmetrically to the wages applied to the production jobs, wages offered to engineers are indexed on the productivity, accounting in the short-run for the firm's ability to attract enough engineers to meet its demand. Formally, wages per job are set as follows:

$$w_{E,f}(t) = w_{E,f}(t-1) \left[1 + \xi_1 \left(\frac{A(t-1)}{A(t-2)} - 1 \right) + \xi_2 \max \left\{ 0; 1 - \frac{L_{E,f}^S(t-1)}{L_{E,f}^D(t-1)} \right\} \right] \quad (2.26)$$

where $\xi_1 \in [0; 1]$ measures the indexation of wages on firm's productivity and ξ_2 reflects the weight of the wage premium.

As for the production jobs, the actual number of engineers hired for the R&D activity $L_{E,f}(t)$ is determined by a job-level matching mechanism:

$$L_{E,f}(t) = \min \left\{ L_{E,f}^D(t); L_{E,f}^S(t) \right\} \quad (2.27)$$

The matching mechanism is symmetric to the one used at the production level described in paragraph 2.2.3.

The R&D activity focuses on the improvement of the efficiency of a specific task in the firm. These improvements are two-sided and aim at both improving the ability of machine/systems potentially automating the task measured by an index $\sigma_{i,f}(t)$, on the one hand, and improving the task level labor productivity $a_{i,m,f}$, on the other. For a given period t , the firms choose to focus their R&D efforts toward the task $\hat{i}(t)$ with the highest index $x_{i,f}(t)$ accounting for the relative labor cost of task i to the total wage bill of the firm:

$$\hat{i}(t) = i \in [1; I] \mid x_{i,f}(t) > x_{i',f}(t) \forall i' \neq \hat{i} \quad (2.28)$$

$$\text{with } x_{i,f}(t) = \frac{\sum_{m=1}^M w_{m,f}(t) B_{i,m,f}(t)}{\sum_{m=1}^M w_{m,f}(t) A_{m,f}(t)}$$

Once the target of the R&D activity is set, the process is assumed to be stochastic. In direct line with the evolutionary models of technical change, we assume here that the probability of success of the R&D activity is a growing function of the number of engineers hired $L_{E,f}(t)$ and their productivity $\frac{1}{A_{E,f}(t)}$:

$$P[\text{Innovation} = 1] = 1 - e^{-\rho \frac{1}{A_{E,f}(t)} L_{E,f}(t)} \quad (2.29)$$

If the R&D process is successful, the output corresponds to both an improvement in the ability of the machines $\sigma_{i,f}(t)$ as well as a modification in the labor intensity of the task $a_{i,m,f}(t)$ (i.e., a gain in the labor productivity embodied in the production

technology) in the various jobs making use of this specific task. In line with the evolutionary literature, we assume that these technical changes result from an incremental improvement through local search Nelson and Winter (1982).

More formally, the improvement in the ability of the machine to be used to automate a task resulting from the R&D activity $\epsilon_{i,f}^\sigma(t)$ is randomly drawn from a Gaussian Law centered on zero and an endogenous variance depending on the distance between the maximum skill level of the task to automate (s_i^{MAX}) and the technological frontier of the firm $\sigma_{i,f}(t-1)$:

$$\begin{aligned} \sigma_{i,f}(t) &= \sigma_{i,f}(t-1)(1 + \max\{0; \epsilon_{i,f}^\sigma(t)\}) \\ \text{with } \epsilon_{i,f}^\sigma(t) &\sim N(0; \beta(s_i^{MAX} - \sigma_{i,f}(t-1))) \end{aligned} \tag{2.30}$$

In doing so, we formally account for the fact that the closer to the frontier the firm is the smaller the possible improvement in the technology.

Symmetrically, the modification in the labor intensity of the task $a_{i,m,f}(t)$ is randomly drawn from a Gaussian Law with centered on zero and an exogenous variance γ_i :

$$\begin{aligned} a_{i,m,f}(t) &= a_{i,m,f}(t-1)(1 + \min\{0; \epsilon_{i,m,f}^a(t)\}) \\ \text{with } \epsilon_{i,m,f}^a(t) &\sim N(0; \gamma_i) \quad \forall m \end{aligned} \tag{2.31}$$

Considering only the modifications reducing labor intensity, we assume that technical change is by essence labor-saving.

These two mechanisms allow us, through the aggregation of individual learning and absorptive capacity to endogenize both the evolution of skills and tasks as well as the structure of task per job. Indeed, Consoli et al. (2019) have emphasized on the importance of the qualitative transformations in job activities to explain the employment dynamics.

2.2.5 Entry/exit mechanism

Given the skill depreciation mechanism, some agents may be pushed out of the labor market because of their inadequate qualifications. An agent who has been unemployed for five consecutive periods leaves the model and is replaced by a new agent. The skills of the new entrants are determined by a random draw among the occupations with labor shortages. Once a job is drawn, the skills of the agent become equal to the required skills to apply to this job plus a 0.5 percent margin. The agent population remains constant over time.

2.2.6 Timeline of the model

1. Based on their production and demand during the previous period, firms set their desired level of production [equation 2.5]
2. Firms set their desired level of R&D by investing a share of their sales made during the previous period. This desired R&D is constraint by their financial reserves [equation 2.24]
3. Based on $Y_f^D(t)$ and $RD_f^D(t)$, firms compute the number of people they want to hire or fire (equation 2.14 for production jobs, and equation 2.25 for engineers).
4. For each job, firms made a wage proposal [equation 2.15].
5. Employed agents have a probability ι to scan the job market looking for a new opportunity, and $\iota = 1$ for the unemployed. An agent scanning the labor market has a positive probability to apply for a job if and only if he is skilled enough, and if the wage proposed by the firm is higher than his previous wage. If the agent was already employed in $t-1$, he has to resign from in job to candidate to another job.
6. If, for a given job, the number of candidates is lower than the number of open positions, all the candidates are hired. If there are more candidates than the number of positions, the firm tries to estimate the skills of each candidate [equation 2.19 and 2.20], rank them and hire only the best ones. If, for a given job, the firm wants to decrease its workforce, it ranks its workers from the less qualified to the most qualified and lays off the first ones.
7. Agents who have not been hired (or have been fired) are now unemployed.
8. Firms compute their cost and set their price. [equation 2.11]
9. Workers set their desired level of consumption based on their consumption during the previous period and their current income. [equation 2.7]
10. The skills levels of each agent are updated. [equation 2.22] Consequently, the agent's productivity on each type of tasks is updated [equation 2.23].
11. The R&D outputs are revealed [equations 2.30 and 2.31].
12. Firms update their minimum skills levels requirement based on the average skills levels of their workers [equation 2.17].
13. Given the evolution of workers' productivity and R&D outputs, firms' productivity is updated. [equations 2.3 and 3.5].
14. Agents who remain unemployed for 5 consecutive periods exit the model and are replaced by new agents with skills that match the labor demand.

2.3 Simulation Results

2.3.1 Initialization and simulation protocol

To initialize the model, we used O*net data to set the skill level of each job. The O*net 23.0 database covers more than 900 occupations and includes for each of them a description of the tasks and skills required. To keep the model tractable, we reduced the number of jobs to 6 by aggregating at the 2 digits SOC-Code. Skills are regrouped in six broad categories presented on the O*net website: basic, complex problem solving, resource management, social, system and technical skills. The skills aggregation within an occupation category is done by averaging the skill level weighted by the employment share of the job within the category. The table containing the data computed by the authors and used for the initialization of the model is available in the appendix. We then assigned a job to the agents, and left 26.5% of them unemployed. We deliberately chose a high initial unemployment rate to avoid a labor shortage right during the first simulation steps of the model. Workers' skills are initialized based on the skills required to perform their job, plus a stochastic component to generate some heterogeneity among workers having the same occupation :

$$s_{j,i} = \bar{s}_{i,m,f}(t)(1 + \max\{0; N(0, k)\}) \quad (2.32)$$

The skills of the unemployed are randomly drawn from a Gaussian distribution :

$$s_{j,i} = \bar{s}_i^u(t)(1 + \max\{0; N(0, k^u)\}) \quad (2.33)$$

We then let the model run for 50 steps to remove noise. During these periods, we allow nominal wages downward adjustments to accelerate the process. We then analyze the model for the next 250 periods. We chose this window period because it is sufficient to show the main model results needed to answer the research questions. The trends observed during these 250 steps continue over the following periods, and adding more periods would not have added more information. The results presented below are the average of 50 simulations, that have been run by using 5 different random seeds. The model contains 2000 agents, 10 firms and 60 different job types (6 per firm). The values for each parameter are presented in Table 3.16 in the Appendix. In this model, profits are used only to fund R&D. The value of the mark-up is calculated as the ratio of the engineers' wage to the total wages, giving a result equal to 8.9%.

2.3.2 Polarization measurement

Labor market polarization is characterized by a polarization of both the wage and the skill structure. The standard way to measure it is to classify each agent in a low/medium/high skill category and to look at the evolution of the wage and employment share of medium-skilled workers compared to the evolution of the share of the low and high skilled workers. We do not use this dichotomy, which could potentially

influence the results via a threshold effect linked to the skill values used for classifying low/medium/high skilled workers, but we do aggregate the skills to have a metric for computing the degree of labor market polarization. To do this, we compute a skill index \hat{s} for each job that takes into account both the average skill level required and the degree of specialization:

$$\hat{s}_{m,f}(t) = \frac{\sum_{i=1}^I \bar{s}_{i,m,f}(t)}{I} \left[1 + \sum_{i=1}^I \left(\frac{\bar{s}_{i,m,f}(t)}{\sum_{i=1}^I \bar{s}_{i,m,f}(t)} \right)^2 \right] \quad (2.34)$$

This index captures the two dimensions of competence: the average level of skill required, and the dimension of expertise through the level of specialization required. We assume that, for two jobs requiring the same average skill level, the more specialized one is considered as the more qualified one. For example, if we assume that there are only 2 different skills in the model, and job A requires a value of 3 for both skills whereas job B requires a value equal to zero for the first skill and 6 for the second, the skill index for job B is higher than the skill index for job A ($6 > 4.5$).

Once we have calculated, within each company and for each job, this skill index, we compute a median relative polarization index following the formula proposed by Wang and Tsui (2000):

$$PI_s(t) = \frac{1}{L_t} \left[\sum_{f=1}^F \sum_{m=1}^M \left[L_{m,f}(t) \frac{(\hat{s}_{m,f}(t) - m(\hat{s}(t)))}{m(\hat{s}(t))} \right]^r \right] \quad (2.35)$$

Where \hat{s} is the required skill index, and $m(\hat{s})$ the median skill index. Wang and Tsui (2000) recommend choosing a value of r between 0 and 1, a value close to zero giving more weight to variations around the median, and a value close to 1 giving the same weight to all distances to the median. We chose a value of 0.5 as this value does not drive the results.

Similarly, we calculate a wages polarization index:

$$PI_w(t) = \frac{1}{L_t} \left[\sum_{f=1}^F \sum_{m=1}^M \left[L_{m,f}(t) \frac{(w_{m,f}(t) - m(\hat{w}(t)))}{m(\hat{w}(t))} \right]^r \right] \quad (2.36)$$

These two indexes allow us to capture both aspects of labor market polarization, and we use them as a measurement tool to estimate the effect of automation on the skills and wages structure.

2.3.3 Empirical validation

In the baseline scenario without automation, the model can reproduce some stylized facts. At the micro-level, some of the properties corresponding to stylized facts have emerged by construction, given the equations chosen to model agents' behavior: as for example the heterogeneity of agents in terms of skills and the pro-cyclical accumulation of skills for workers which are guaranteed by Equation 2.22. Agents may be overqualified due to a partly stochastic selection process on the labor market (part 2.3), and

others agents can be under-qualified if their skills evolve more slowly than those of their colleagues (equation 2.17).

Other stylized facts generated by the model are the results of complex micro-dynamics. At the micro-level, we observe in figure 2.2 a strong heterogeneity in the dynamics of firms' productivity (Bartelsman and Doms (2000) and Dosi (2008)).




Figure 2.2: Evolution of firms' productivity (panel (a)) and Real GDP (panel (b)) - Base 100 in step 0 - 10-periods moving average - Baseline scenario

At the macro level, some stylized facts emerge from the aggregation of the endogenous micro dynamics of the model (Gatti et al. (2011), Dosi and Roventini (2019)). For example, as shown in Figure 2.2, real GDP is growing and fluctuating endogenously. Figure 2.3 shows a positive correlation between the change in unemployment and inequality, consistent with empirical observations (Mocan (1999), Pontusson and Weisstanner (2016), Deyshappriya (2017)).




Figure 2.3: The unemployment-inequality nexus

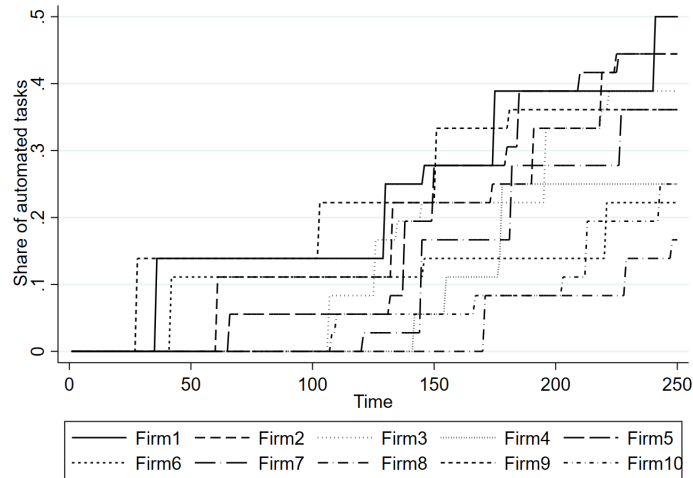


Figure 2.4: Automation spikes

Once we introduce automation, the model successfully generates automation spikes at the firm level (Figure 2.4).

In the next section, we focus on the main stylized fact of interest in this paper: the polarization of the labor market. To do so, we compare three scenarios: one without automation, one with automation and an intermediate case with slow automation. Based on the results, we attempt to answer the two research questions presented earlier in section 2.1.

2.3.4 Baseline scenario

In this baseline scenario, there is no automation. To generate this scenario, we set $\beta = 0$ in equation 2.30. The values of the other parameters of the model are given in the Appendix. All the results presented in this section are from the demand side, i.e., the wages offered by firms and the skills required by employers. The results are presented in figure 2.5, and show that the labor market exhibits a polarization pattern in the wage structure, but not in the skill structure.

With regard to wages, this result is explained by the dynamics of wage-setting (equation 2.15), which depends on two factors: the evolution of the firm's productivity (ξ_1) and the difficulty of recruitment (ξ_2). For the first parameter, we saw in the section on stylized facts that the productivity of firms is heterogeneous and that the dispersion of productivity levels tends to increase over time. This dynamic is driven by two factors: the level of productivity embodied in the capital, and the level of skills of its employees, which measures their ability to fully exploit the capital at their disposal. Regarding the second parameter, given the dynamics of agents' skill evolution (equation 2.22) and skill demand (equation 2.17), heterogeneity in labor supply and labor demand increase across time. As a result, tensions can appear in the labor market, especially for jobs

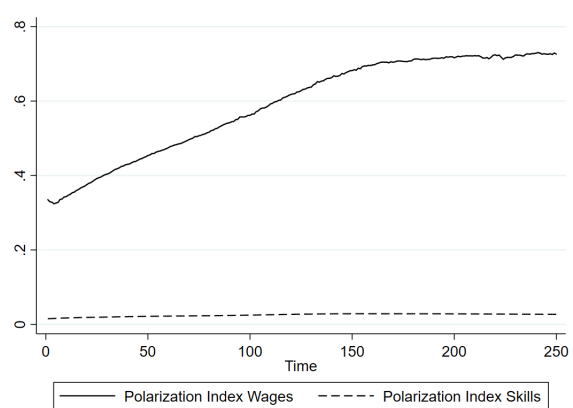


Figure 2.5: Polarization indexes (equations 3.12 and 3.11) - Baseline scenario

requiring high and varied skills, thus accelerating the growth of high wages.




Figure 2.6: Wage distribution Step 1 (panel (a)) vs. Step 250 (panel (b)) - Baseline Scenario

The presence of wage polarization is confirmed by the evolution of the wage distribution between the first and last period of the simulation (Figure 2.6). This polarization is mainly driven by high wages: while in step 1 the wages are between 0.75 and 2.5 times the median wage, at the end of the simulation the distribution exhibits a much wider dispersion, with some wages exceeding 20 times the median wage. Regarding skills, we observe in Figure 2.5 that the polarization index remains flat during the whole simulation. This result is confirmed by the absence of significant change in the distribution of the required skills between the beginning and the end of the simulation (Figure 2.7).

These observations about the distribution of skills and wages tell us something about the degree of polarization in the labor market, but they do not give us any information about the dynamics of the evolution of median values. Figure 2.8 shows the evolution of the median skill index, the real median wage and the average real wage. Over the whole simulation, the real median wage declines while the average real wage remains almost

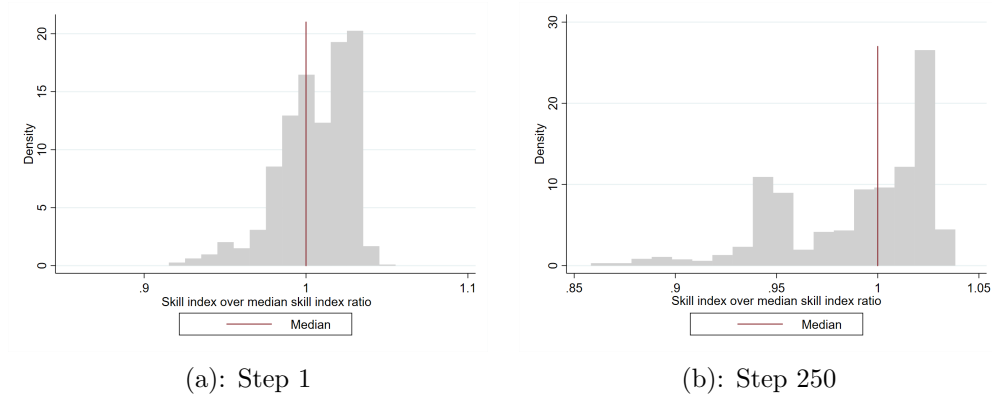


Figure 2.7: Skill index distribution Step 1 (panel (a)) vs. Step 250 (panel (b)) - Baseline Scenario

constant. This decrease in the median and stabilization in the mean is consistent with the changes highlighted in Figure 2.7. The wage distribution also shows a right-skewed distribution by the end of the simulations.

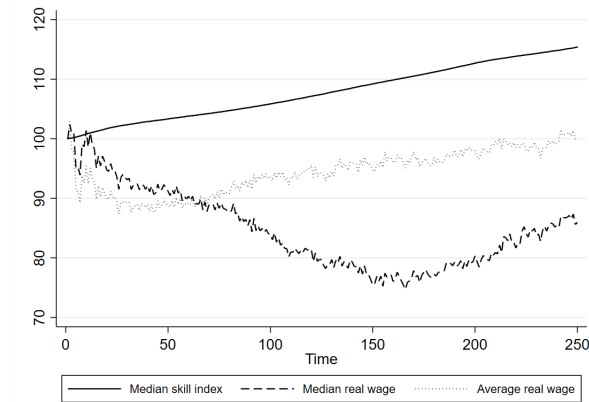


Figure 2.8: Median skill, median and average real wage evolution - Baseline - Base 100 on step 0

The evolution of skills shows that there is a general rise in workers' qualifications. Consequently, the agents on the left of the median (Figure 2.7) are not necessarily less qualified than at step 1, but may have experienced periods of unemployment or underemployment that have slowed down their learning dynamics (equation 2.22). This smooth and continuous rise in median skill level is not a surprise: in this baseline scenario, there is no mechanism to exert downward pressure on skills requirements, so while some agents may face a decline in skills while unemployed, qualifications are, overall, increasing.

This first set of results allows us to answer the first research question (RQ1). With-

out automation, an incomplete polarization of the labor market can be generated, in the sense that only the wages distribution is polarized but not the skills distribution. Finally, the distribution of skills does not seem to be significantly affected, and the polarization observed in many countries is characterized by a polarization of both wages and skills. As a result, this first scenario, without automation, does not seem sufficient to explain the polarization observed by the empirical studies.

2.3.5 Automation scenario

In this second scenario, we introduce automation by setting $\beta = 0.1$ in equation 2.30. Firms can now perform R&D not only to improve the efficiency of existing capital, as in the baseline scenario, but also to automate tasks. Firms follow a simple rule of cost reduction by trying to automate the most costly tasks. Figure 2.9 shows that there is a clear polarization in the labor market, both in terms of wages and skills. However, the polarization of the wage structure does not appear to be more severe than in the baseline scenario, with a value of the wages polarization index at the end of the simulation very close to the one we obtained in the baseline scenario.

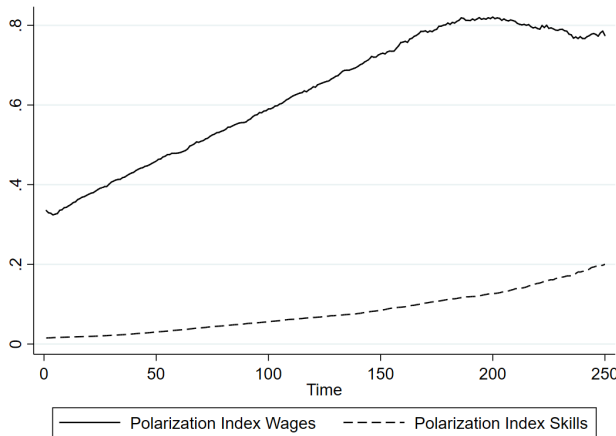


Figure 2.9: Polarization indexes - Automation scenario

A look at the dynamics of the Gini indices, presented in Figure 2.10, computed only on workers to neutralize the effect of unemployment which would mechanically impact the value of the index, underlines that even if automation does not have a significant impact on the shape of the wage distribution, it still has an impact on inequality.

The upward trend in polarization index for skills in Figure 2.9 and the atomized skill distribution presented in Figure 2.11 highlight the effect of automation on the skill structure. The changes in the skills distribution testify of the impact of automation on the qualitative aspect of labor demand. We observe a strong increase in the density of jobs having skills requirements corresponding to less than 80% of the median skill index, and similarly for the right-hand side of the distribution requiring a skill level 1.2 times higher than the median level.

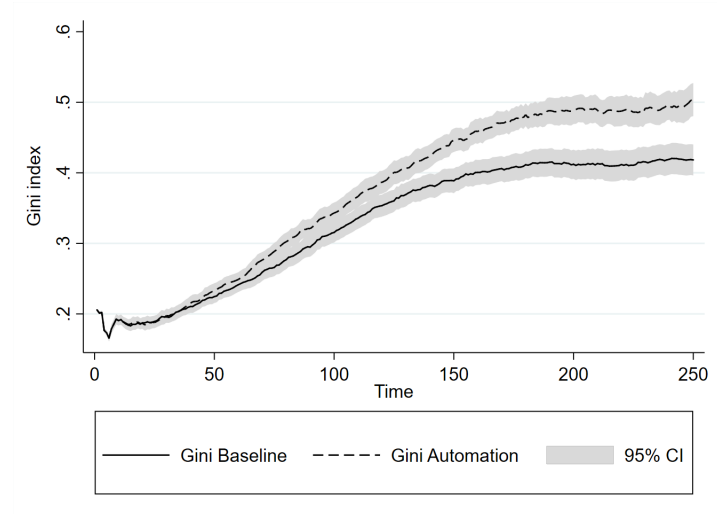



Figure 2.10: Dynamics of the Gini indices



(a): Wage distribution

(b): Skill index distribution

Figure 2.11: Wage distribution (panel (a)) and Skill index distribution (panel (b)) - Automation scenario - Step 250

In contrast to the baseline scenario, the median skill decreases over time (Figure 2.12). It initially follows a similar upward dynamic to that of the scenario without automation until step 50, and then turns around to engage in a continuous downward trend. This inflection point corresponds to the moment when firms begin to automate tasks. As agents improve their skills through a process of learning by doing, they tend to forget how to perform some of the tasks no longer needed in their job. In some extreme cases where almost all tasks are automated in a job, workers lose most of their skills and only retain a kind of “supervisory skill” that allows them to perform only one task of checking if the automated production process is working properly. This result seems to give some credit to the “deskilling hypothesis”, according to which technological change is skill-saving. The literature provides mixed evidence to back this hypothesis: Kunst (2019) use a panel of more than 160 countries shows that automation in the

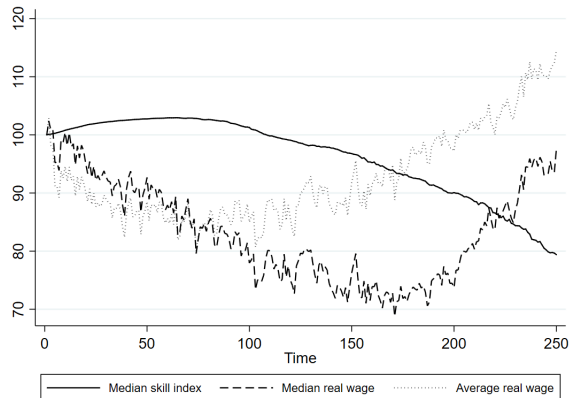


Figure 2.12: Median skill and median real wage evolution - Automation scenario

manufacturing industry has been deskilling since 1950; while the results obtained by McGuinness et al. (2021) using European micro-data tend to show that automation leads to an increase in qualifications, and so is not a de-skilling process.

To confirm that these findings are robust and not simply a matter of randomness, we have carried out a t-test presented in Table 3.1. We find a statistically significant difference between the skills polarization index of the baseline scenario and the one of the automation scenario, but no difference in the wages polarization index. The observation of the Gini coefficients tells us, however, that even if automation does not significantly distort the wage structure, it does have an impact on inequality, with an increase of almost 25% in the Gini index between the two scenarios. Automation has a positive effect on the median real wage, but negatively impacts the median skill level. Interestingly, automation also impacts employment from a quantitative standpoint, with an unemployment rate twice as high as in the scenario without automation. This phenomenon can be explained by the increased dispersion of required skills in the scenario with automation (Figure 2.11) compared to the baseline scenario (Figure 2.7). As a result, the matching between labor supply and demand is more difficult, increasing friction and thus unemployment. Finally and quite logically, aggregate labor productivity is higher in the scenario with automation than in the baseline scenario.

	Automation	Baseline	Difference	P-value
Skills polarization index	0.200	0.027	0.173	5.51e-38
Wages polarization index	0.775	0.726	0.048	0.0657
Gini index	0.504	0.418	0.085	0.0000
Median skill index	3.232	4.698	-1.466	7.89e-46
Real median wage	764.5	673.9	90.66	0.0127
Unemployment rate	0.289	0.153	0.135	9.36e-16
Labor productivity	1172.3	947.1	225.2	2.02e-13

Table 2.1: *Automation and baseline scenarios*

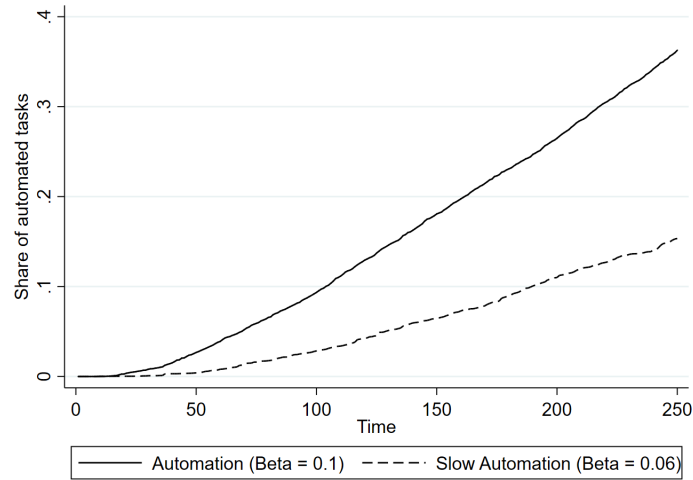


Figure 2.13: Share of automated tasks - Automation and slow automation scenarios

We have seen that automation has an impact on the polarization of the skills structure and wage inequality. However, the question arises whether slower automation, expressed in the model by a lower value of the parameter β , leads to similar conclusions. To answer this question, we consider a third scenario in which the automation of production proceeds more slowly.

2.3.6 Slow automation scenario

In this scenario, we set the value of the β parameter in equation (2.30) to 0.06. The choice of this value is motivated by the rate of automation it generates. Figure 2.13 shows the evolution of the percentage of automated tasks, and a value of 0.06 results in a percentage of automated tasks about half that of the previous scenario, offering an excellent intermediate case. A first comparison with the baseline scenario, presented in Table 2.2, shows that even when automation is relatively slow, the labor market remains polarized but, once again, automation does not seem to increase the degree of polarization of the wage structure. However, we still notice an effect on the Gini index. Slow automation seems to be sufficient to increase the unemployment rate, but not the median wage. Finally, we once again notice a negative effect on the median skill index.

A comparison between the two scenarios with automation is presented in table 2.2. We can observe a relation between the pace of automation and the deformation of the skill structure, with a skills polarization index that more than double between the two scenarios. The median skill index is also negatively impacted by the pace of automation. The wage structure stays stable, as indicated by the non-statistically significance of the difference between the two scenarios for the wage polarization and Gini indexes. On the other hand, the median real wage is positively impacted, which can be explained by the increase in labor productivity, as wages are indexed to the productivity of firms

	Slow	Auto.	Baseline	Difference	P-value	Difference	P-value
	(1)	(2)	(3)	(1) - (2)	(1) - (2)	(1) - (3)	(1) - (3)
Skills polarization index	0.076	0.200	0.027	-0.124	4.18e-26	0.049	2.57e-37
Wages polarization index	0.779	0.775	0.726	0.004	0.869	0.053	0.042
Gini index	0.470	0.504	0.418	-0.034	0.0449	0.052	0.00179
Median skill index	4.129	3.232	4.698	0.897	1.59e-26	-0.569	5.18e-40
Real median wage	646.7	764.5	673.9	-117.8	0.00255	-27.19	0.363
Unemployment rate	0.197	0.289	0.153	-0.092	4.56e-09	0.043	0.0000
Labor productivity	1000.6	1172.3	947.1	-171.7	1.05e-08	53.48	0.0000

Table 2.2: *Slow automation scenario versus automation and baseline scenarios*

(equation 2.15). We observed that an acceleration of automation leads to an increase of the unemployment rate, which is consistent with the fragmentation of the skill structure leading to an increase in the mismatch between labor supply and labor demand.

Finally, we observe in figure 2.14 that the higher the rate of automation, the more specialized the jobs become. This phenomenon can be explained by the dynamics of the model: when a task is automated, the relative working time spend on other tasks tends to increase, leading mechanically to a growing specialization of jobs. The faster the automation, the more specialized and therefore more productive are the workers on the tasks that have not been automated yet. But this increasing specialization is a double-edged sword: given the loss of versatility, it becomes more complicated for a worker to find another job if laid-off. A task automated in a given firm might not be necessarily automated in another firm. This second firm might require candidates to have sufficient skills to perform this task, excluding the laid-off workers that have lost this specific skill due to the automation of this task in his previous job.

The faster the automation, the higher the specialization but the lower the median skill index (Table 2.2). To resolve this apparent contradiction, we need to refer to the formula used to compute the skill index, described in equation 3.13. To observe simultaneously an increase in the specialization index and a decrease in the median skill level, the median value of the average skill level of workers has to decline faster than the skill index; implying that even if agents tend to be more specialized, they are, on average, less skilled than in the scenario without automation. The learning gains resulting from the increase in the relative working time spent on non-automated tasks do not, on average, compensate for the loss of skills resulting from the automation of the other tasks. This result is consistent with the “deskilling hypothesis” mentioned in the precedent section, and implies that the severity of this deskilling process is correlated with the pace of automation.

2.3.7 Summary of the results

Through the three scenarios studies with this model, we have tried to provide answers to the following research questions:

- RQ 1: Can we generate a polarization of the labor market without automation?

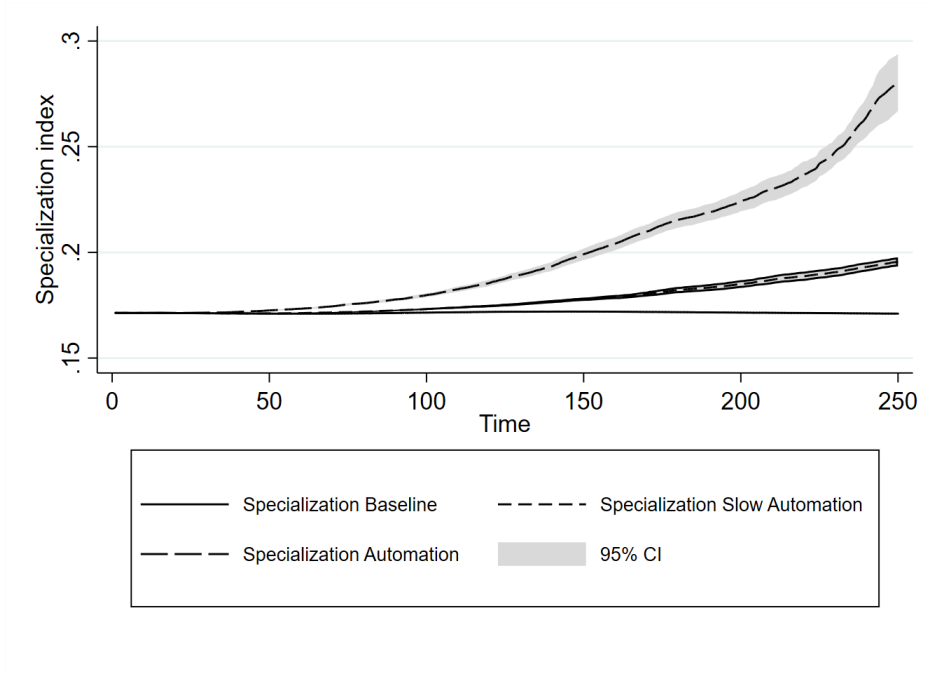


Figure 2.14: Dynamics of specialization indexes

- RQ 2: In the scenario with automation, can we still generate a polarization without using the routine-biased technical change hypothesis?

Regarding the first question, the answer is partly affirmative. Indeed, in the baseline scenario, the wage structure tends to polarize naturally, but the skill distribution remains stable. In the second scenario, the introduction of automation leads to a clear polarization of the skills structure, but does not seem to have a strong impact on the polarization of the wage distribution.

Regarding the second question, the answer is affirmative. In the scenario with automation, we do not make any assumption on the degree of routinization of tasks, but instead companies target their R&D efforts on the most labor-intensive tasks. As these tasks are often the most complex and therefore require a high level of skill, companies struggle to automate them but manage to increase their technological frontier (represented by the variable $\sigma_{i,f}(t-1)$ in equation (2.30)) to automate medium-skilled routine tasks. There are also feedbacks between the labor market conditions and the R&D choices of the firms: if a skill is abundant on the labor market, the cost of labor remains relatively low and therefore firms might not seek to automate the tasks associated with this skill, even if they are easy to automate.

These results contribute to the literature on the impact of automation on the labor market in several ways. First, they provide further evidence of the key role that automation plays in labor market polarization. Without automation, the model fails to generate polarization in the skill structure. Even though the wage structure does not

vary much between the two scenarios, the introduction of automation in the model lead to a rising wages inequality illustrated by the increase of almost 25% of the Gini index. (table 3.1).

These results also indicate that the routine-biased technical change hypothesis is not a theoretical necessity to successfully generate polarization on skill structure and to observe an effect on inequalities. Indeed, a simple cost-reducing rule is, in this model, sufficient. We do not make any assumption about the nature of the tasks impacted by automation, nor follow the standard dichotomy between low-skilled, medium-skilled and high-skilled workers, and our model generates nevertheless a polarization with little restrictive assumptions. Our results do not imply that the routine-biased technical change hypothesis is wrong, but it emphasizes the importance of more classic economic factors, as labor costs, in the explanation of the polarization process. Further theoretical and empirical works should focus more on these factors.

Finally, in this model, automation appears to be a deskilling process: while the degree of specialization increases in the scenario with automation (2.14), the median skill index decreases as the number of automated tasks increases (Figure 2.12). This apparent contradiction is explained by the fact that while automation tends to make jobs more specialized, i.e., intensive in tasks that are difficult to automate or tasks that can be performed by cheap labor, it also tends to make workers less versatile and therefore, on average, less skilled. This result contributes to the debate on the deskilling or skill-enhancing effect of automation, where no clear consensus emerges from the literature (Kunst (2019) and McGuinness et al. (2021)).

2.4 Conclusions

This paper study the effect of automation on labor market polarization. We have shown that without automation, the skill structure does not polarize. The theoretical framework we have developed is less restrictive than those based on the routine-biased technical change hypothesis, as it can generate a polarization in the skill demand structure and an increase in wages inequalities without the need to make ex-ante assumptions about the nature of the tasks and the skill level of the workers targeted by automation. A simple cost reduction rule with a learning and unlearning process related to the time spent on a task seems to be sufficient to generate a polarization in the skills structure and an increase in inequalities. Because the most labor-intensive tasks are also very often the most complex, firms do not manage to automate them but do manage to automate slightly less expensive but much less complex tasks. On the other hand, firms do not necessarily devote their R&D efforts to a task that could be easily automated but that remains relatively cheap to perform due to an abundance of labor with the required skills to execute this task. This result provides some support to the complex-task biased technical change hypothesis developed by Caines et al. (2017).

We have also highlighted a paradoxical effect of automation: by pushing workers to specialize in tasks that are not yet automated for reasons of cost or feasibility, automation makes workers less versatile and, on average, less skilled than in the scenario

without automation. This has the effect of reducing mobility in the labor market and leads to an increase in the unemployment rate.

A severe polarization of the labor market would lead to two problems: increasing difficulties in entering the labor market for people without adequate skills, and deepening wage inequalities. Future investigations should focus on the policy tools, from distributive policies to training policies, that could be used to effectively reduce this polarization and ensure that labor market transformations induced by technology do not lead to a “winners take all” scenario that would significantly increase inequalities.

In this second chapter, we have highlighted the underlying mechanisms that can explain the emergence of technological unemployment and labor market polarization. In particular, we have shown that it is possible to reproduce these two phenomena using a less restrictive hypothesis than the routine biased technical change hypothesis.

The observations we made in chapters 1 and 2 are rather pessimistic: robotics, AI and automation, in general, seem to lead to an increase in unemployment and a widening of inequalities. We can therefore wonder if these effects are irreversible or if, on the contrary, it is possible to counteract them by introducing appropriate policies. This is what chapter 3 proposes to address, extending the model presented in chapter 2 to test four policies: a regulation of layoffs and resignations to stabilize the labor market, the introduction of a minimum wage, unemployment insurance and, finally, a training system.

Chapter 3

Is labor market polarization inevitable? Some answers from a policy simulation exercise.

In this chapter, we extend the agent-based model presented in Chapter 2 to analyze separately the impact of labor market flexibility, a minimum wage, an unemployment insurance and a training system on labor market polarization and some keys macroeconomic variables. Our results highlight that three of these policies help to attenuate the polarization of the labor market, but also involved some trade-offs. For example, a minimum wage leads to lower income inequality and a wage distribution less polarized, but at the cost of a higher unemployment rate. In a final scenario, we test the simultaneous combination of these four policies, and our results indicate positive interactions between them, offsetting the negative effects observed when they are analyzed separately.

3.1 Introduction

Recent technological advances, particularly in the fields of robotics and artificial intelligence, have raised questions about their effects on the labor market. The main concerns relate to technological unemployment, with robots and algorithms capable of performing an increasing range of tasks. The prospective study by Frey and Osborne (2017), which concludes that 47% of jobs in the United States are at high risk of being automated, has fuelled fears of a future technological tsunami that would make human labor obsolete. Attractive as it may be, this narrative is strongly attenuated by the empirical literature on this topic. Bessen (2016), Dauth et al. (2018) and Graetz and Michaels (2018) found no effect of automation on overall employment. Acemoglu and Restrepo (2020), Chiacchio et al. (2018) found that automation decreased aggregate employment (or increase unemployment rate), but the magnitude of the effect is small, as shown in Chapter 1 of this thesis.

Automation has a quantitative impact on jobs, but also a qualitative one. Indeed,

automation modifies the composition of intra-job tasks, by changing the nature of the tasks and the share of working time devoted to each of them. As a result, technical progress generates shocks on the composition of labor demand, and the qualitative adjustment of labor supply may take some time; generating both frictional unemployment and a distortion of the employment distribution.

Indeed, further articles have pointed out that these technological shocks are one of the factors explaining the polarization of the labor market observed in developed countries: Autor et al. (2006) in the US, Salvatori (2018) in the United Kingdom, Furukawa and Toyoda (2018) in Japan and Goos et al. (2014) in Europe. There are some debates on the importance of factors such as offshoring, immigration, population aging, women's entry into the labor market on the degree of labor market polarization, but there is a consensus that technical change is an important variable. Indeed, advances in robotics and artificial intelligence broaden the range of tasks that can be automated, pushing the set of jobs where workers are in direct competition with automation technologies.

Consistent with these findings, we have shown in Chapter 2 that the pace of progress in automation technologies has a direct effect on employment and inequality. As technical progress accelerates, the polarization of the skills distribution tends to increase. In the long-run, the structure of the labor market takes the shape of an hourglass: most workers are highly or low-skilled, and the share of medium-skilled workers becomes marginal. Indeed, the mismatch between workers and skill requirements tends to increase: medium-skilled workers who manage to access higher-quality jobs may be temporarily under-qualified, while those who have failed to access these jobs move to low-skilled occupations for which they are overqualified. This qualitative misadjustment in the labor market generates underemployment and unemployment.

This theoretical result, leading to an overqualification of a certain number of workers, is validated by empirical facts. On the supply side, Sparreboom and Tarvid (2016) have shown that there is an upward trend in overeducated workers in 16 of the 26 European countries. Green et al. (2016) have highlighted that technical progress is robustly involved in the rise of jobs skills requirements in the UK, leading to a growing cost in terms of wages for overeducated workers. On the demand side, Brunello and Wruuck (2019) show, using European Investment Bank surveys, that the share of firms reporting that their investments are constrained by a lack of skills within their organization increased between 2016 and 2018, with more than three-quarters of respondents concerned by this skills shortage.

For workers, being over or underqualified can have consequences in terms of well-being. Based on a survey, Zhu and Chen (2016) found that being overskilled has negative mental impacts. Mavromaras et al. (2013) used Australian micro-data and found that being over-qualified harms job satisfaction. It is interesting to note that the recent literature on the link between labor market mismatch and mental well-being makes a clear distinction between being over-educated and being over-skilled : although there is no consensus on the effect of over-education on well-being, there are clear evidence that being over-skilled has a negative impact on well-being (McGuinness and Sloane (2011) Mavromaras et al. (2013), Zhu and Chen (2016)); a result that can be explained by the

fact that education is only an imperfect proxy for skills. These effects on mental well-being are reinforced by the growing anxiety of economic insecurity, especially for workers performing tasks that could be easily automated. Apart from social and health concerns, the fear of falling down the social ladder, coupled with an increase in inequalities, leads to a reinforcement of the dichotomy between the “winners” and the “losers” of the competition on the labor market. This phenomenon has a concrete political impact that is reflected in the rise of radical right-wing parties. Winkler (2019) analyzed the relationship between inequality and votes in 25 European countries from 2002 to 2014, and found evidence that increasing inequality leads to an increase in support for far-right political parties. The effect is robust even after controlling for migration and job destruction in the manufacturing sector, pointing out that factors like financial liberalization, institutional changes and technical progress could be the main drivers of this pattern. Other studies also highlight the link between the rise in inequality and the vote for radical right-wing parties (Albertini et al. (2017), Oesch and Rennwald (2018), Anelli et al. (2019) and Kurer (2020)).

Based on the literature, we have seen that automation is a key element at the origin of labor market polarization, leading to increasing skill mismatches and inequalities, which in turn can lead to mental health, political, social and economic issues. To address these issues, it is necessary to think about the policies that could attenuate the negative effects of automation on the labor market, reducing inequalities and, ultimately, reducing polarization and unemployment.

A first policy option is to deregulate the labor market to increase its flexibility. According to the proponents of this approach, unemployment is mainly structural and linked to labor market regulations. This point of view was notably defended by Debrun (2003) at the IMF, the OECD (1994) and a whole set of labor economists (Scarpetta (1996) Siebert (1997), Blanchard and Wolfers (2000), Nickell et al. (2005)). The OECD (2004) employment outlook report offers a good summary of the rationale used by the advocates of this position: “Indeed, in deciding whether to hire a worker, employers will take into account the likelihood that firing costs will be incurred in the future. In sum, EPL¹ leads to two opposite effects on labor market dynamics: it reduces inflows into unemployment, while also making it more difficult for job seekers to enter employment (i.e. lower outflows from unemployment)” (OECD (1994) p. 63). While this OECD report concludes that employment protection laws (EPL) tend to reduce the chances of re-employment of the unemployed, it also points out that the overall effect on employment is ambiguous, in contrast to the position originally taken in 1994.

Other studies have challenged the so-called “OECD-IMF consensus” (Howell (2004), Amable et al. (2011), Brancaccio et al. (2018)). Brancaccio et al. (2020) conduct a meta-analysis and find that EPL does not affect employment levels. Among the 53 academic articles identified by the authors to conduct their analysis, only 28% support the “OECD-IMF consensus” while 51% reject it. Regarding the link between EPL and inequality, Brancaccio et al. (2018) find that a decrease in the EPL index between 0.4 and 0.5 points in OECD countries is associated with a 4 to 5 points decrease in wage

¹Employment Protection Laws

shares in the following five years. Ciminelli et al. (2018) share a similar conclusion, finding that deregulation of employment protection may have caused a decline of about 15 percent in the wage share in advanced countries. Pak and Schweltnus (2019) highlight two contradictory effects of EPL on the labor share: while strong employment protection may increase wages, it may also accelerate the substitution of capital for labor by increasing labor costs. On average, the authors find that the second effect outweighs the first, and thus that EPLs tend to have a negative impact on the wage share, but this effect is only significant at the 10% level.

Another axis of labor market flexibility is related to wage setting. One policy is to introduce a minimum wage to act on the left side of the wage distribution to reduce inequality. Boockmann (2010) conducts a meta-analysis of 55 empirical studies and finds that the effects are very heterogeneous across countries, thus it is not possible to draw a general conclusion about the relationship between minimum wage and employment. Another meta-analysis by Nataraj et al. (2014) finds that the minimum wage impacts negatively formal employment while increasing non-formal employment. Using a meta-sample of 77 studies from 18 countries, Chletsos and Giotis (2015) find no effect of minimum wage on employment.

The case of Germany, where a national minimum wage has been implemented in 2015, offers an interesting analytical framework. Caliendo et al. (2019) conducted a literature review and, based on studies conducted between 2016 and 2018, conclude that the minimum wage had a beneficial effect on average hourly wages and a slightly negative effect on employment. Using micro-data, Bossler and Gerner (2020) run a difference-in-difference analysis and find that the introduction of a minimum wage in Germany has increased average wages between 3.8% and 6.3%, while employment loss is estimated to be approximately equal to 1.7%. This last result is challenged by Dustmann et al. (2020), who find that while the minimum wage has a positive impact on wages, it had no significant effects on employment. The empirical literature seems to reach a consensus on the beneficial effect of such a minimum wage on wages, but the effects on employment are still debated.

Theoretical studies have also addressed the question of the link between labor market regulation, employment and inequality. Dosi et al. (2018) extend the K+S model (Dosi et al. (2010)) to test two scenarios: the first one, called the “Fordist regime”, is characterized by, among others characteristics, a full indexation of minimum wage on productivity gains; and the second one, called the “Competitive regime”, characterized by a minimum wage is only partially indexed on labor productivity. The Fordist regime leads to both higher GDP growth and labor productivity, and a lower unemployment rate and Gini index, implying that a higher minimum wage has beneficial macroeconomics effects. Fierro et al. (2021) develop a multi-sector ABM with workers differentiated by three education levels. The authors conduct a policy experiment introducing a minimum wage indexed to the wage of high-skilled workers, and find that it has a positive impact on the pace of automation, while decreasing the wage gap between low- and high-skilled workers.

Another possible policy would be the implementation of an unemployment benefit

system. The objective of this policy would be to reduce the volatility of demand by providing income to job seekers when economic conditions are difficult. The empirical literature shows a positive (negative) correlation between unemployment benefits and the employment (unemployment) rate (Scarpetta (1996), Elmeskov et al. (1998), Nunziata et al. (2002), OECD (2004), Fredriksson and Söderström (2008), Hagedorn et al. (2013)), but the magnitude of the effect tend to be low (Rothstein (2011), Farber and Valletta (2015), Chodorow-Reich et al. (2019), Johnston and Mas (2018)). On the wage side, unemployment benefits should theoretically increase the reservation wage, giving more bargaining power to the unemployed and thus leading to an increase in the average wage. Lancaster and Chesher (1983) conducted surveys of British unemployed and found a positive elasticity between the reservation wage and unemployment benefits. Addison et al. (2010) used European microdata to calculate the elasticity of the reservation wage with respect to the level of unemployment benefits, and find a positive and statistically significant relationship between these two variables. Le Barbanchon et al. (2019) perform a similar analysis on French microdata and find that the effect of potential benefit duration on the reservation wage is almost zero.

The purpose of this chapter is to analyze the tools that policymakers could use to curb labor market polarization. The policies we are going to test fall into 4 distinct categories. The first block is related to labor market reforms, which aim to impose some rules to stabilize the labor market, with a specific focus on the impact of a regulation of the firings and resignations. The second block will be dedicated to the redistributive policy, which aims at correcting the loss of income due to the labor market transformations induce by automation. To do so, we will introduce unemployment benefits funded by social security contributions. The third block will be on pay-setting institutional reforms, with the introduction of a minimum wage into the model. This policy choice is motivated in particular by the findings of Kristal and Cohen (2017), who have shown that computer technologies account for a significant part of the rise of hourly wage inequalities in the private sector in the United States between 1988 and 2012, but they also stressed that institutional parameters are very important, and especially minimum wage that seems to explain a large part of the increase in wage inequality. The relative importance of these factors seems to be country-specific: Albertini et al. (2017) look at the difference between the United States and France, and found that polarization in the United States is mainly due to technology, whereas in France it is mainly related to changes in labor market institutions. Finally, the last block will address the qualitative mismatch issues on the labor market by introducing a training system allowing unemployed agents to upgrade their skills to increase their chance to find a job. In addition to addressing the issue of labor market polarization, these four policies (contract regulation, unemployment benefits, minimum wage and training) cover the three dominant approaches to technological unemployment: the neoclassical approach, according to which technological unemployment is indistinguishable from any other form of unemployment and is the result of labor market rigidities; the Keynesian approach for which technological unemployment is involuntary unemployment due to a lack of aggregate demand, and the structuralist approach according to which techno-

logical unemployment stems from a qualitative mismatch between labor demand and supply.

3.2 The core of the model

The equations, presented in this section, show the main features of the model developed in Chapter 2 that we will then extend to test the policies. The parts concerning the selection mechanism in the consumer goods market and the R&D process will not be developed here because they are not affected by the extensions of the model.

To produce a homogeneous good, a firm f uses a specific combination of labor:

$$Y_f(t) = \min \left\{ \frac{L_{1,f}(t)}{A_{1,f}(t)}; \dots; \frac{L_{m,f}(t)}{A_{m,f}(t)}; \dots; \frac{L_{M,f}(t)}{A_{M,f}(t)} \right\} \quad (3.1)$$

Where $L_{m,f}(t)$ represent the different jobs in firm f at time t , and A the corresponding job intensity. Each job embodied a set of tasks of intensity B :

$$\frac{1}{A_{m,f}(t)} = \min \left\{ \frac{1}{B_{1,m,f}(t)}; \dots; \frac{1}{B_{i,m,f}(t)}; \dots; \frac{1}{B_{I,m,f}(t)} \right\} \quad (3.2)$$

When a task i is automated, B takes a value close to zero (ε). If it is not automated, the intensity of the task changes as a function of the productivity of capital and the skills of workers, which change as a function of the time spent on each task:

$$s_{j,i}(t) = s_{j,i}(t-1) \left[1 - \delta_1 \left(1 - \frac{B_{i,m,f}(t-1)}{\sum_{i=1}^I B_{i,m,f}(t-1)} \right) + \delta_2 \frac{B_{i,m,f}(t-1)}{\sum_{i=1}^I B_{i,m,f}(t-1)} (s_i^{MAX} - s_{j,i}(t-1)) \right] \quad (3.3)$$

Where $s_{j,i}(t)$ is the skill level of agent j to realize a task i , δ_1 and δ_2 are the depreciation and learning parameters respectively, and s_i^{MAX} is a boundary representing the maximum level of expertise attainable. The skill level of workers directly affects their ability to exploit the productive potential of the firm's capital:

$$b_{j,i,m,f}(t) = e^{\kappa(s_{j,i,m,f}(t-1) - s_i^{MAX})} \quad (3.4)$$

Where b is then used to compute the intensity of task i in job m within firm f :

$$B_{i,m,f}(t) = \begin{cases} \varepsilon & \text{if the task is automated} \\ \left(a_{i,m,f}(t-1) \sum_{j=1}^{L_{m,f}(t)} \frac{b_{j,i,m,f}(t)}{L_{m,f}(t-1)} \right)^{-1} & \text{if the task is performed by workers} \end{cases} \quad (3.5)$$

The production dynamic is Keynesian: firms choose a desired level of production based on effective demand, computed based on the quantity produced and the past level of demand:

$$Y_f^D(t) = \alpha Y_f^D(t-1) + (1-\alpha)D_f(t-1) - (Y_f(t-1) - D_f(t-1)) \quad (3.6)$$

Then, for each job m , firms open a number of positions corresponding to their desired production level and the productivity of their employees during the previous period :

$$L_{m,f}^D(t) = A_{m,f}(t-1)Y_f^D(t) + T_{m,f}(t-1) - L_{m,f}(t-1) \quad (3.7)$$

Where $T_{m,f}(t-1)$ represents the turnover of the previous period. Unemployed people scan the job market and position themselves in a queue. Agents who are already employed have a probability ι of leaving their job to apply for an offer with a better wage. Each job is characterized by a set of required skills that represent the minimum level of qualification required, and an agent can only apply if he meets these conditions. The skills required (\bar{s}) for a job are dynamic, and evolve according to the average qualification level of the employees occupying this job:

$$\bar{s}_{i,m,f}(t) = \begin{cases} 0 & \text{if the task is automated} \\ \frac{\sum_{j=1}^{L_{m,f}(t-1)} s_{j,i,m,f}(t-1)}{L_{m,f}(t-1)} & \text{if the task is performed by labor} \end{cases} \quad (3.8)$$

To lower their production costs, firms invest in R&D to both improve the productivity of existing capital and to try to automate tasks. R&D activities are self-funded and are carried out by engineers that the firms have to recruit by using a similar recruitment process to the one for production jobs. As a result, a firm's R&D may be constrained by its ability to attract a sufficient number of engineers; and R&D efficiency is directly linked to the skill levels of the firm's engineers. When a task is automated, the level of skill required to realize this task takes a value close to zero, impacting both the qualitative and quantitative aspects of employment.

The wage proposed by the firms for each job is non-negotiable, and evolves according to the productivity of the firm and the difficulties encountered in the recruitment process during the previous period:

$$w_{m,f}(t) = w_{m,f}(t-1) \left[1 + \xi_1 \left(\frac{A_f(t-1)}{A_f(t-2)} - 1 \right) + \xi_2 \max \left\{ 0; 1 - \frac{L_{m,f}^S(t-1)}{L_{m,f}^D(t-1)} \right\} \right] \quad (3.9)$$

Where A_f is the productivity of the firm, $\xi_1 \in [0;1]$ represents the bargaining power of workers, L^D and L^S are the number of open positions and the number of applicants respectively and ξ_2 a parameter reflecting the degree of responsiveness of wages to a labor shortage.

Once the matching on the labor market is done, firms produce and set their prices by using a fixed markup based on the unitary cost of production. The market selection occurs through a replicator dynamics, which allocate a market share to each firm based

on their prices and their ability to satisfy the demand. Firms compute their profit as follow :

$$\pi_f(t) = p_f(t)Q_f(t) - \sum_{m=1}^M w_{m,f}(t)L_{m,f}(t) \quad (3.10)$$

Workers consume part of their wages, R&D output is revealed and the period ended. If an agent remains unemployed for 5 consecutive periods, he or she is removed from the model and replaced by a new agent whose skills are randomly drawn from the skill levels required for jobs in a labor shortage. There is no exit/entry mechanism for the firms.

The polarization of wages is measured in the same way as in Chapter 2, using a polarization index proposed by Wang and Tsui (2000), which measures the distance to the median:

$$PI_w(t) = \frac{1}{L_t} \left[\sum_{f=1}^F \sum_{m=1}^M \left[L_{m,f}(t) \frac{(w_{m,f}(t) - m(\hat{w}(t)))}{m(\hat{w}(t))} \right]^r \right] \quad (3.11)$$

We proceed similarly to estimate the polarization of the skill structure:

$$PI_s(t) = \frac{1}{L_t} \left[\sum_{f=1}^F \sum_{m=1}^M \left[L_{m,f}(t) \frac{(\hat{s}_{m,f}(t) - m(\hat{s}(t)))}{m(\hat{s}(t))} \right]^r \right] \quad (3.12)$$

Where \hat{s} is a skill index based on the average skill level weighted by a Herfindahl-Hirschmann index that reflects the degree of specialization::

$$\hat{s}_{m,f}(t) = \frac{\sum_{i=1}^I \bar{s}_{i,m,f}(t)}{I} \left[1 + \sum_{i=1}^I \left(\frac{\bar{s}_{i,m,f}(t)}{\sum_{i=1}^I \bar{s}_{i,m,f}(t)} \right)^2 \right] \quad (3.13)$$

3.3 Extensions and policy experiments

In this section, we present the different extensions we made to the original model to test the four policies. Each part is independent and each policy will be tested separately to disentangle their effects. In a final scenario, we will test all policies simultaneously.

3.3.1 Labor market flexibility

The baseline scenario is characterized by a flexible labor market: firms are free to lay off excess labor, and workers have a non-zero probability of resigning to apply for another job. To test the effect of labor market flexibility, we set up an alternative scenario in which layoffs and resignations are regulated: only companies that have made losses in the previous period are allowed to lay off in period t, and employees remain loyal to their company so that only unemployed agents scan the labor market to find a job. To implement this scenario, equation 3.7 is modified as follows:

$$L_{m,f}^D(t) = \begin{cases} A_{m,f}(t)Y_f^D(t) + T_{m,f}(t-1) - L_{m,f}(t-1) & \text{if negative profit in } t-1 \\ \max(0, A_{m,f}(t)Y_f^D(t) + T_{m,f}(t-1) - L_{m,f}(t-1)) & \text{otherwise} \end{cases} \quad (3.14)$$

The parameter ι , which determines the probability that an already employed agent will resign to apply elsewhere, is set to zero. A decrease in labor mobility could reduce the volatility of macro-aggregates by stabilizing the system. On the supply side, given the complementarity between jobs (equation 3.1), a decrease in turnover due to the regulation of resignations will reduce drastically the case where output is constrained by a lack of workers in one or more types of jobs. On the demand side, a ban on layoffs for firms that are making a profit will stabilize consumption and thus firms' desired level of output, which depends on the past level of demand (equation 3.6). Based on these two simultaneous effects, a decrease in labor mobility can be expected to increase real GDP and thus reduce unemployment.

While the stabilization of the labor market should lead to a decrease in the vacancy rate and thus to a reduction in wage inequality due to the decrease in labor shortages (parameter ξ_2 in equation 3.9), the net effect on the distribution of skills is more difficult to predict. As shown by equation 3.3, there is a direct relationship between the employment status of an agent and his skills. On the one hand, greater job stability should decrease the probability for a worker to become unemployed, but on the other hand, the reduction in turnover also makes it more difficult for unemployed agents to find a job, introducing an insider/outsider dichotomy into the model. The decrease in the profit share due to the layoff constraint could also have a negative impact on firms' R&D activities which are self-financed by profits, thus slowing down automation and its impact on labor market polarization. Consequently, we formulate the following research hypothesis:

- RH 1: A decrease in labor mobility (FRR scenario) will increase real GDP, decrease the unemployment rate and attenuate the polarization of the labor market.

3.3.2 Redistributive policy

Another axis focuses on redistributive policies. We introduce into the model an unemployment insurance financed by a contribution on wages.

To implement it, equation (7) is modified by introducing a contribution rate τ_1 :

$$w_{m,f}(t) = (1 - \tau_1)w_{m,f}(t-1) \left[1 + \xi_1 \left(\frac{Af(t-1)}{Af(t-2)} - 1 \right) + \xi_2 \max \left\{ 0; 1 - \frac{L_{m,f}^S(t-1)}{L_{m,f}^D(t-1)} \right\} \right] \quad (3.15)$$

Unemployed agents can now receive an income equal to the level of the unemployment benefit (w^u):

$$R_j(t) = \begin{cases} w_j^u(t) & \text{if } k(t) = 0 \\ w_{m=k,f}(t) & \text{if } k \neq 0 \end{cases} \quad (3.16)$$

Where $R_j(t)$ is the income of agent j at time t , and k the working contract ID of an agent, equal to zero if he is unemployed. Unemployment benefits are calculated based on the agent's last wage:

$$w_j^u(t) = \begin{cases} w_j^u(t-1) * U_j(t) & \text{if } R_j(t-1) = w_j^u(t-1) \\ \tau_2 w_j(t-1) * U_j(t) & \text{if } R_j(t-1) = w_j(t-1) \end{cases} \quad (3.17)$$

Where τ_2 is a parameter between 0 and 1 determining the degree of generosity of unemployment benefits. To maintain a financial balance, an agent is only eligible for unemployment insurance if he or she has contributed for at least three consecutive periods: U_j is a binary variable equal to 1 if agent j has contributed during at least three consecutive periods before being unemployed, and is equal to zero otherwise.

The introduction of unemployment benefits into the model can potentially have two effects: to support demand by limiting the drop in income when an agent becomes unemployed, and to increase the segmentation of the labor market. Indeed, in the initial model, given that an agent's reservation wage is equal to his current income, an agent who had a high income in $t-1$ but who becomes unemployed in period t can potentially apply to all jobs for which he is sufficiently qualified. Although the probability of applying is proportional to the level of the wage proposal, there is a non-zero probability that a highly skilled unemployed agent will apply for a very low-skilled job. With unemployment benefits, the agent can afford to be more selective and will not apply to low-skilled jobs that pay less than the benefits he currently receives. For this scenario, we formulate the following research hypothesis:

- RH 2: An unemployment benefit system (UB scenario) will impact positively the real wages, but will not reduce the polarization of the labor market.

3.3.3 Pay-settings institutions policy

The third block deals with institutional policies, a key component of our analysis. As pointed out by Kristal and Cohen (2016), changes in pay-settings institutions seem to play a major role in the labor market polarization observed in the US. In this section, we will implement a minimum wage into the model.

To do so, we start from equation (6) and we add a variable (\bar{w}) representing the level of the minimum wage:

$$w_{m,f}(t) = \max \left\{ \bar{w}(t); w_{m,f}(t-1) \left[1 + \xi_1 \left(\frac{Af(t-1)}{Af(t-2)} - 1 \right) + \xi_2 \max \left\{ 0; 1 - \frac{L_{m,f}^S(t-1)}{L_{m,f}^D(t-1)} \right\} \right] \right\} \quad (3.18)$$

The minimum wage is partially indexed on labor productivity and can't decrease in nominal value:

$$\bar{w}(t) = \max \left\{ \bar{w}(t-1); \bar{w}(t-1) * \left[1 + \xi_3 \left(\frac{A(t-1)}{A(t-2)} - 1 \right) \right] \right\} \quad (3.19)$$

Where A is the aggregate labor productivity and $\xi_3 \in [0; 1]$ the degree of indexation of the minimum wage on productivity.

By setting a wage floor, the minimum wage will mechanically raise the lowest wages. As a result, the wage gap between the highest and lowest wages will narrow, reducing the polarization of the labor market, leading us to the following research hypothesis:

- RH 3: The introduction of a minimum wage (MW scenario) will increase real GDP, decrease the unemployment rate and the polarization of the wage distribution.

3.3.4 Training policy

Finally, the last block is dedicated to training policy. The idea is to implement a training system enabling unemployed agents to improve their skills to increase their probability of finding a job. Concretely, we introduce the following algorithm:

- Unemployed agents with open unemployment rights ($U_j(t) = 1$) scan the labor market to identify job opportunities that offer a higher salary, but for which they are not sufficiently qualified to apply.
- These offers are then ranked in ascending order according to the skill distance (SD) that we compute by summing only the positive distances:

$$SD_{j,m}(t) = \frac{\sum_{i=1}^I \max(0; \bar{s}_{i,m,f}(t) - s_{j,i,m,f}(t))}{I} \quad (3.20)$$

- The offer ranked number 1 is selected, and the agent then undergoes training to acquire the skills required to apply for the selected position.
- At the end of the training period, the skills of the trained agents are updated to the levels required to apply for the job targeted by the training, plus a bonus of 1% to compensate for the increase in skills required in $t+1$ due to the learning dynamic (equation 3.8). For each type of skill, the post-training level can only increase. Therefore, if for a given skill, the agent has a higher level than the one transmitted by the training, he keeps his pre-training skill level.

We assume that this training system is financed by a 1% contribution rate (ξ_4) on wages. Moreover, we set a maximum average skill distance above which an agent will not be able to follow the training. This threshold δ^{MAX} is set to 0.25 and reflect the fact that an agent cannot be trained in a short period for a job that requires very different skills than those he or she currently possesses (for example, it is impossible to train an economist to become a neurosurgeon in one month).

A first effect of this policy should be to boost labor productivity. Indeed, by upgrading the skills of unemployed agents, the median level of qualification should increase and thus the productivity (equations 3.4 and 2.4). A second effect should be to reduce the polarization of the skill distribution by reducing the gap between low skilled and high skilled workers, and by enabling unemployed middle-skilled workers who have lost

their jobs due to automation to upgrade their skills so they can apply for high-skilled jobs. Finally, this policy should also reduce the share of unemployment related to the qualitative mismatch between labor supply and demand. Based on these potentials effects, we formulate the following research hypothesis:

- RH 4: The training policy (TP scenario) will help to fight against the deskilling effect of automation, reduce unemployment and the degree of polarization of the skills distribution.

3.4 Results

3.4.1 Initialization and baseline scenario

The procedure for initializing the model is similar to the one followed in chapter 2, where we have shown that the introduction of automation in the model had led to an increase in inequality, in the unemployment rate, a polarization of the skill demand, a decrease in the median skill index, an increase of labor productivity and real median wage. The main results found in chapter 2 are sum up in table 3.1:

	Automation	Baseline	Difference	P-value
Skills polarization index	0.200	0.027	0.173	5.51e-38
Wages polarization index	0.775	0.726	0.048	0.0657
Gini index	0.504	0.418	0.085	0.0000
Median skill index	3.232	4.698	-1.466	7.89e-46
Real median wage	764.5	673.9	90.66	0.0127
Unemployment rate	0.289	0.153	0.135	9.36e-16
Labor productivity	1172.3	947.1	225.2	2.02e-13

Table 3.1: *Automation and baseline scenarios*

In this chapter, the baseline scenario now corresponds to the scenario with automation, characterized by a flexible labor market and by a total absence of policy intervention. We will then compare each policy to this scenario to estimate their effectiveness on the reduction of the polarization of the labor market and their impact on some keys macroeconomic variables. For each scenario, we run the model 50 times on 250 periods.

3.4.2 Firing and resignation regulation (FRR)

In this scenario, we test the policy described in Section 3.1. A firm is forbidden to lay off workers if it made a profit in the previous period, and employees remain loyal to their firm by not looking for a better paying job. Unemployed agents are not impacted by this policy, and continue to look for a job until their application is accepted by a firm or are removed from the model and replaced by a new agent (see the paragraph dedicated to the entry/exit mechanism in part 3.2).

Results are presented in Table 3.2. As illustrated by the differences between the two scenarios for the wages polarization index and the Gini index, this policy has a positive effect on reducing wage inequality. This result is explained by the decrease in the dispersion of productivity gains at the firm level, since only firms with losses in the previous period are allowed to lay off workers. Wages offer being indexed on firm's productivity growth (equation 3.9), it mechanically leads to a reduction in wage gaps. Despite the decline in labor productivity, the median real wage is higher than in the baseline scenario. This apparent paradox is explained by lower inflation, as illustrated in Figure 3.1(a). The increase in real wages is therefore not so much due to an increase in nominal wages than to a decrease in inflation, which is explained by a lower vacancy rate than in the baseline scenario (figure 3.1(b)). This increase in the median real wage is associated with an increase in real GDP, which gains more than 18% between the two scenarios, ultimately generating a decrease of nearly 6 points in the unemployment rate.

Finally, the median skill level is higher than in the baseline scenario. This result is explained by the decrease in the unemployment rate, limiting the number of agents suffering a net skill loss due to periods of inactivity (equation 3.3). Finally, the level of polarization of the skill structure remains unchanged from the baseline scenario.

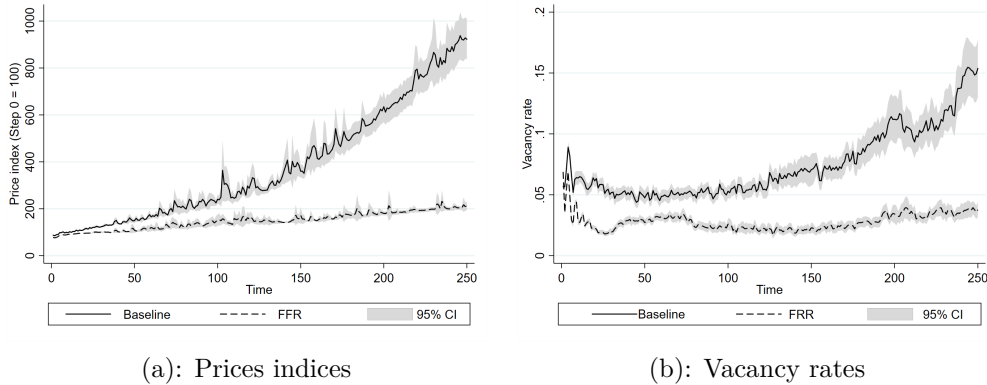


Figure 3.1: Prices indices (panel (a)) and vacancy rates (panel (b)) - Baseline and FRR scenarios

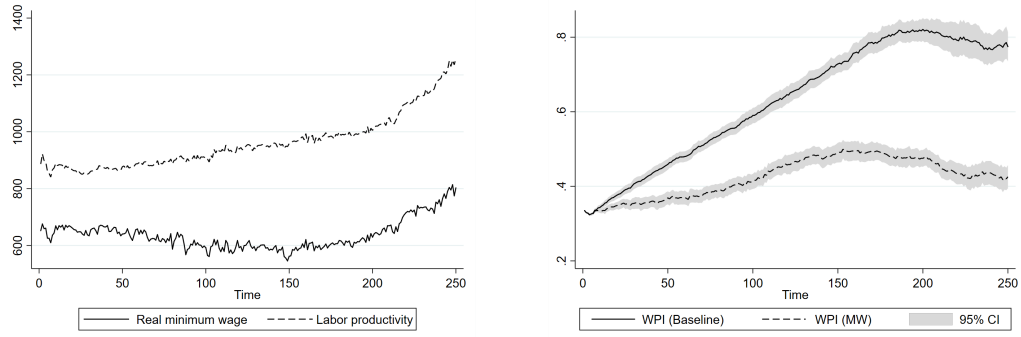
	FRR	Baseline	Difference	P-value
Skills polarization index	0.183	0.200	-0.0175	0.0680
Wages polarization index	0.618	0.775	-0.157	2.30e-10
Gini index	0.297	0.504	-0.207	2.27e-28
Median skill index	3.519	3.232	0.287	0.0000
Real median wage	864.7	764.5	100.2	0.00788
Unemployment rate	0.106	0.289	-0.183	2.71e-24
Labor productivity	1094.8	1172.3	-77.51	0.00599
Share of automated tasks	0.414	0.448	-0.0336	0.0668
Real GDP	1955384.0	1646495.9	308888.1	2.31e-17

Table 3.2: *FRR and baseline scenarios*

3.4.3 Minimum wage (MW)

In this scenario, we introduce a minimum wage partially indexed on labor productivity (equation 3.19). Results are presented in table 3.3. This policy has a beneficial effect on the degree of polarization of the labor market, as illustrated by the sharp decrease in the wages polarization index. As a logical consequence of this decrease in wage dispersion, wage inequality also decreases, with a reduction of 16.7 points of the Gini index. The median real wage is also positively impacted, with an increase of about 15%; but this increase in labor cost comes at the expense of the unemployment rate, which rose by 3.6 points. Some caution is needed in interpreting this last result, since the difference in unemployment rates between the two scenarios is not significant at the 1% level. The percentage of automated tasks is lower in the baseline scenario. This can be explained by the increase in the cost of labor generating two opposite effects: the first is by increasing the cost of labor for the least qualified workers, the minimum wage has the effect of making the automation of tasks performed by these workers more attractive for the firms; and the second is to increase the wages of engineers, thus increasing the costs of R&D and thus slowing down technological progress. In this scenario, the second effect seems to outweigh the first, but we can't draw any clear conclusion as the difference is not significant at the 1% level.

These results seem to corroborate that a minimum wage is an effective policy for reducing inequality, but at the expense of an increase of the unemployment rate. We then test a scenario (LMW) where the indexation of the minimum wage to productivity is lower. To do so, we set ξ_3 to 0.25. The results are presented in Table 3.4. A lower real minimum wage than in the previous scenario still reduces wage polarization and the resulting inequality, as shown by the decrease in the polarization index and the Gini coefficient compared to the baseline scenario; but to a lesser extent. We also notice that the unemployment rate is not affected, implying that the effect on employment is only felt above a certain minimum wage level.



(a): Real minimum wage and labor productivity

(b): WPI dynamics

Figure 3.2: Real minimum wage and labor productivity (panel (a)) Wage polarization indices (WPI) dynamics (panel (b)) - Baseline and MW scenarios

	MW	Baseline	Difference	P-value
Skills polarization index	0.210	0.200	0.009	0.391
Wages polarization index	0.426	0.775	-0.349	1.72e-24
Gini index	0.337	0.504	-0.167	1.27e-16
Median skill index	3.207	3.232	-0.0251	0.732
Real median wage	882.4	764.5	117.8	0.00441
Unemployment rate	0.325	0.289	0.0362	0.0307
Labor productivity	1258.5	1172.3	86.25	0.0183
Share of automated tasks	0.410	0.448	-0.0379	0.0381
Real GDP	1680562.4	1636495.9	34066.5	0.307

Table 3.3: *Automation and baseline scenarios*

	LMW (1)	MW (2)	Baseline (3)	Difference (1) - (2)	P-value (1) - (2)	Difference (1) - (3)	P-value (1) - (3)
Real minimum wage	325.0	802.7	-	-477.7	6.40e-29	-	-
Skills polarization index	0.210	0.210	0.200	0.000	0.951	0.0102	0.346
Wages polarization index	0.701	0.426	0.775	0.276	1.88e-22	-0.0731	0.004
Gini index	0.453	0.337	0.504	0.116	1.59e-10	-0.0507	0.00280
Median skill index	3.191	3.207	3.232	-0.0162	0.820	-0.0413	0.594
Real median wage	811.2	882.4	764.5	-71.20	0.0604	46.62	0.248
Unemployment rate	0.288	0.325	0.289	-0.0375	0.0171	-0.00129	0.939
Labor productivity	1236.1	1258.5	1172.3	-22.38	0.703	63.88	0.279
Share of automated tasks	0.412	0.410	0.448	0.002	0.825	-0.036	0.0564
Real GDP	1729984.6	1680562.4	1646495.9	49422.2	0.342	83488.7	0.118

Table 3.4: *LMW = Low Minimum Wage scenario*

3.4.4 Unemployment benefits (UB)

In this scenario, we introduce unemployment benefits into the model, as detailed in equations 3.16 and 3.17. We test two scenarios: in the first one, the level of compensation represents 50% of the agent's previous wage (UB scenario), while in the second one it represents 80% (HUB scenario). This unemployment benefit system is fully funded by a contribution levied on wages, whose rate has been set so that the system is close to equilibrium at the final step of the simulation; with an average debt to GDP ratio equal to 0.32% and 0.27% for the UB and HUB scenario respectively (figure 3.3).

The results are presented in Table 3.22. A first observation is that unemployment benefits do not seem to have an impact on the polarization of the labor market or inequality. On the other hand, this policy has the effect of significantly increasing the unemployment rate (+ 5.3 points), which is explained by the fact that by introducing a reservation wage equal to the amount of unemployment benefits received, unemployed agents are more selective in their job search and therefore more likely to remain unemployed for a long time. This increase in unemployment should result in a decrease in the median skill index (equation 3.3), but this effect is offset by the slowdown in automation compared to the baseline scenario (-4.1 points in the percentage of automated tasks). Finally, labor productivity increases, but as for the share of automated tasks, the difference with the baseline scenario is not significant at the 1% level.

An increase in unemployment benefits from 50 percent to 80 percent (table 3.6) reinforces the previously observed dynamics: the real median wage increases, but the increase in the unemployment rate is even more severe (+ 12.9 points in comparison with the baseline scenario). The polarization indices are still not impacted by this policy. A policy of implementing unemployment insurance or increasing unemployment benefits does not, therefore, seem to be the most appropriate policy to mitigate labor market polarization in this model.

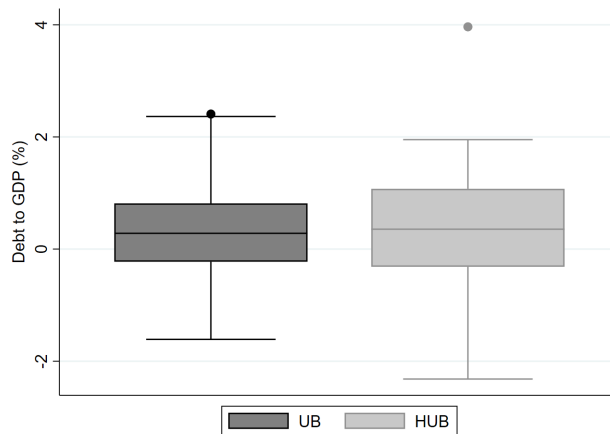


Figure 3.3: Debt ratio - UB and HUB scenarios

	UB	Baseline	Difference	P-value
Skills polarization index	0.198	0.200	-0.00174	0.906
Wages polarization index	0.758	0.775	-0.0165	0.557
Gini index	0.532	0.504	0.0279	0.104
Median skill index	3.233	3.232	0.0008	0.992
Real median wage	776.2	764.5	11.62	0.785
Unemployment rate	0.342	0.289	0.0529	0.00152
Labor productivity	1302.6	1172.3	130.4	0.0128
Share of automated tasks	0.407	0.448	-0.0412	0.0274
Real GDP	1691814.3	1646495.9	45318.3	0.353

Table 3.5: *UB and baseline scenarios*

	HUB (1)	UB (2)	Baseline (3)	Difference (1) - (2)	P-value (1) - (2)	Difference (1) - (3)	P-value (1) - (3)
Skills polarization index	0.195	0.198	0.200	-0.00386	0.772	-0.0056	0.572
Wages polarization index	0.783	0.758	0.775	0.0249	0.343	-0.00843	0.750
Gini index	0.577	0.532	0.504	0.0454	0.00454	0.0733	0.000
Median skill index	3.284	3.233	3.232	0.0507	0.497	0.0516	0.458
Real median wage	876.6	776.2	764.5	91.42	0.0459	103.0	0.0258
Unemployment rate	0.418	0.342	0.289	0.0765	0.0000	0.129	4.39e-13
Labor productivity	1435.3	1302.6	1172.3	132.7	0.0136	263.1	5.96e-08
Share of automated tasks	0.412	0.407	0.448	0.00561	0.517	-0.0356	0.0543
Real GDP	1655434.1	1691814.3	1646495.9	-36380.2	0.458	8938.1	0.801

Table 3.6: *HUB, UB and baseline scenarios*

3.4.5 Training policy (TP)

In this scenario, we introduce a training system that allows the unemployed to update their skills to better match employers' expectations. Details about the implementation into the model are presented in section 3.4. and results are presented in Table 3.7. The training system has a clear impact on skills, both in terms of reducing polarization and increasing skill levels. A better matching between the supply and demand for skills results in an unemployment rate that is almost six points lower than in the baseline scenario, which supports the structuralist theory that technological unemployment is mainly generated by a qualitative mismatch between labor supply and demand. The wage structure and the real median wage are not affected.

The increase in the median skill level slows down the speed of automation. Indeed, a task is only automated when the technological frontier of the firm exceeds the required qualification, which in turn depends on the average qualification level of workers (equation 3.8). By increasing the average level of qualification, the training system also increases the level of skill required which makes their automation more difficult. This slowdown in automation could result in a decline in productivity, but the increase in the average required skill level, which allows agents to better exploit the productive capacities of the capital at their disposal (equation 3.4), offsets this effect and results in a null net impact on productivity.

	TP	Baseline	Difference	P-value
Skills polarization index	0.154	0.200	-0.0459	0.0000
Wages polarization index	0.741	0.775	-0.0331	0.216
Gini index	0.453	0.504	-0.0506	0.00228
Median skill index	3.860	3.232	0.628	3.11e-14
Real median wage	770.1	764.5	5.570	0.896
Unemployment rate	0.229	0.289	-0.0595	0.0000
Labor productivity	1114.7	1172.3	-57.53	0.0659
Share of automated tasks	0.327	0.448	-0.121	1.48e-09
Real GDP	1708612.4	1646495.9	62116.5	0.0564

Table 3.7: *TP and baseline scenarios*

3.5 Results summary

The results of the four policy scenarios are summarized in table 3.23. In this model, the only policy that does not affect the polarization of the labor market is the introduction of unemployment benefits (UB), which nevertheless has a positive effect on wages and productivity at the cost of a decline in the level of employment. The introduction of a regulation on firing and resignations (FRR) or of a minimum wage (MW) has the effect of reducing wage polarization, but has no impact on the skill structure which is only affected by the introduction of a training policy (TP).

	FRR	MW	UB	TP
Skills polarization index	/	/	/	-
Wages polarization index	-	-	/	/
Gini index	-	-	/	-
Median skill index	+	/	/	+
Real median wage	+	+	/	/
Unemployment rate	-	+	+	-
Labor productivity	-	+	+	/
Share of automated tasks	/	-	-	-
Real GDP	+	/	/	/

Table 3.8: *Results summary - "/" = No significant effect (p -value > 0.05)*

Based on these results, we can validate or invalidate the four research hypotheses previously formulated:

- RH 1: A decrease in labor mobility (FRR scenario) will increase real GDP, decrease the unemployment rate and attenuate the polarization of the labor market.
- RH 2: An unemployment benefit system (UB scenario) will impact positively the real wages, but will not reduce the polarization of the labor market.

- RH 3: The introduction of a minimum wage (MW scenario) will increase real GDP, decrease the unemployment rate and the polarization of the wage distribution.
- RH 4: The training policy (TP scenario) will help to fight against the deskilling effect of automation, reduce unemployment and the degree of polarization of the skills distribution.

RH 1 is partially validated, as this policy does not seem to have an impact on the skills polarization index. RH 2 is partially validated, as the median real wage is not impacted. RH3 is partially validated, since the minimum wage contributes to reducing the value of the wages polarization index, but does not seem to impact real GDP. Finally, RH4 is fully validated, and it is the only policy that has an impact on the skills distribution.

It is interesting to note that these policies have complementary effects (FRR, MW and TP to reduce labor market polarization) but also opposite effects: while FRR lowers the unemployment rate, the introduction of a minimum wage has the opposite effects. Similarly, the negative impact of the firing regulation on labor productivity, could be offset by the productivity boost obtained in the MW and UB scenarios. To test these potential complementarities among these different scenarios, we test a final scenario (PC) that combines all policies. We rerun the model 50 times over 250 periods with all four policies activated simultaneously. The results are presented in Table 3.9:

	PC	Baseline	Difference	P-value
Skills polarization index	0.158	0.200	-0.0424	0.0000150
Wages polarization index	0.357	0.775	-0.418	1.47e-32
Gini index	0.235	0.504	-0.268	1.36e-35
Median skill index	3.847	3.232	0.615	1.36e-35
Real median wage	1036.4	764.5	271.8	1.36e-16
Unemployment rate	0.250	0.289	-0.0384	0.00407
Labor productivity	1302.6	1172.3	130.4	0.0000
Share of automated tasks	0.339	0.448	-0.109	3.20e-08
Real GDP	1951565.0	1646495.9	305069.1	1.03e-18

Table 3.9: *PC and baseline scenarios*

We observe that the combination of these four policies allows us to obtain the advantages without the disadvantages. Labor market polarization is attenuated, with a particularly pronounced effect on wages (0.357 vs. 0.775 in the baseline scenario) coupled with a sharp fall in inequality, as shown by the drop in the Gini index (0.235 vs. 0.504). The general level of qualifications increases, which is explained by the training system coupled with a decrease in the share of automated tasks (33.9% vs 44.8%), generating a net positive effect on labor productivity (+11%). The combination of the four policies benefits wages with a 35% increase in the real median wage, but also to jobs with a 3.9 decrease of the unemployment rate.

A look at the dynamics of different variables presented in figure 3.4 provides some interesting insights. In particular, we observe that the combination of these policies has only a limited effect on skills, merely slowing down the polarization of skills without succeeding in reversing the underlying trend. A similar phenomenon can be observed in the evolution of the median skill index, which is certainly higher at the end of the simulation than in the baseline scenario, but which remains lower than the initial level. Conversely, these policies are very effective in increasing wages and reducing inequality, as illustrated by the drop in the wages polarization index, the Gini index and the rise in the real median wage.

3.6 Conclusions

In this chapter, we have analyzed the effectiveness of different public policies to address the social and economic issues raised by the effects of automation on the labor market. We have tested 4 policies: one aiming at stabilizing the labor market via a regulation of the firings and resignations (FRR scenario), a minimum wage indexed on aggregate productivity (MW scenario), an unemployment benefit system financed by a tax on wages (UB) and a training system (TP scenario) for the unemployed, also financed by a contribution on wages. We have seen that each policy has its advantages and disadvantages (Table 3.9), and that none of them can, individually, reduce both the polarization of the wage and the skill distribution.

However, the combination of these four policies offers interesting results: the negative effects of some policies are offset by the other policies, leading to net positive effects on the distribution of skills and wages, the level of employment, real wages, productivity and the level of real output. This encouraging result highlights the importance of having a holistic approach to public policies, that allow for the observation of dynamic interactions that can lead to a more positive outcome that would be missed by an individual analysis of each of these policies.

This work is part of an emerging literature that seeks public policy solutions to the economic and social problems raised by labor market polarization. Our results related to the MW scenario confirm the results obtained by Fierro et al. (2021), who observe that the introduction of a minimum wage reduces the wage gap between high-skilled and low-skilled workers. Our results must be interpreted bearing in mind the analytical framework in which they were obtained: a closed economy, which is therefore not subject to international competition. In the case of an open economy, it is likely that the combination of these four policies, which have the effect of increasing wages and therefore labor costs, could have negative feedback effects on employment linked to an increase of the costs differentials with the rest of the world. This issue is out of the scope of this chapter but could be considered in future work.

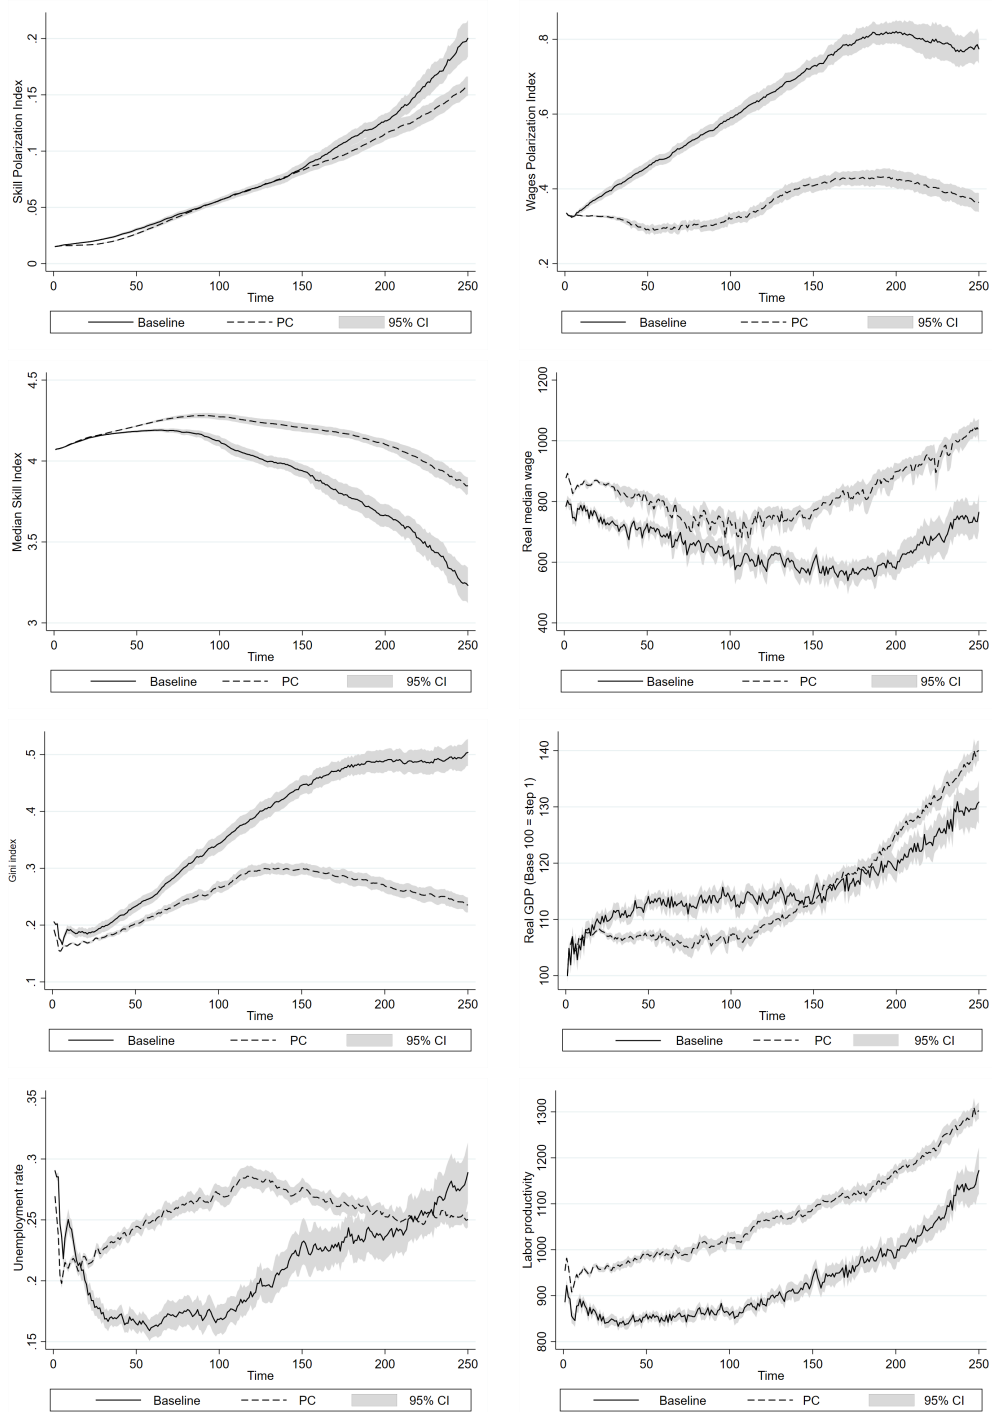


Figure 3.4: Dynamics of polarization indexes and keys macroeconomic variables - PC vs Base-line scenarios

General conclusion

The purpose of this thesis is to analyze the effect of robots, artificial intelligence and, more broadly, automation on employment and inequality. Chapter 1 is dedicated to an empirical analysis that aims to answer a very simple question: is there any statistical evidence that robotics and AI technologies could increase unemployment? The results obtained on a sample of 33 OECD countries between 2005 and 2017 allow us to answer in the affirmative: a 10% increase in the stock of industrial robots is associated with a 0.42 point increase in the unemployment rate; and the AI variable constructed from patents is also positively correlated with the unemployment rate, even if the relationship is less robust than for robots.

We continue the analysis by regressing the robots and AI variables on the unemployment rates differentiated by education level, and the results show that there is a lot of heterogeneity between the different groups. For example, the effect of robots is strongest on the unemployment rate of people with upper secondary or post-secondary non-tertiary education, followed closely by the unemployment rate of people with below upper secondary education, and then the weakest effect is on the unemployment rate of people with tertiary education. This first result supports the thesis of a polarization of the labor market induced by automation. Finally, we continue the disaggregation of the unemployment rate by analyzing the effects of these technologies on the unemployment rates by level of education and age groups. Here again, the results highlight strong heterogeneity within a group with the same level of education. For example, the positive correlation between robots and the unemployment rate of individuals with below upper secondary education is found only for young people (25-34 years-olds), and no statistically significant effect is found for other age groups with a similar level of education. These results contribute to the literature by providing additional evidence of the potentially harmful effect of AI and robots on employment at the aggregate level, while highlighting effects that turn out to be positive depending on the education level and age of workers.

In the second chapter, we focus on the theoretical mechanisms explaining the polarization of the labor market. This phenomenon, observed in many developed countries, is characterized by an increase in the share of employment and wages of low- and high-skilled workers, and a decrease in the share of employment and wages of medium-skilled workers. The theoretical literature explaining the link between automation and labor market polarization is based on the routine-biased technical change hypothesis (RBTC), which postulates that automation mainly targets routine tasks performed mostly by

medium-skilled workers. The hypothesis is intuitive at first sight, but it also tends to condition the result obtained: if we postulate that automation targets routine tasks, which are themselves performed by medium-skilled workers, then we indirectly postulate that automation targets medium-skilled workers first, and the dynamic can only lead to a polarization of the labor market. This raises questions about the relevance of this assumption, and about the possibility of generating labor market polarization from less restrictive assumptions.

To answer this question, we have created and programmed a multi-agent model allowing us to simulate the effect of automation on employment and wages. This model is populated by heterogeneous agents, each endowed with different skills that they improve during their professional activities, and that deteriorate during periods of inactivity. Evolving in an uncertain environment that prevents the resolution of an intertemporal maximization program, the agents in the model adopt simple adaptive behaviors that allow them to fulfill an objective: to increase their profit for firms, and to increase their wage for workers. To produce one unit of a good, firms combine different types of jobs following a Leontief-type production function. Agents can only apply for a job if their skills match those required for the job in question. The model includes six different skills rated between zero and seven, the latter being the maximum level of expertise. The wages and skills required for each of these occupations are initialized from U.S. data obtained by aggregating information from the O*net database and the Bureau of Labor Statistics. Finally, firms seek to improve the efficiency of their production process by hiring engineers who conduct research and development activities to improve the level of productivity embodied in capital and to automate some tasks.

First, we test a scenario without automation. The model reproduces some stylized facts, such as endogenous real GDP growth, heterogeneity in the level of productivity of firms, and a positive correlation between the unemployment rate and income inequality. This first scenario leads to wage polarization, linked in particular to a competitive labor market that offers higher wages to the rarest skills, but the distribution of the demand for skills remains unpolarized. Empirically, labor market polarization is characterized by a distortion of the distribution of wages and the skill level of jobs; so this first scenario is not able to fully reproduce this phenomenon.

In a second step, we test a scenario with automation. Unlike the RBTC explanation, we do not make any assumptions about the types of tasks targeted by automation; we simply assume that firms automate some tasks to reduce their production costs. This scenario reproduces the stylized facts listed above, as well as the erratic nature of automation, characterized by spikes. Keeping the same parameters values as in the first scenario, the introduction of automation leads to a polarization of wages and of the distribution of skill demand. This result is explained by the deskilling aspect of automation: when a task is automated, the skills needed to perform it are no longer required, and consequently workers no longer develop this skill. The working time saved is reallocated to other non-automated tasks, allowing the development of skills related to these tasks to be accelerated, because the learning mechanism is based on experience effects. The result is that workers are more specialized but, on average, less qualified

because their skills that are no longer mobilized have depreciated, leading to a loss of versatility. In this model, automation generates technological unemployment, with an increase of 13.6 points in the unemployment rate between the two scenarios. Finally, we test an intermediate scenario in which automation is slower, and show that the degree of labor market polarization and the magnitude of technological unemployment are proportional to the pace of automation.

After having studied the empirical link between robotics, artificial intelligence and unemployment (Chapter 1), and after having highlighted the theoretical mechanisms that can explain the emergence of technological unemployment and labor market polarization (Chapter 2), we focus in chapter 3 on policies that can counteract these negative effects of automation. Technological progress has historically been a vector of economic development; it is therefore not a question of trying to slow it down, but of deploying policy tools to accompany the changes it brings about to the labor market. To do this, we extend the model developed in Chapter 2 to test four policies: regulation of layoffs and resignations to stabilize the labor market, a minimum wage, an unemployment insurance and a training system. The results indicate that these policies have heterogeneous effects that are sometimes complementary, such as the regulation of labor contracts and the introduction of a minimum wage, which reduce wage polarization, while the training system reduces skill polarization; and sometimes opposite, such as the regulation of labor contracts, which negatively impacts productivity, while the minimum wage and unemployment insurance increase it.

Some of the policies tested appear to be effective in reducing the polarization of either the wage distribution or the demand for skills, but none of them manages to act on both simultaneously. In a final scenario, we test all four policies together, and the results are very positive, with a reduction in wage and skill polarization, an increase in the median real wage, in the general level of skills, in labor productivity and in real GDP, which leads to a decrease in the unemployment rate. These results show that technological unemployment and labor market polarization can be effectively tackled with appropriate policies; and that policymakers need to think about public policies as a whole and not separately, at the risk of missing potential positive synergies.

We have shown in Chapter 1 that there is a statistically significant link between robots, AI and unemployment. In Chapter 2, we explained the theoretical mechanisms behind the emergence of technological unemployment and labor market polarization. In Chapter 3, we tested different policies to fight against these two phenomena and showed that they can be mitigated through public intervention. These three chapters contribute to the literature on the link between technological progress, employment and inequality by covering both the empirical, theoretical and policy aspects of the subject; but they also suffer from several limitations.

In Chapter 1, the effects of robots on unemployment are easily interpretable, while the effects of artificial intelligence are less robust and more difficult to interpret. While for the “robots” variable we use data on the stock deployed in firms, the “AI” variable is built using data from patents related to artificial intelligence. Consequently, our data on AI are more a measure of innovation in this field than a measure of the stock of

software and algorithms used in production processes, a discrepancy that may explain the lack of robustness of the results related to this variable. The best solution would be to have data on the stock of AI-related capital, as for industrial robots; but such data are not, to our knowledge, available at the macroeconomic level.

In chapter 2, firms compete in a single market by producing a homogeneous good. Although the heterogeneity of wages and skills required for a same type of occupation already emerges from the properties of the model, it would be interesting to observe whether the dynamics of these two variables differ in a multi-sectoral model. Indeed, data from the Bureau of Labor Statistics indicates that there are large disparities in wages for the same job between different sectors; and it would be interesting to study whether the heterogeneity of technological trajectories between sectors can be explained by these differences. Finally, a sectoral approach would make it possible, through an input-output matrix, to include intermediate consumption in the production function to study the inter-sectoral compensation mechanisms. Indeed, automation in sector A leads to productivity gains and thus to a decrease in unit production costs and hence in prices. This fall in price is also reflected in the other sectors that use the output of sector A as an input in their production function, allowing firms in these sectors to be more competitive, to gain market share, to increase their production and, finally, to create jobs. Finally, opening the model to international trade could mitigate the extent of technological unemployment and labor market polarization generated by the model. By lowering costs, automation allows domestic firms to reduce the price-competitiveness gap with countries where labor is cheap, thus preserving jobs.

The last limitation raised for chapter 2 can also be applied to chapter 3, since the combination of the four policies tested has the effect of increasing the median real wage by 35 percent, which is not an issue in a closed economy, but could be detrimental to the competitiveness of domestic firms in an open economy. The net effect on employment, which is slightly positive in the case of a closed economy, would then be much more uncertain and could even turn negative.

These limitations provide avenues for future research. The combination of more precise data and an extended model constitutes a fertile field of research for a more detailed analysis of the transformations of the labor market induced by technological progress.

Appendix

	$U_{i=1}^{k=25-34}$	$U_{i=1}^{k=35-44}$	$U_{i=1}^{k=45-54}$	$U_{i=1}^{k=55-64}$
$U_i^k(t-1)$	-0.191* (0.099)	-0.178** (0.067)	0.010 (0.085)	-0.176* (0.070)
$Y(t)$	-0.555*** (0.057)	-0.523*** (0.149)	-0.598*** (0.140)	-0.300** (0.084)
$Y(t-1)$	-0.548*** (0.093)	-0.334*** (0.101)	-0.208*** (0.093)	-0.226*** (0.037)
$IR_i^k(t)$	0.165* (0.084)	-0.024 (0.086)	0.172 (0.105)	0.155* (0.088)
$IR_i^k(t-1)$	0.003 (0.120)	0.023 (0.073)	0.017 (0.144)	0.106* (0.058)
$AI(t)$	-0.087 (0.058)	-0.058 (0.055)	-0.000 (0.081)	-0.096 (0.058)
$AI(t-1)$	0.043 (0.057)	0.002 (0.035)	0.038 (0.024)	-0.025 (0.035)
$AI(t-2)$	0.033 (0.054)	-0.006 (0.041)	0.030 (0.020)	-0.001 (0.024)
$Robots(t)$	0.024 (0.015)	0.004 (0.011)	-0.003 (0.006)	0.029** (0.014)
$Robots(t-1)$	0.005 (0.008)	0.033* (0.016)	0.004 (0.015)	-0.001 (0.009)
$Robots(t-2)$	0.007 (0.012)	0.024 (0.028)	-0.002 (0.007)	0.021** (0.009)
<i>Intercept</i>	-0.277 (1.515)	0.498 (1.134)	-0.192 (0.790)	1.420* (0.795)
Time dummies	Yes	Yes	Yes	Yes
Observations	319	319	319	319
R^2	0.5179	0.4378	0.4353	0.3526

Notes: Robust standard errors in parenthesis, clustered at country-level. * $p < 0.10$
 ** $p < 0.05$ *** $p < 0.01$

Table 3.10: $U_{i=1}$ differentiated by age group - FE estimates

	$U_{i=2}^{k=25-34}$	$U_{i=2}^{k=35-44}$	$U_{i=2}^{k=45-54}$	$U_{i=2}^{k=55-64}$
$U_i^k(t-1)$	0.091 (0.063)	0.022 (0.074)	0.041 (0.076)	-0.077 (0.086)
$Y(t)$	-0.480*** (0.051)	-0.384*** (0.051)	-0.347*** (0.062)	-0.262** (0.045)
$Y(t-1)$	-0.146*** (0.069)	-0.166*** (0.035)	-0.155*** (0.043)	-0.259*** (0.067)
$IR_i^k(t)$	0.025 (0.075)	-0.106 (0.111)	0.024 (0.165)	-0.049 (0.047)
$IR_i^k(t-1)$	-0.149* (0.075)	-0.008 (0.110)	-0.064 (0.101)	-0.041 (0.038)
$AI(t)$	-0.015 (0.028)	-0.021 (0.026)	-0.019 (0.024)	-0.001 (0.020)
$AI(t-1)$	-0.038* (0.020)	-0.017 (0.021)	-0.012 (0.025)	-0.043** (0.019)
$AI(t-2)$	0.043* (0.024)	0.024 (0.016)	0.012 (0.013)	0.017 (0.016)
$Robots(t)$	0.010* (0.005)	0.020** (0.009)	0.020*** (0.005)	0.002 (0.004)
$Robots(t-1)$	-0.006 (0.006)	-0.021*** (0.004)	-0.002 (0.005)	0.011 (0.008)
$Robots(t-2)$	-0.008 (0.005)	-0.004 (0.010)	-0.003 (0.008)	-0.003 (0.010)
<i>Intercept</i>	0.542 (0.653)	0.898 (0.664)	0.680* (0.368)	0.409 (0.527)
Time dummies	Yes	Yes	Yes	Yes
Observations	319	319	319	319
R^2	0.6294	0.5760	0.5895	0.5173

Notes: Robust standard errors in parenthesis, clustered at country-level. * $p < 0.10$
** $p < 0.05$ *** $p < 0.01$

Table 3.11: $U_{i=2}$ differentiated by age group - FE estimates

	$U_{i=3}^{k=25-34}$	$U_{i=3}^{k=35-44}$	$U_{i=3}^{k=45-54}$	$U_{i=3}^{k=55-64}$
$U_i^k(t-1)$	0.009 (0.094)	-0.106 (0.099)	-0.105** (0.046)	-0.218** (0.084)
$Y(t)$	-0.273*** (0.072)	-0.147*** (0.037)	-0.129*** (0.020)	-0.058*** (0.017)
$Y(t-1)$	-0.093** (0.042)	-0.140*** (0.031)	-0.102*** (0.028)	-0.173*** (0.047)
$IR_i^k(t)$	0.082 (0.085)	0.101 (0.123)	-0.109* (0.062)	-0.028 (0.034)
$IR_i^k(t-1)$	-0.014 (0.073)	0.020 (0.081)	-0.056 (0.065)	0.034 (0.033)
$AI(t)$	0.014 (0.018)	0.013 (0.012)	-0.008 (0.010)	0.020 (0.021)
$AI(t-1)$	0.002 (0.020)	-0.005 (0.010)	0.011 (0.013)	-0.023 (0.017)
$AI(t-2)$	0.010* (0.014)	-0.002 (0.018)	0.016 (0.010)	0.018 (0.018)
$Robots(t)$	-0.002 (0.012)	0.000 (0.005)	0.011* (0.006)	0.017*** (0.005)
$Robots(t-1)$	0.001 (0.005)	0.007 (0.007)	0.001 (0.004)	0.017* (0.010)
$Robots(t-2)$	0.010* (0.005)	-0.002 (0.005)	-0.003 (0.003)	-0.006 (0.007)
<i>Intercept</i>	0.159 (0.393)	0.253 (0.319)	-0.122 (0.208)	-0.426* (0.216)
Time dummies	Yes	Yes	Yes	Yes
Observations	319	319	319	319
R^2	0.4423	0.3990	0.4955	0.3174

Notes: Robust standard errors in parenthesis, clustered at country-level. * $p < 0.10$
** $p < 0.05$ *** $p < 0.01$

Table 3.12: $U_{i=3}$ differentiated by age group - FE estimates

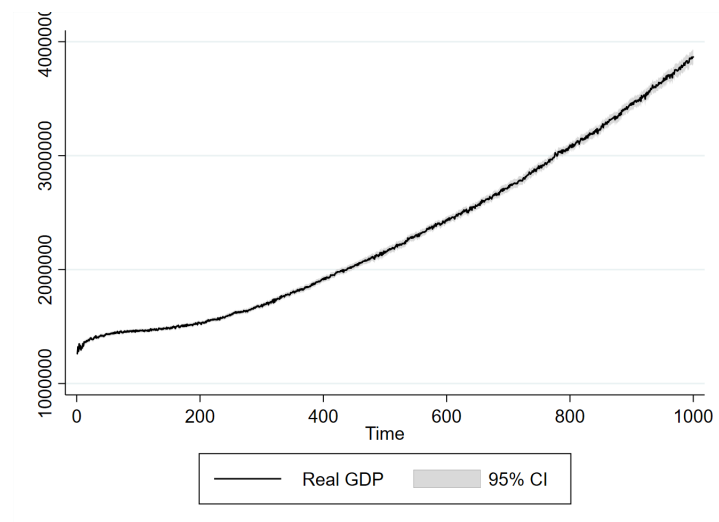


Figure 3.5: Real GDP - Baseline scenario - 1000 steps and 100 replications

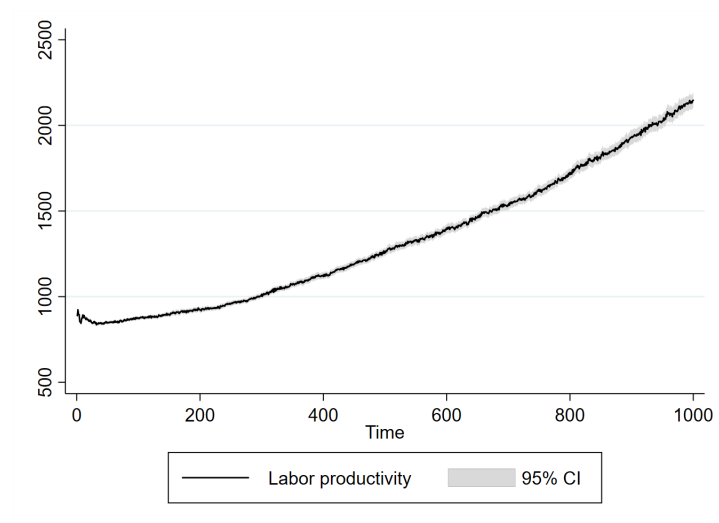


Figure 3.6: Labor productivity - Baseline scenario - 1000 steps and 100 replications

Agent	Component	Equation
Firms	Production	$Y_f(t) = \min \left\{ \frac{L_{1,f}(t)}{A_{1,f}(t)}; \dots; \frac{L_{m,f}(t)}{A_{m,f}(t)}; \dots; \frac{L_{M,f}(t)}{A_{M,f}(t)} \right\} \quad (2.1)$
	Desired production	$Y_f^D(t) = \alpha Y_f^D(t-1) + (1-\alpha) D_f(t-1) - (Y_f(t-1) - D_f(t-1)) \quad (2.5)$
	Job productivity	$\frac{1}{A_{m,f}(t)} = \min \left\{ \frac{1}{B_{1,m,f}(t)}; \dots; \frac{1}{B_{i,m,f}(t)}; \dots; \frac{1}{B_{I,m,f}(t)} \right\} \quad (2.3)$
	Task intensity	$B_{i,m,f}(t) = \begin{cases} 0 & \text{if the task is automated} \\ \left(a_{i,m,f}(t-1) \sum_{j=1}^{L_{m,f}(t)} \frac{b_{j,i,m,f}(t)}{L_{m,f}(t-1)} \right)^{-1} & \text{if the task is performed by workers} \end{cases} \quad (2.4)$
	Labor demand	$L_{m,f}^D(t) = A_{m,f}(t) Y_f^D(t) + (1-\lambda)(T_{m,f}(t-1) - L_{m,f}(t-1)) \quad (2.14)$
	Wage proposal	$w_{m,f}(t) = w_{m,f}(t-1) \left[1 + \xi_1 \left(\frac{A_f(t-1)}{A_f(t-2)} - 1 \right) + \xi_2 \max \left\{ 0; 1 - \frac{L_{m,f}^S(t-1)}{L_{m,f}^D(t-1)} \right\} \right] \quad (2.15)$
	R&D	$R_f(t) = \min \left\{ \eta p_f(t-1) * Q_f(t-1); \sum_{\tau=1}^{t-1} (\pi_f(\tau) - R_f(\tau)) \right\} \quad (2.24)$
	Technological frontier	$\sigma_{i,f}(t) = \sigma_{i,f}(t-1) (1 + \max\{0; \epsilon_{i,f}^\sigma(t)\}) \quad (2.30)$ with $\epsilon_{i,f}^\sigma(t) \sim N(0; \beta(s_i^{MAX} - \sigma_{i,f}(t-1)))$
	Capital productivity	$a_{i,m,f}(t) = a_{i,m,f}(t-1) (1 + \min\{0; \epsilon_{i,m,f}^a(t)\}) \quad (2.31)$ with $\epsilon_{i,f}^a(t) \sim N(0; \gamma_i) \forall m$
	Workers	Skills
Worker's efficiency		$b_{j,i,m,f}(t) = e^{\kappa(s_{j,i,m,f}(t-1) - s_i^{MAX})} \quad (3.4)$
Desired level of consumption		$C_j^D(t) = \pi * C_j(t-1) + (1-\pi) * W_j(t) \quad (2.7)$

Table 3.13: Main equations of the model

Parameter	Description	Value
F	Number of firms	10
J	Number of agents	2000
δ_1	Skills decline rate	0.00125
δ_2	Skills accumulation rate	0.003
κ	Skills-productivity elasticity	0.03
s_i^{MAX}	Maximum skill level	7
α	Degree of adaptation of the firm	0.5
ν	Sensitivity of the mark-up to market dynamics	0
ϕ	Degree of competition among firms	0.02
ω	Importance of prices in consumers' behaviour	0.5
λ	Institutional frame of the labor market	0
ξ_1	Wage indexation on firm's productivity growth	0.1
ξ_2	Weight of wage premium	0.0005
η	Share of sales invested in R&D	0.0815
ρ	Elasticity between the number of engineers (adjusted by quality) and the probability to innovate	$3 \cdot 10^{-6}$
β	Magnitude of technological progress	0
γ_i	Magnitude of embodied technical progress	0.25
k	standard deviation for the initialization of workers' skills	0.005
k^u	standard deviation for the initialization of unemployed agents' skills	$1 \cdot 10^{-8}$
ι	Probability for a worker to look for job opportunities	0.1
ψ	Indexation of consumption on past consumption level	0.8

Table 3.14: Parameters setting.

Occupation group	Technical	System	Social	Management	Basic	CPS	Wage	Share (%)
Management, Business and Finance	0.87	3.43	3.47	3.05	3.46	3.57	1926	11.04
Engineers, Mathematicians and computer scientists	2.13	3.67	3.07	2.58	3.71	3.88	1754	4.84
Other services	0.81	2.38	2.85	1.83	2.77	2.74	929	37.81
Sales, office and administrative	0.64	2.38	2.97	1.88	2.76	2.72	809	23.09
Primary sector and maintenance	2.25	2.39	2.50	1.95	2.61	2.86	974	8.45
Production occupations, transportation and materials	1.71	2.17	2.29	1.77	2.39	2.58	747	14.77

Table 3.15: Data used to initialize the model.

Parameter	Description	Baseline	FRR	MW (LMW)	UB (HUB)	TP
F	Number of firms	10	-	-	-	-
J	Number of agents	2000	-	-	-	-
δ_1	Skills decline rate	0.00125	-	-	-	-
δ_2	Skills accumulation rate	0.003	-	-	-	-
κ	Skills-productivity elasticity	0.03	-	-	-	-
s_i^{MAX}	Maximum skill level	7	-	-	-	-
α	Degree of adaptation of the firm	0.5	-	-	-	-
ν	Sensitivity of the mark-up to market dynamics	0	-	-	-	-
ϕ	Degree of competition among firms	0.02	-	-	-	-
ω	Importance of prices in consumers' behaviour	0.5	-	-	-	-
λ	Institutional frame of the labor market	0	-	-	-	-
ξ_1	Wage indexation on firm's productivity growth	0.1	-	-	-	-
ξ_2	Weight of the wage premium	0.0005	-	-	-	-
η	Share of sales invested in R&D	0.0815	-	-	-	-
ρ	Elasticity between the number of engineers (adjusted by quality) and the probability to innovate	$3 \cdot 10^{-6}$	-	-	-	-
β	Magnitude of technological progress	0.1	-	-	-	-
γ_i	Magnitude of embodied technical progress	0.25	-	-	-	-
k	standard deviation for the initialization of workers' skills	0.005	-	-	-	-
k^u	standard deviation for the initialization of unemployed agents' skills	$1 \cdot 10^{-8}$	-	-	-	-
ι	Probability for a worker to look for job opportunities	0.1	0	-	-	-
ξ_3	Indexation of the minimum wage on aggregate productivity	/	/	0.5 (0.25)	/	/
τ_1	Contribution rate to the unemployment insurance	/	/	/	0.1 (0.19)	/
τ_2	Indexation of unemployment benefits on the last wage	/	/	/	0.5 (0.8)	/
δ^{MAX}	Training threshold	/	/	/	/	0.25
ξ_4	Contribution rate to the training system	/	/	/	/	0.01

Table 3.16: **Parameters setting.** - = same value than in the baseline scenario
/ = absence of the variable in this scenario

	FRR	TP	Difference	P-value
Skills polarization index	0.183	0.154	0.0284	0.0000
Wages polarization index	0.618	0.741	-0.123	1.63e-08
Gini index	0.297	0.453	-0.156	1.02e-22
Median skill index	3.519	3.860	-0.341	6.11e-10
Real median wage	864.7	770.1	94.59	0.0109
Unemployment rate	0.106	0.229	-0.123	2.86e-21
Labor productivity	1094.8	1114.7	-19.98	0.319
Share of automated tasks	0.414	0.327	0.0872	8.91e-20
Real GDP	1955384.0	1708612.4	246771.6	3.93e-15

Table 3.17: *FRR and TP scenarios*

	FRR	MW	Difference	P-value
Skills polarization index	0.183	0.210	-0.0271	0.0026
Wages polarization index	0.618	0.426	0.192	2.1e-17
Gini index	0.297	0.337	-0.0403	0.0023
Median skill index	3.519	3.207	0.312	5.10e-08
Real median wage	864.7	882.4	-17.66	0.605
Unemployment rate	0.106	0.325	-0.219	1.86e-34
Labor productivity	1094.8	1258.5	-163.8	2.67e-08
Share of automated tasks	0.414	0.410	0.00439	0.562
Real GDP	1955384.0	1680562.4	274821.6	1.86e-16

Table 3.18: *FRR and MW scenarios*

	FRR	UB	Difference	P-value
Skills polarization index	0.183	0.198	-0.0158	0.228
Wages polarization index	0.618	0.758	-0.140	4.59e-09
Gini index	0.297	0.532	-0.0403	2.37e-32
Median skill index	3.519	3.233	0.286	0.0000
Real median wage	864.7	776.2	88.55	0.0173
Unemployment rate	0.106	0.342	-0.236	2.44e-38
Labor productivity	1094.8	1302.6	-207.9	0.0000
Share of automated tasks	0.414	0.407	0.00761	0.361
Real GDP	1955384.0	1691814.3	263569.7	6.38e-08

Table 3.19: *FRR and UB scenarios*

	MW	UB	Difference	P-value
Skills polarization index	0.210	0.198	0.0113	0.429
Wages polarization index	0.426	0.758	-0.332	1.69e-23
Gini index	0.337	0.532	-0.195	3.17e-20
Median skill index	3.207	3.233	-0.0259	0.740
Real median wage	882.4	776.2	106.2	0.00941
Unemployment rate	0.325	0.342	-0.0167	0.260
Labor productivity	1258.5	1302.6	-44.12	0.390
Share of automated tasks	0.410	0.407	0.0032	0.695
Real GDP	1680562.4	1691814.3	-11251.9	0.812

Table 3.20: *MW and UB scenarios*

	MW	TP	Difference	P-value
Skills polarization index	0.210	0.154	0.0555	3.65e-08
Wages polarization index	0.426	0.741	-0.316	9.43e-24
Gini index	0.337	0.532	-0.116	6.91e-11
Median skill index	3.207	3.860	-0.653	3.82e-17
Real median wage	882.4	770.1	112.3	0.0060
Unemployment rate	0.325	0.229	0.0956	6.57e-10
Labor productivity	1258.5	1114.7	143.8	0.0000
Share of automated tasks	0.410	0.327	0.0828	6.67e-19
Real GDP	1680562.4	1708612.4	-28050.0	0.355

Table 3.21: *MW and TP scenarios*

	UB	TP	Difference	P-value
Skills polarization index	0.198	0.154	0.0442	0.0013
Wages polarization index	0.758	0.741	0.0167	0.527
Gini index	0.532	0.453	0.0785	0.0000
Median skill index	3.233	3.860	-0.627	6.64e-13
Real median wage	776.2	770.1	6.048	0.886
Unemployment rate	0.342	0.229	0.112	6.94e-13
Labor productivity	1302.6	1114.7	187.9	0.0002
Share of automated tasks	0.407	0.327	0.0796	6.99e-16
Real GDP	1691814.3	1708612.4	-16798.1	0.719

Table 3.22: *UB and TP scenarios*

Bibliography

- Aboal, D., Garda, P., Lanzilotta, B. and Perera, M.: 2015, Innovation, firm size, technology intensity, and employment generation: Evidence from the uruguayan manufacturing sector, *Emerging Markets Finance and Trade* **51**(1), 3–26.
- Acemoglu, D. and Autor, D.: 2011, Skills, tasks and technologies: Implications for employment and earnings, *Handbook of labor economics*, Vol. 4, Elsevier, pp. 1043–1171.
- Acemoglu, D., Autor, D., Hazell, J. and Restrepo, P.: 2020, Ai and jobs: Evidence from online vacancies, *Technical report*, National Bureau of Economic Research.
- Acemoglu, D., Lelarge, C. and Restrepo, P.: 2020, Competing with robots: Firm-level evidence from france, *AEA Papers and Proceedings*, Vol. 110, pp. 383–88.
- Acemoglu, D. and Restrepo, P.: 2020, Robots and jobs: Evidence from us labor markets, *Journal of Political Economy* **128**(6), 2188–2244.
- Addison, J. T., Centeno, M. and Portugal, P.: 2010, Unemployment benefits and reservation wages: key elasticities from a stripped-down job search approach, *Economica* **77**(305), 46–59.
- Aghion, P., Antonin, C., Bunel, S. and Jaravel, X.: 2020, What are the labor and product market effects of automation? new evidence from france.
- Albertini, J., Hairault, J.-O., Langot, F. and Sopraseuth, T.: 2017, A tale of two countries: A story of the french and us polarization.
- Amable, B., Demmou, L. and Gatti, D.: 2011, The effect of employment protection and product market regulation on labour market performance: substitution or complementarity?, *Applied Economics* **43**(4), 449–464.
- Amendola, M. and Gaffard, J. L.: 1998, *Out of Equilibrium*, Clarendon Press, Oxford.
- Anelli, M., Colantone, I. and Stanig, P.: 2019, We were the robots: Automation and voting behavior in western europe, *BAFFI CAREFIN Centre Research Paper* (2019-115).

- Antonucci, T.: 2007, Innovation and employment in europe: A sectoral perspective, *Innovation, industrial dynamics and structural transformation*, Springer, pp. 255–279.
- Antonucci, T. and Pianta, M.: 2002, Employment effects of product and process innovation in europe, *International Review of Applied Economics* **16**(3), 295–307.
- Arntz, M., Gregory, T. and Zierahn, U.: 2016, The risk of automation for jobs in oecd countries.
- Autor, D. H., Katz, L. F. and Kearney, M. S.: 2006, The polarization of the us labor market, *The American economic review* **96**(2), 189–194.
- Autor, D. H., Levy, F. and Murnane, R. J.: 2003, The skill content of recent technological change: An empirical exploration, *The Quarterly journal of economics* **118**(4), 1279–1333.
- Autor, D. and Salomons, A.: 2018, Is automation labor-displacing? productivity growth, employment, and the labor share, *Technical report*, National Bureau of Economic Research.
- Autor, H. and 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.
- Babecky, J., Caju, P. D., Kosma, T., Lawless, M., Messina, J. and Rõõm, T.: 2010, *Downward nominal and real wage rigidity: Survey evidence from European firms*, The World Bank.
- Baensch, L., Lanzaot, M. L., Lotti, G. and Stucchi, R.: 2019, Do labor market regulations affect the link between innovation and employment? evidence from latin america, *The BE Journal of Economic Analysis & Policy* **19**(3).
- Barbieri, L., Piva, M. and Vivarelli, M.: 2019, R&d, embodied technological change, and employment: evidence from italian microdata, *Industrial and Corporate Change* **28**(1), 203–218.
- Bartelsman, E. J. and Doms, M.: 2000, Understanding productivity: Lessons from longitudinal microdata, *Journal of Economic literature* **38**(3), 569–594.
- Benavente, J. M. and Lauterbach, R.: 2008, Technological innovation and employment: complements or substitutes?, *The European Journal of Development Research* **20**(2), 318–329.
- Bertani, F., Raberto, M. and Teglio, A.: 2020, The productivity and unemployment effects of the digital transformation: an empirical and modelling assessment, *Review of Evolutionary Political Economy* **1**(3), 329–355.
- 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).

- Bessen, J. E., Goos, M., Salomons, A. and Van den Berge, W.: 2019, Automatic reaction-what happens to workers at firms that automate?, *Boston Univ. School of Law, Law and Economics Research Paper* .
- Bessen, J., Goos, M., Salomons, A. and van den Berge, W.: 2020, Firm-level automation: Evidence from the netherlands, *AEA Papers and Proceedings*, Vol. 110, pp. 389–93.
- Bianchini, S. and Pellegrino, G.: 2019, Innovation persistence and employment dynamics, *Research Policy* **48**(5), 1171–1186.
- Blanas, S., Gancia, G. and Lee, S. Y.: 2019, Who is afraid of machines?, *Economic Policy* **34**(100), 627–690.
- Blanchard, O. and Wolfers, J.: 2000, The role of shocks and institutions in the rise of european unemployment: the aggregate evidence, *The Economic Journal* **110**(462), C1–C33.
- Blundell, R. and Bond, S.: 1998, Initial conditions and moment restrictions in dynamic panel data models, *Journal of econometrics* **87**(1), 115–143.
- Bogliacino, F. and Pianta, M.: 2010, Innovation and employment: a reinvestigation using revised pavitt classes, *Research Policy* **39**(6), 799–809.
- Bogliacino, F., Piva, M. and Vivarelli, M.: 2012, R&d and employment: An application of the lsdvc estimator using european microdata, *Economics Letters* **116**(1), 56–59.
- Bonfiglioli, A., Crinò, R., Fadinger, H. and Gancia, G.: 2020, Robot imports and firm-level outcomes.
- Boockmann, B.: 2010, The combined employment effects of minimum wages and labor market regulation: A meta-analysis.
- Bossler, M. and Gerner, H.-D.: 2020, Employment effects of the new german minimum wage: Evidence from establishment-level microdata, *ILR Review* **73**(5), 1070–1094.
- Brancaccio, E., De Cristofaro, F. and Giammetti, R.: 2020, A meta-analysis on labour market deregulations and employment performance: no consensus around the imf-oecd consensus, *Review of Political Economy* **32**(1), 1–21.
- Brancaccio, E., Garbellini, N. and Giammetti, R.: 2018, Structural labour market reforms, gdp growth and the functional distribution of income, *Structural Change and Economic Dynamics* **44**, 34–45.
- Broadberry, S., Campbell, B. M., Klein, A., Overton, M. and Van Leeuwen, B.: 2015, *British economic growth, 1270–1870*, Cambridge University Press.
- Brunello, G. and Wruuck, P.: 2019, Skill shortages and skill mismatch in europe: a review of the literature.

- Burstein, A., Morales, E. and Vogel, J.: 2019, Changes in between-group inequality: computers, occupations, and international trade, *American Economic Journal: Macroeconomics* **11**(2), 348–400.
- Caines, C., Hoffmann, F. and Kambourov, G.: 2017, Complex-task biased technological change and the labor market, *Review of Economic Dynamics* **25**, 298–319.
- Caliendo, M., Wittbrodt, L. and Schröder, C.: 2019, The causal effects of the minimum wage introduction in germany—an overview, *German Economic Review* **20**(3), 257–292.
- Calvino, F.: 2019, Technological innovation and the distribution of employment growth: a firm-level analysis, *Industrial and Corporate Change* **28**(1), 177–202.
- Calvino, F. and Virgillito, M. E.: 2018, The innovation-employment nexus: a critical survey of theory and empirics, *Journal of Economic surveys* **32**(1), 83–117.
- Carbonero, F., Ernst, E. and Weber, E.: 2020, Robots worldwide: The impact of automation on employment and trade.
- Caselli, M., Fracasso, A., Scicchitano, S., Traverso, S. and Tundis, E.: 2021, Stop worrying and love the robot: An activity-based approach to assess the impact of robotization on employment dynamics, *Technical report*, GLO Discussion Paper.
- Castillo, V., Maffioli, A., Rojo, S. and Stucchi, R.: 2014, The effect of innovation policy on smes’ employment and wages in argentina, *Small Business Economics* **42**(2), 387–406.
- Chiacchio, F., Petropoulos, G. and Pichler, D.: 2018, The impact of industrial robots on eu employment and wages: A local labour market approach, *Technical report*, Bruegel working paper.
- Chletsos, M. and Giotis, G. P.: 2015, The employment effect of minimum wage using 77 international studies since 1992: A meta-analysis.
- Chodorow-Reich, G., Coglianesi, J. and Karabarbounis, L.: 2019, The macro effects of unemployment benefit extensions: a measurement error approach, *The Quarterly Journal of Economics* **134**(1), 227–279.
- Chui, M. and Malhotra, S.: 2018, Ai adoption advances, but foundational barriers remain, *Mckinsey and Company* .
- Ciarli, T., Lorentz, A., Savona, M. and Valente, M.: 2010, The effect of consumption and production structure on growth and distribution. a micro to macro model, *Metroeconomica* **61**(1), 180–218.
- Ciminelli, G., Duval, R. and Furceri, D.: 2018, Employment protection deregulation and labor shares in advanced economies, *The Review of Economics and Statistics* pp. 1–44.

- Cirera, X. and Sabetti, L.: 2019, The effects of innovation on employment in developing countries: evidence from enterprise surveys, *Industrial and Corporate Change* **28**(1), 161–176.
- Clark, C.: 1940, The conditions of economic progress, mcmillan.
- Clark, S. L., Iceland, J., Palumbo, T., Posey, K. and Weismantle, M.: 2003, Comparing employment, income, and poverty: Census 2000 and the current population survey, *Bureau of the Census, US Department of Commerce, September* .
- Coad, A. and Rao, R.: 2010, Firm growth and r&d expenditure, *Economics of Innovation and New Technology* **19**(2), 127–145.
- Coad, A. and Rao, R.: 2011, The firm-level employment effects of innovations in high-tech us manufacturing industries, *Journal of Evolutionary Economics* **21**(2), 255–283.
- Compagnucci, F., Gentili, A., Valentini, E. and Gallegati, M.: 2019, Robotization and labour dislocation in the manufacturing sectors of oecd countries: a panel var approach, *Applied Economics* **51**(57), 6127–6138.
- Consoli, D., Marin, G., Rentocchini, F. and Vona, F.: 2019, Routinization, within-occupation task changes and long-run employment dynamics, *Technical report, LEM Working Paper Series*.
- Crespi, G. and Tacsir, E.: 2012, Effects of innovation on employment in latin america," idb publications 78759, inter-american development bank.
- Crespi, G., Tacsir, E. and Pereira, M.: 2019, Effects of innovation on employment in latin america, *Industrial and Corporate Change* **28**(1), 139–159.
- Dauth, W., Findeisen, S., Suedekum, J. and Woessner, N.: 2021, The Adjustment of Labor Markets to Robots, *Journal of the European Economic Association* . jvab012. **URL:** <https://doi.org/10.1093/jeea/jvab012>
- Dauth, W., Findeisen, S., Suedekum, J., Woessner, N. et al.: 2018, Adjusting to robots: Worker-level evidence, *Unpublished manuscript, Julius-Maximilians-Universität Würzburg* .
- Dawid, H. and Hepp, J.: 2021, Distributional effects of technological regime changes: Hysteresis, concentration and inequality dynamics.
- De Elejalde, R., Giuliadori, D. and Stucchi, R.: 2015, Employment and innovation: Firm-level evidence from argentina, *Emerging Markets Finance and Trade* **51**(1), 27–47.
- De Vries, G. J., Gentile, E., Miroudot, S. and Wacker, K. M.: 2020, The rise of robots and the fall of routine jobs, *Labour Economics* **66**, 101885.

- Debrun, X.: 2003, Unemployment and labor market institutions: Why reforms pay off, *World economic outlook* pp. 129–150.
- Dengler, K. and Matthes, B.: 2018, The impacts of digital transformation on the labour market: Substitution potentials of occupations in germany, *Technological Forecasting and Social Change* **137**, 304–316.
- Deyshappriya, N.: 2017, Impact of macroeconomic factors on income inequality and income distribution in asian countries.
- Dixon, J., Hong, B. and Wu, L.: 2021, The robot revolution: Managerial and employment consequences for firms, *Management Science* .
- Domini, G., Grazzi, M., Moschella, D. and Treibich, T.: 2021, Threats and opportunities in the digital era: automation spikes and employment dynamics, *Research Policy* **50**(7), 104137.
- Dosi, G.: 2008, Statistical regularities in the evolution of industries. a guide through some evidence and challenges for the theory, *L'industria* (2), 185–220.
- Dosi, G., Fagiolo, G. and Roventini, A.: 2006, An evolutionary model of endogenous business cycles, *Computational Economics* **27**, 3–34.
- Dosi, G., Fagiolo, G. and Roventini, A.: 2010, Schumpeter meeting keynes: A policy-friendly model of endogenous growth and business cycles, *Journal of Economic Dynamics and Control* **34**(9), 1748–1767.
- Dosi, G., Pereira, M. C., Roventini, A. and Virgillito, M. E.: 2017, When more flexibility yields more fragility: the microfoundations of keynesian aggregate unemployment, *Journal of Economic Dynamics and Control* **81**, 162–186.
- Dosi, G., Pereira, M. C., Roventini, A. and Virgillito, M. E.: 2018, Causes and consequences of hysteresis: aggregate demand, productivity, and employment, *Industrial and Corporate Change* **27**(6), 1015–1044.
- Dosi, G., Pereira, M. C., Roventini, A., Virgillito, M. E. et al.: 2021, Technological paradigms, labour creation and destruction in a multi-sector agent-based model, *Technical report*, Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies.
- Dosi, G., Piva, M., Virgillito, M. E. and Vivarelli, M.: 2021, Embodied and disembodied technological change: The sectoral patterns of job-creation and job-destruction, *Research Policy* **50**(4), 104199.
- Dosi, G. and Roventini, A.: 2019, More is different... and complex! the case for agent-based macroeconomics, *Journal of Evolutionary Economics* **29**(1), 1–37.
- Dottori, D.: 2021, Robots and employment: evidence from italy, *Economia Politica* pp. 1–57.

- Duby, G. and Wallon, A.: 1976, Histoire de la France rurale, t. 3. de 1789 à 1914.
- Dustmann, C., Lindner, A., Schönberg, U., Umkehrer, M. and Vom Berge, P.: 2020, Reallocation effects of the minimum wage, *Centre for Research and Analysis of Migration Discussion Paper* **7**, 20.
- Elmeskov, J., Martin, J. P. and Scarpetta, S.: 1998, Key lessons for labour market reforms: Evidence from OECD countries' experience, *Swedish economic policy review* **5**(2).
- Evangelista, R., Guerrieri, P. and Meliciani, V.: 2014, The economic impact of digital technologies in Europe, *Economics of Innovation and New Technology* **23**(8), 802–824.
- Evangelista, R. and Vezzani, A.: 2012, The impact of technological and organizational innovations on employment in European firms, *Industrial and Corporate Change* **21**(4), 871–899.
- Fagiolo, G., Dosi, G. and Gabriele, R.: 2004, Matching, Bargaining, and Wage Setting in an Evolutionary Model of Labor Market and Output Dynamics, *Advances in Complex Systems* **14**, 237–273.
- Farber, H. S. and Valletta, R. G.: 2015, Do extended unemployment benefits lengthen unemployment spells? Evidence from recent cycles in the US labor market, *Journal of Human Resources* **50**(4), 873–909.
- Feldmann, H.: 2013, Technological unemployment in industrial countries, *Journal of Evolutionary Economics* **23**(5), 1099–1126.
- Felten, E. W., Raj, M. and Seamans, R.: 2018, A method to link advances in artificial intelligence to occupational abilities, *AEA Papers and Proceedings*, Vol. 108, pp. 54–57.
- Felten, E. W., Raj, M. and Seamans, R.: 2019, The occupational impact of artificial intelligence: Labor, skills, and polarization, *NYU Stern School of Business*.
- Fierro, L. E., Caiani, A., Russo, A. et al.: 2021, Automation, job polarisation, and structural change, *Technical report*.
- Fossen, F. M. and Sorgner, A.: 2019, New digital technologies and heterogeneous employment and wage dynamics in the United States: Evidence from individual-level data.
- Fredriksson, P. and Söderström, M.: 2008, Do unemployment benefits increase unemployment? New evidence on an old question, *Technical report*, Working Paper.
- Frey, C. B. and Osborne, M. A.: 2014, Agiletown: The relentless march of technology and London's response, *London Futures*.

- Frey, C. B. and Osborne, M. A.: 2017, The future of employment: How susceptible are jobs to computerisation?, *Technological forecasting and social change* **114**, 254–280.
- Fuei, L. K.: 2017, Automation, computerization and future employment in singapore, *Journal of Southeast Asian Economies* pp. 388–399.
- Furukawa, Y. and Toyoda, H.: 2018, The recent process of job polarization in japan: evidence from 1980 to 2010, *Applied Economics Letters* **25**(7), 456–460.
- Gaffard, J.-L.: 2018, Towards a non-walrasian macroeconomics, *Revue de l'OFCE* (3), 235–256.
- Gardberg, M., Heyman, F., Norbäck, P.-J. and Persson, L.: 2020, Digitization-based automation and occupational dynamics, *Economics Letters* **189**, 109032.
- Garfinkel, J.: 2020, Gartner survey of more than 3,000 cios reveals that enterprises are entering the third era of it.
- Gasparetto, A. and Scalera, L.: 2019, A brief history of industrial robotics in the 20th century, *Advances in Historical Studies* **8**(1), 24–35.
- Gatti, D. D., Desiderio, S., Gaffeo, E., Cirillo, P. and Gallegati, M.: 2011, *Macroeconomics from the Bottom-up*, Vol. 1, Springer Science & Business Media.
- Georgieff, A. and Milanez, A.: 2021, What happened to jobs at high risk of automation?
- Giuliodori, D. and Stucchi, R.: 2012, Innovation and job creation in a dual labor market: evidence from spain, *Economics of Innovation and New Technology* **21**(8), 801–813.
- Goos, M., Manning, A. and Salomons, A.: 2014, Explaining job polarization: Routine-biased technological change and offshoring, *American economic review* **104**(8), 2509–26.
- Graetz, G. and Michaels, G.: 2018, Robots at work, *Review of Economics and Statistics* **100**(5), 753–768.
- Green, F., Felstead, A., Gallie, D. and Henseke, G.: 2016, Skills and work organisation in britain: a quarter century of change, *Journal for Labour Market Research* **49**(2), 121–132.
- Gunaratne, S. A.: 2001, Paper, printing and the printing press: A horizontally integrative macrohistory analysis, *Gazette (Leiden, Netherlands)* **63**(6), 459–479.
- Hagedorn, M., Karahan, F., Manovskii, I. and Mitman, K.: 2013, Unemployment benefits and unemployment in the great recession: the role of macro effects, *Technical report*, National Bureau of Economic Research.
- Hall, B. H., Lotti, F. and Mairesse, J.: 2009, Innovation and productivity in smes: empirical evidence for italy, *Small business economics* **33**(1), 13–33.

- Harrison, R., Jaumandreu, J., Mairesse, J. and Peters, B.: 2014, Does innovation stimulate employment? a firm-level analysis using comparable micro-data from four european countries, *international Journal of industrial organization* **35**, 29–43.
- Hein, E. and Lavoie, M.: 2019, Post-keynesian economics, *The Elgar Companion to John Maynard Keynes*, Edward Elgar Publishing.
- Hou, J., Huang, C., Licht, G., Mairesse, J., Mohnen, P., Mulkey, B., Peters, B., Wu, Y., Zhao, Y. and Zhen, F.: 2019, Does innovation stimulate employment? evidence from china, france, germany, and the netherlands, *Industrial and Corporate Change* **28**(1), 109–121.
- Howell, D. R.: 2004, *Fighting unemployment: the limits of free market orthodoxy*, Oxford University Press.
- Humlum, A.: 2019, Robot adoption and labor market dynamics, *Princeton University*.
- IFR: 2016, Industrial robots—definition and classification, *World Robotics* pp. 25–34.
- Jo, Y. J.: 2019, Downward nominal wage rigidity in the united states, *Job Market Paper*.
- Johnston, A. C. and Mas, A.: 2018, Potential unemployment insurance duration and labor supply: The individual and market-level response to a benefit cut, *Journal of Political Economy* **126**(6), 2480–2522.
- Kaldor, N.: 1932, A case against technical progress?, *Economica* (36), 180–196.
- Kaldor, N.: 1957, A model of economic growth, *Economic Journal* **December**, 65–94.
- Keynes, J. M.: 1930, Economic prospects of our grandchildren, *Essays in persuasion* pp. 191–202.
- Khan, W. and Fatima, S.: 2016, An assessment of sectoral dynamics and employment shift in indian and chinese economy, *South Asian Survey* **23**(2), 119–134.
- Klenert, D., Fernandez-Macias, E. and Antón Pérez, J. I.: 2020, Do robots really destroy jobs? evidence from europe, *Technical report*, JRC Working Papers Series on Labour, Education and Technology.
- Knight, F. H.: 1921, *Risk, uncertainty and profit*, Vol. 31, Houghton Mifflin.
- Koch, M., Manuylov, I. and Smolka, M.: 2021, Robots and firms, *The Economic Journal* **131**(638), 2553–2584.
- Kristal, T. and Cohen, Y.: 2017, The causes of rising wage inequality: the race between institutions and technology, *Socio-Economic Review* **15**(1), 187–212.

- Kunst, D.: 2019, Deskilling among manufacturing production workers.
- Kurer, T.: 2020, The declining middle: Occupational change, social status, and the populist right, *Comparative Political Studies* **53**(10-11), 1798–1835.
- Kurz, H. D.: 1984, Ricardo and lowe on machinery, *Eastern Economic Journal* **10**(2), 211–229.
- Lachenmaier, S. and Rottmann, H.: 2011, Effects of innovation on employment: A dynamic panel analysis, *International journal of industrial organization* **29**(2), 210–220.
- Laguna, H. and Bianchi, C.: 2020, Firm’s innovation strategies and employment: new evidence from uruguay, *Serie Documentos de Trabajo; 06/20* .
- Lancaster, T. and Chesher, A.: 1983, An econometric analysis of reservation wages, *Econometrica: Journal of the Econometric Society* pp. 1661–1676.
- Le Barbanchon, T., Rathelot, R. and Roulet, A.: 2019, Unemployment insurance and reservation wages: Evidence from administrative data, *Journal of Public Economics* **171**, 1–17.
- Lederer, E.: 1931, *Technischer Fortschritt und Arbeitslosigkeit*, JCB Mohr.
- Lee, K.: 2016, Artificial intelligence, automation, and the economy, *Executive Office of the President of the USA* **20**.
- Leontief, W.: 1979, Is technological unemployment inevitable?, *Challenge* **22**(4), 48–50.
- Leontief, W., Leontief, W. W., Duchin, F. et al.: 1986, *The future impact of automation on workers*, New York: Oxford University Press.
- Lim, J. and Lee, K.: 2019, Employment effect of innovation under different market structures: Findings from korean manufacturing firms, *Technological Forecasting and Social Change* **146**, 606–615.
- Llerena, P. and Lorentz, A.: 2004, Cumulative causation and evolutionary micro-founded technical change: On the determinants of growth rates differences, *Revue Economique* **55**(6), 1191–1214.
- Llerena, P., Lorentz, A., Marengo, L. and Valente, M.: 2014, Knowledge, production and time: An attempted evolutionary model of production, *15th International Conference of the International Joseph A. Schumpeter Society*, Jena, Germany.
- Lorentz, A.: 2016, *Nicholas Kaldor-Faits stylisés, progrès technique et croissance cumulative*, Éditions EMS.
- Lorentz, A. and Savona, M.: 2008, Evolutionary micro-dynamics and changes in the economic structure, *Journal of Evolutionary Economics* **18**(3-4), 389–412.

- Lucchese, M. and Pianta, M.: 2012, Innovation and employment in economic cycles, *Comparative Economic Studies* **54**(2), 341–359.
- MacBride, R.: 1967, *The automated state: computer systems as a new force in society*, Philadelphia: Chilton Book Company.
- Mann, K. and Püttmann, L.: 2018, Benign effects of automation: New evidence from patent texts, *Available at SSRN 2959584* .
- Marcolin, L., Miroudot, S. and Squicciarini, M.: 2016, Gvcs, jobs and routine content of occupations.
- Marx, K.: 1867, *Capital: A Critique of Political Economy. Volume I: The Process of Production of Capital*, Giuseppe Castrovilli.
- Matuzeviciute, K., Butkus, M. and Karaliute, A.: 2017, Do technological innovations affect unemployment? some empirical evidence from european countries, *Economies* **5**(4), 48.
- Mavromaras, K., McGuinness, S., O’Leary, N., Sloane, P. and Wei, Z.: 2013, Job mismatches and labour market outcomes: panel evidence on university graduates, *Economic Record* **89**(286), 382–395.
- McGuinness, S., Pouliakas, K. and Redmond, P.: 2021, Skills-displacing technological change and its impact on jobs: challenging technological alarmism?, *Economics of Innovation and New Technology* pp. 1–23.
- McGuinness, S. and Sloane, P. J.: 2011, Labour market mismatch among uk graduates: An analysis using reflex data, *Economics of Education Review* **30**(1), 130–145.
- Mellacher, P. and Scheuer, T.: 2020, Wage inequality, labor market polarization and skill-biased technological change: An evolutionary (agent-based) approach, *Computational Economics* pp. 1–46.
- Meriküll, J.: 2010, The impact of innovation on employment: firm-and industry-level evidence from a catching-up economy, *Eastern European Economics* **48**(2), 25–38.
- Metcalf, J. S.: 1994, Competition, fisher’s principle and increasing returns in the selection process, *Journal of Evolutionary Economics* **4**, 327–346.
- Mill, J. S.: 1848, Principles of political economy with some of their applications, *Social Philosophy* .
- Mitra, A.: 2020, Technological progress, innovation and employment: firm-level evidence from india’s manufacturing sector, *Innovation and Development* **10**(3), 451–465.
- Mocan, H. N.: 1999, Structural unemployment, cyclical unemployment, and income inequality, *Review of Economics and Statistics* **81**(1), 122–134.

- Montobbio, F., Staccioli, J., Virgillito, M. E. and Vivarelli, M.: 2022, Robots and the origin of their labour-saving impact, *Technological Forecasting and Social Change* **174**, 121122.
- Náplava, R.: 2019, Changing structure of employment in europe: Polarization issue, *Review of Economic Perspectives* **19**(4), 307–318.
- Nataraj, S., Perez-Arce, F., Kumar, K. B. and Srinivasan, S. V.: 2014, The impact of labor market regulation on employment in low-income countries: A meta-analysis, *Journal of Economic Surveys* **28**(3), 551–572.
- Nedelkoska, L. and Quintini, G.: 2018, Automation, skills use and training.
- Nelson, R. and Winter, S.: 1982, *An Evolutionary Theory of Economic Change*, Harvard University Press.
- Nickell, S.: 1981, Biases in dynamic models with fixed effects, *Econometrica: Journal of the econometric society* pp. 1417–1426.
- Nickell, S., Nunziata, L. and Ochel, W.: 2005, Unemployment in the oecd since the 1960s. what do we know?, *The Economic Journal* **115**(500), 1–27.
- Nunziata, L. et al.: 2002, *Unemployment, labour market institutions and shocks*, Nuffield College Oxford, UK.
- Nuvolari, A. and Russo, E.: 2019, Technical progress and structural change: a long-term view, *Technical report*, LEM Working Paper Series.
- Öcalan-Özel, S. and Pénin, J.: 2016, Exclusive or open? an economic analysis of university intellectual property patenting and licensing strategies, *Journal of Innovation Economics* (3), 133–153.
- OECD: 1994, *The OECD jobs study: facts, analysis, strategies*, Organisation for Economic Co-operation and Development.
- OECD: 2004, Employment protection regulation and labour market performance.
- OECD, B. P.: 2017, Data sources, *OECD Science, Technology and Industry Scoreboard 2017*.
- Oesch, D. and Rennwald, L.: 2018, Electoral competition in europe’s new tripolar political space: Class voting for the left, centre-right and radical right, *European journal of political research* **57**(4), 783–807.
- Okumu, I. M., Bbaale, E. and Guloba, M. M.: 2019, Innovation and employment growth: Evidence from manufacturing firms in africa, *Journal of Innovation and Entrepreneurship* **8**(1), 1–27.

- Okun, A. M.: 1963, *Potential GNP: its measurement and significance*, Cowles Foundation for Research in Economics at Yale University.
- Pajarinen, M., Rouvinen, P., Ekeland, A. et al.: 2015, Computerization threatens one-third of finnish and norwegian employment, *Etla Brief* **34**, 1–8.
- Pak, M. and Schwellnus, C.: 2019, Labour share developments over the past two decades: The role of public policies.
- Pellizzari, M. and Fichen, A.: 2017, A new measure of skill mismatch: theory and evidence from piaac, *IZA Journal of Labor Economics* **6**(1), 1.
- Peluffo, A.: 2020, Does innovation affect the demand for employment and skilled labor?, *Serie Documentos de Trabajo; 13/20* .
- Peters, B., Dachs, B., Dünser, M., Hud, M., Köhler, C. and Rammer, C.: 2014, Firm growth, innovation and the business cycle: Background report for the 2014 competitiveness report, *Technical report*, ZEW Gutachten/Forschungsberichte.
- Pontusson, J. and Weisstanner, D.: 2016, The political economy of compensatory redistribution: Unemployment, inequality and policy choice, *Technical report*, LIS Working Paper Series.
- Pretty, J. N.: 1991, Farmers' extension practice and technology adaptation: Agricultural revolution in 17–19th century britain, *Agriculture and Human Values* **8**(1), 132–148.
- Ricardo, D.: 1821, *On the principles of political economy*, J. Murray London.
- Rifkin, J.: 1995, *The end of work: The decline of the global labor force and the dawn of the post-market era.*, ERIC.
- Rifkin, J.: 2014, *The zero marginal cost society: The internet of things, the collaborative commons, and the eclipse of capitalism*, St. Martin's Press.
- Rothstein, J.: 2011, Unemployment insurance and job search in the great recession, *Technical report*, National Bureau of Economic Research.
- Salvatori, A.: 2018, The anatomy of job polarisation in the uk, *Journal for labour market research* **52**(1), 8.
- Sauvy, A.: 1980, La machine et le chômage. le progrès technique et l'emploi, *Paris, Dunod* .
- Say, J.-B.: 1803, A treatise on political economy, 3 rd us edition, *Philadelphia: John Grigg* .
- Scarpetta, S.: 1996, Assessing the role of labour market policies and institutional settings on unemployment: A cross-country study, *OECD Economic studies* **26**(1), 43–98.

- Scholl, K. and Hanson, R.: 2020, Testing the automation revolution hypothesis, *Economics Letters* **193**, 109287.
- Schumpeter, J. A.: 1954, History of economic analysis, ed, *EB Schumpeter. London* .
- Siebert, H.: 1997, Labor market rigidities: at the root of unemployment in europe, *Journal of Economic perspectives* **11**(3), 37–54.
- Smith, A.: 1776, An inquiry into the nature and causes of the wealth of nations: Volume one, London: printed for W. Strahan; and T. Cadell, 1776.
- Smith, T. B.: 1999, Assistance and repression: Rural exodus, vagabondage and social crisis in france, 1880-1914, *Journal of Social History* pp. 821–846.
- Sparreboom, T. and Tarvid, A.: 2016, Imbalanced job polarization and skills mismatch in europe, *Journal for Labour Market Research* **49**(1), 15–42.
- Stam, E. and Wennberg, K.: 2009, The roles of r&d in new firm growth, *Small business economics* **33**(1), 77–89.
- Suetonius, S.: 2019, *The lives of the twelve Caesars*, BoD–Books on Demand.
- Thaler, R. H. and Ganser, L.: 2015, Misbehaving: The making of behavioral economics.
- Triguero, A., Córcoles, D. and Cuerva, M. C.: 2014, Persistence of innovation and firm’s growth: evidence from a panel of sme and large spanish manufacturing firms, *Small business economics* **43**(4), 787–804.
- Van Roy, V., Vertesy, D. and Vivarelli, M.: 2015, Innovation and employment in patenting firms: empirical evidence from europe.
- Van Roy, V., Vivarelli, M. and Vertesy, D.: 2015, *The employment impact of innovation: evidence from European patenting companies*, Vita e Pensiero.
- Vermeulen, B., Kesselhut, J., Pyka, A. and Saviotti, P. P.: 2018, The impact of automation on employment: just the usual structural change?, *Sustainability* **10**(5), 1661.
- Vivarelli, M.: 2014, Innovation, employment and skills in advanced and developing countries: A survey of economic literature, *Journal of Economic Issues* **48**(1), 123–154.
- Wagniart, J.-F.: 1998, Les migrations des pauvres en france à la fin du xix e siècle: le vagabondage ou la solitude des voyages incertains, *Genèses* pp. 30–52.
- Wang, Y.-Q. and Tsui, K.-Y.: 2000, Polarization orderings and new classes of polarization indices, *Journal of Public Economic Theory* **2**(3), 349–363.
- Warwick, K.: 2013, *Artificial intelligence: the basics*, Routledge.

- Webb, M.: 2019, The impact of artificial intelligence on the labor market, *Available at SSRN 3482150* .
- Winkler, H.: 2019, The effect of income inequality on political polarization: Evidence from european regions, 2002–2014, *Economics & Politics* **31**(2), 137–162.
- WIPO, W.: 2019, Technology trends 2019: Artificial intelligence, *Geneva: World Intellectual Property Organization* .
- Zamalloa, I., Kojcev, R., Hernández, A., Muguruza, I., Usategui, L., Bilbao, A. and Mayoral, V.: 2017, Dissecting robotics-historical overview and future perspectives, *arXiv preprint arXiv:1704.08617* .
- Zhu, C., Qiu, Z. and Liu, F.: 2021, Does innovation stimulate employment? evidence from china, *Economic Modelling* **94**, 1007–1017.
- Zhu, R. and Chen, L.: 2016, Overeducation, overskilling and mental well-being, *The BE Journal of Economic Analysis & Policy* **16**(4).
- Zimmermann, V.: 2009, The impact of innovation on employment in small and medium enterprises with different growth rates, *Jahrbücher für Nationalökonomie und Statistik* **229**(2-3), 313–326.

Résumé en français

Cette thèse s'articule autour de trois chapitres couvrant les aspects empiriques, théoriques et politiques de l'impact de la robotique et de l'intelligence artificielle sur l'emploi et les inégalités. Le lien entre l'innovation de procédé et emploi est ancien et complexe : celle-ci est un vecteur de développement économique, mais aussi une source d'instabilité qui peut se matérialiser en préoccupations politiques. Par exemple, dans la Rome antique, l'empereur Vespasien aurait refusé d'utiliser une machine permettant de transformer de lourdes colonnes vers le site de construction du Capitole, de peur que cette innovation mette au chômage les ouvriers en charge du transport. Plus récemment, la question du remplacement de l'Homme par la machine a été au cœur de la campagne présidentielle française de 2017 de Benoît Hamon, candidat du parti socialiste, qui proposait l'instauration d'un revenu universel afin d'assurer à tous un filet de sécurité face à la menace que ferait peser l'automatisation sur les emplois. Aux États-Unis, le candidat à la présidence Andrew Yang a mené campagne sur une idée similaire, baptisée « Freedom dividend », visant à garantir à chaque américain un revenu inconditionnel d'un montant de mille dollars par mois.

La science économique s'est penchée sur la question depuis l'époque des fondateurs de l'économie politique. Pour Adam Smith (1776), les destructions d'emplois liés à la substitution de l'homme par la machine sont compensées par une augmentation de la productivité, aboutissant à une baisse du coût unitaire de production, à une baisse des prix, à une hausse de la demande engendrant à son tour une hausse de la production et donc de la demande de travail. Jean-Baptiste Say (1803) adopte une position similaire, soulignant que les gains de productivité permettent d'améliorer la compétitivité des manufactures nationales et donc, in fine, de préserver des emplois. Ricardo (1821) adopte une opinion plus nuancée : si le financement des machines provient du « fond de salaire », autrement dit de la masse salariale, alors il est probable que l'automatisation affectera négativement l'emploi. Pour Marx (1867), la mécanisation permet aux capitalistes de renforcer leur pouvoir aux dépens de la classe ouvrière. Marx critique les mécanismes de compensation, en soulignant par exemple que la quantité de travail incorporé dans une machine est nécessairement inférieure à la quantité de travail économisée par l'entreprise acquéreuse, sinon le capitaliste possédant cette dernière n'aurait aucun intérêt financier à investir dans cette machine. Marx reconnaît néanmoins que les gains de productivité permettent de faire baisser les prix et donc de booster l'activité. Finalement, Keynes (1930) adopte une vision positive en soulignant que, même si à court terme le progrès technique peut engendrer une hausse du chômage, à long terme

il bénéficie au développement de l'activité et permettrait à l'humanité de résoudre son « problème économique », i.e. la capacité à produire suffisamment pour répondre aux besoins de l'ensemble des êtres humains. Keynes propose de lutter contre le chômage technologique en réduisant le temps de travail, prophétisant sur une semaine de 15h. Même si cette prédiction ne s'est pas encore réalisée, Keynes avait remarquablement bien capturé les deux grandes tendances du milieu du XXe et début du XXIe siècle, caractérisées par une baisse du temps de travail (dans les pays de l'OCDE, celui-ci a baissé de 11% entre 1970 et 2019) et une hausse du PIB réel par habitant (de 145% dans les pays de l'OCDE sur la même période).

Contrairement aux économistes du XVIIIe, XIXe et de la première moitié du XXe siècle, les économistes contemporains ont la possibilité de profiter de l'augmentation de la quantité et de la qualité des données, des progrès des techniques économétriques ainsi que de l'essor de l'informatique pour tester statistiquement l'efficacité de ces mécanismes de compensation. Calvino et Virgillito (2018) recensent deux types de mécanismes de compensation : classiques et keynésien-schumpetériens. Le premier type englobe quatre mécanismes différents : une augmentation de la demande de travail dans les entreprises qui produisent les innovations de procédé, une baisse des prix due aux gains de productivité qui entraîne une augmentation de la demande globale, une diminution des salaires qui augmente à son tour la demande de travail, et une augmentation de l'investissement financé par la baisse du coût unitaire de production, déclenchant ainsi de nouvelles embauches pour accroître la production. La deuxième catégorie comprend deux mécanismes : une augmentation des salaires due aux gains de productivité pour les travailleurs qui restent dans l'entreprise, ce qui stimule la demande de main-d'œuvre ; et la création de nouveaux produits qui ouvrent de nouveaux marchés et augmentent ainsi la main-d'œuvre.

Une première littérature cherche à estimer l'impact sur l'emploi des innovations de procédé, c'est-à-dire les innovations visant à améliorer l'efficacité des processus de production. Au niveau de la firme, les résultats sont ambivalents : alors que certains auteurs observent un effet positif sur l'emploi (Zimmermann (2009), Lachenmaier et Rottmann (2011), Giuliadori et Stucchi (2012), Triguero et al. (2014), Peluffo (2020), Laguna et Bianchini (2020), Okumu et al. (2019), Zhu et al. (2021)), d'autres concluent à un effet négatif (Peters et al. (2014)), Aboal et al. (2015), Evangelista et Vezzani (2012)) ou à une absence d'impact statistiquement significatif (Hall et al. (2009), Benavente et Lauterbach (2008), Bianchini et Pellegrino (2019), Baensch et al. (2019), Cirera et Sabetti (2019), Mitra (2020), Castillo et al. (2014), Barbieri et al. (2019), Lim et Lee (2019), De Elejalde et al. (2015), Benavente et Lauterbach (2008), Hou et al. (2019)). Un inconvénient majeur de l'analyse au niveau de la firme est qu'elle ne permet pas d'étudier l'efficacité des mécanismes de compensation. Dans le cas des études concluant que les innovations de procédé ont un effet positif sur l'emploi, il n'y a aucun moyen de savoir si ces créations d'emplois ne se font pas au détriment des pertes d'emplois chez leurs concurrents moins innovants. Inversement, les études qui concluent à un effet négatif ne permettent pas d'observer les effets sur, par exemple, les entreprises qui produisent ces innovations de procédé.

Pour résoudre ce problème, certains auteurs ont porté leur cadre d'analyse au niveau sectoriel. Antonucci (2007) utilise des données sur 22 secteurs manufacturiers dans 10 pays européens couvrant la période 1994-2001 et constate que les innovations de procédé a un impact négatif sur l'emploi, mais cet effet est compensé par une augmentation de la demande générée par les gains de productivité. Ce résultat n'est pas partagé par Bogliacino et Pianta (2010), qui analysent un échantillon de 38 industries dans huit pays européens (Allemagne, Espagne, France, Italie, Pays-Bas, Portugal, Royaume-Uni et Norvège) sur la période 1996-2004, et mettent en évidence une corrélation négative entre les dépenses en machines innovantes par employé et les heures travaillées. Une corrélation négative entre innovation de procédé et emploi est également observée par Meriküll (2010), Lucchese et Pianta (2012) (uniquement durant les périodes récessions), et Dosi et al. (2021).

Enfin, un dernier niveau d'analyse se situe au niveau macroéconomique. Feldmann (2013) analyse un échantillon de 21 pays industriels sur la période 1985 à 2009 et utilise les brevets triadiques (c'est-à-dire, les brevets déposés auprès des trois principaux offices de brevets : l'Office européen des brevets, le Bureau américain des brevets et des marques de commerce et l'Office japonais des brevets) comme indicateur de l'évolution technologique. À noter que les brevets ne soient pas un indicateur parfait des innovations de procédé étant donné que l'innovation de produit peut également être brevetée et que les stratégies en matière de brevets peuvent gonfler artificiellement le nombre de brevets. Feldmann (2013) trouve une corrélation positive et significative entre les brevets triadiques et le taux de chômage. Ce résultat n'est pas partagé par Matuzeviciute et al. (2017) qui utilisent également les brevets triadiques comme proxy, mais ne trouvent aucune relation significative avec le taux de chômage dans 25 pays européens sur la période 2000 - 2012. Evangelista et al. (2014) utilisent un indicateur basé sur les technologies de l'information de la communication (TIC) et trouvent une corrélation positive avec l'emploi ; un résultat partagé par Marcolin et al. (2016) pour 28 pays de l'OCDE sur la période 2000-2011 en utilisant comme variable les brevets au lieu de l'intensité des TIC. L'une des principales limites de cette littérature macroéconomique est que, en raison d'un manque de données, les innovations de procédé ne sont pas clairement définies, ce qui conduit les auteurs à utiliser des proxys. Il est donc difficile de savoir si les résultats obtenus permettent de conclure à l'effet positif ou négatif des innovations de procédé sur l'emploi.

Parmi les innovations de procédé, deux technologies attirent particulièrement l'attention : la robotique et l'intelligence artificielle. Les machines robotisées se distinguent des autres machines traditionnelles par leur degré élevé d'autonomie, qui permet de réduire considérablement la quantité de main-d'œuvre requise dans le processus de production. Bien qu'il ne soit pas le seul facteur, le déploiement de robots industriels dans le secteur manufacturier de nombreux pays développés est un bon exemple de ce phénomène, qui s'illustre par une diminution du nombre d'emplois malgré une augmentation de la production. Selon les données de la base de données STAN de l'OCDE et de la Fédération internationale de robotique, la valeur ajoutée produite par le secteur manufacturier français a augmenté de 48%, l'emploi a diminué de 28% et le stock de

robots industriels a triplé entre 1971 et 2017. Au cours de la même période, la valeur ajoutée produite par le secteur manufacturier allemand a augmenté de 54%, l'emploi a diminué de 12% et le stock de robots industriels a été multiplié par 2,9. Des tendances similaires peuvent être observées en Italie, en Espagne, au Royaume-Uni et d'autres pays de l'OCDE. L'intelligence artificielle se distingue de la robotique car elle rend possible l'automatisation des tâches cognitives qui sont, par essence, dématérialisées. En ce sens, elle est similaire aux algorithmes classiques, mais diffère en ce sens que les algorithmes d'apprentissage automatique sont évolutifs et donc adaptatifs. Ils permettent d'automatiser des tâches à forte intensité de connaissances tacites, comme la conduite d'une voiture, car il n'est pas nécessaire pour les programmeurs de coder toutes les connaissances nécessaires à l'exécution d'une tâche. Ce changement de paradigme, qui repousse la frontière des tâches automatisables, soulève des questions quant à une éventuelle accélération de l'automatisation dans le secteur des services, qui a été relativement épargné jusqu'à présent.

C'est à partir de ce constat qu'a émergé une nouvelle littérature s'intéressant à l'impact de la robotique, de l'IA et, de manière plus large, de l'automatisation sur l'emploi. L'étude de Frey et Osborne (2017), qui recourent à l'évaluation d'experts et déterminent que 47% des emplois américains sont menacés par l'automatisation, a ouvert la voie à des études à la méthodologie analogue donnant un pourcentage d'emplois menacés par l'automatisation de 25% à Singapour (Fuei (2017)), de 35% en Finlande et de 33% en Norvège (Pajarinen et al. (2015)), de 42% en Allemagne et de 35% au Royaume-Uni (Frey et Osborne (2014)). Arntz et al. (2016) analysent le risque d'automatisation non plus au niveau de la profession, mais au niveau des tâches, et trouvent que seulement 9% des emplois risquent d'être automatisés aux États-Unis ; un chiffre bien inférieur aux 47% trouvés par Frey et Osborne (2014). Dengler et Matthes (2018) suivent une méthodologie similaire et constatent que 15% des emplois allemands risquent d'être prochainement automatisés. Toutes ces études ont le mérite de quantifier, avec un éventail de résultats relativement large, la proportion d'emplois qui pourraient être automatisés, mais elles souffrent de deux inconvénients majeurs. Le premier est qu'elles reposent sur l'opinion d'experts, qui sont vulnérables à tout un ensemble de biais cognitifs bien connus en économie comportementale (Thaler et Ganser (2015)), et deuxièmement, elles traitent l'automatisation de manière binaire et homogène, sans distinguer les différentes sources d'automatisation.

Pour pallier à ces problèmes, d'autres auteurs utilisent des données sur les robots et l'IA pour les régresser directement sur la variable d'intérêt. Pour les analyses au niveau de la firme, les résultats sur l'emploi sont relativement unanimes et mettent en valeur un effet positif des robots (Koch et al. (2021), Humlum (2019), Bessen et al. (2020), Acemoglu et al. (2020), Domini et al. (2021), Dixon et al. (2021), Aghion et al. (2020)), alors que les effets sur les salaires sont mitigés : Acemoglu et al. (2020) et Humlum (2019) observent un impact positif, tandis que Koch et al. (2021), Bessen et al. (2020) et Aghion et al. (2020) n'observent pas d'effets statistiquement significatifs.

Au niveau sectoriel, les effets sur l'emploi sont moins clairs qu'au niveau de la firme. Klenert et al. (2020) et Aghion et al. (2020) trouvent un effet positif des robots

industriels sur l'emploi total, Compagnucci et al. (2019) et Carbonero et al. (2020) observent un impact négatif ; tandis que Graetz et Michaels (2018) et De Vries et al. (2020) ne trouvent aucun effet statistiquement significatif. Du côté des salaires, Graetz et Michaels (2018) et Aghion et al. (2020) n'observent pas d'impact statistiquement significatif, tandis que Blanas et al. (2019) mettent en évidence un effet positif sur les salaires des travailleurs hautement qualifiés.

Enfin, les études au niveau régional peinent également à trouver un consensus. En ce qui concerne l'emploi, Mann et Püttmann (2018) constatent que les brevets en lien avec des technologies d'automatisation sont positivement corrélés avec l'emploi total. Chiacchio et al. (2018) et Acemoglu et Restrepo (2020) utilisent des données sur le stock de robots industriels et observent un impact négatif sur l'emploi. Dauth et al. (2021) fournissent une conclusion plus nuancée, montrant que les robots ont un impact positif sur l'emploi dans le secteur et les services et négatif dans le secteur manufacturier ; tandis que Caselli et al. (2021) et Dottori (2021) ne trouvent aucun effet statistiquement significatif sur l'emploi total. Du côté des salaires, Acemoglu et Restrepo (2020) trouvent un impact négatif, tandis que Chiacchio et al. (2018) observent qu'il n'y a pas d'effet.

Une limite de la littérature existante est que les études macroéconomiques sont rares et que celles existantes n'aboutissent pas à un consensus. Pourtant, c'est ce niveau d'analyse qui importe aux décideurs politiques : savoir si les robots détruisent des emplois au niveau de certaines entreprises et de certains secteurs est une information utile, mais elle n'aura pas la même importance pour la politique économique si l'effet net sur l'emploi global est positif ou négatif. La littérature manque également de travaux macroéconomiques sur l'impact de l'IA sur l'emploi, en raison notamment de la difficulté à obtenir des données quantitatives.

Le premier chapitre de cette thèse est un exercice empirique qui vient compléter cette revue de littérature, et vise à répondre à une question très simple : il y a-t-il des preuves statistiques que la robotique et l'IA peuvent augmenter le chômage ? Pour y répondre, nous menons une analyse économétrique à partir de données macroéconomiques couvrant 33 pays de l'OCDE sur la période 2005 – 2017. Les données sur les robots proviennent de la base de données de la fédération internationale de robotique, et les données sur l'IA sont collectées à partir de brevets en appliquant la méthodologie développée par l'organisation mondiale de la propriété intellectuelle. Nos résultats mettent en valeur une corrélation positive entre les robots, l'IA et le taux de chômage : une augmentation de 10% du stock de robots industriels est associée à une hausse 0.42 point du taux de chômage, et une augmentation de 10% de notre variable AI est associée à une hausse de 0.27 point du taux de chômage deux ans plus tard.

Nous poursuivons l'analyse en régressant les variables robots et IA sur les taux de chômage différenciés par niveau d'étude, et les résultats indiquent une forte hétérogénéité entre les différents groupes. Par exemple, l'effet des robots est le plus fort sur le taux de chômage des personnes ayant un niveau d'étude compris entre le bac et bac +2, suivi de près par le taux de chômage des personnes avec un niveau d'étude inférieur au bac, puis l'effet le plus faible est sur le taux de chômage des personnes avec un niveau

d'étude supérieur ou égal à bac +2. Ce premier résultat vient supporter la thèse d'une polarisation du marché du travail induite par le progrès technique. Finalement, nous poursuivons la désagrégation du taux de chômage en analysant les effets de ces technologies sur les taux de chômage par niveau d'étude et tranches d'âge. Là encore, les résultats mettent en relief une forte hétérogénéité au sein d'un groupe ayant le même niveau d'éducation. Par exemple, la corrélation positive entre les robots et le taux de chômage des individus ayant un niveau d'étude inférieur au bac porte exclusivement sur les jeunes (25 – 34 ans), et aucun effet statiquement significatif n'est relevé sur les autres tranches d'âge ayant un niveau d'étude similaire. Ces résultats contribuent à la littérature en apportant des preuves supplémentaires de l'effet potentiellement néfaste de l'IA et des robots sur l'emploi au niveau agrégé, tout en mettant en valeur des effets qui s'avèrent être positifs selon le niveau d'éducation et de l'âge des travailleurs.

Dans le second chapitre, nous nous intéressons aux mécanismes théoriques expliquant la polarisation du marché du travail. Ce phénomène, observé dans de nombreux pays développés, se caractérise par une hausse de la part de l'emploi et des salaires des travailleurs faiblement et hautement qualifiés, et une baisse de la part de l'emploi et des salaires des travailleurs moyennement qualifiés. La littérature théorique expliquant le lien entre l'automatisation et la polarisation du marché du travail repose sur l'hypothèse d'un changement technique biaisé en défaveur des tâches routinières (routine-biased technical change hypothesis (RBTC)), qui postule que l'automatisation cible principalement les tâches routinières réalisées majoritairement par des travailleurs moyennement qualifiés. L'hypothèse est à priori intuitive, mais tend également à conditionner le résultat obtenu : si l'on postule que l'automatisation cible les tâches routinières, elles-mêmes exécutées par des travailleurs moyennement qualifiés, alors on postule indirectement que l'automatisation cible prioritairement les travailleurs moyennement qualifiés, et la dynamique du modèle ne peut aboutir qu'à une polarisation du marché du travail. Dès lors, on peut s'interroger sur la pertinence de cette hypothèse, et sur la possibilité de générer une polarisation du marché du travail à partir de postulats moins restrictifs.

Pour répondre à cette question, nous avons créé et programmé un modèle multi-agents nous permettant de simuler l'effet de l'automatisation sur l'emploi et les salaires. Ce modèle est peuplé d'agents hétérogènes, dotés chacun de différentes compétences qu'ils améliorent au gré de leurs activités professionnelles, et qui se détériorent durant les périodes d'inactivité. Évoluant dans un environnement incertain empêchant toute résolution d'un programme de maximisation intertemporelle, les agents du modèle adoptent des comportements adaptatifs simples leur permettant de remplir un objectif : augmenter son profit pour les firmes, et augmenter son salaire pour les travailleurs. Pour produire une unité d'un bien, les firmes combinent différents types d'emplois en suivant une fonction de production de type Leontief. Les agents ne peuvent candidater à un emploi que si leurs compétences correspondent à celles requises pour l'emploi en question. Le modèle comprend six compétences différentes notées entre zéro et sept, cette dernière valeur correspondant au niveau maximum d'expertise. Les salaires et compétences requises pour chacun de ces métiers sont initialisés à partir de données américaines obtenues en agrégeant les informations fournies par la base de données

O*Net et le Bureau of Labor Statistics. Enfin, les entreprises cherchent à améliorer l'efficacité de leur processus de production en embauchant des ingénieurs qui mènent des activités de recherche et développement visant à améliorer le niveau de productivité incorporé dans le capital et d'automatiser certaines tâches. La structure du modèle est représentée par le schéma suivant :

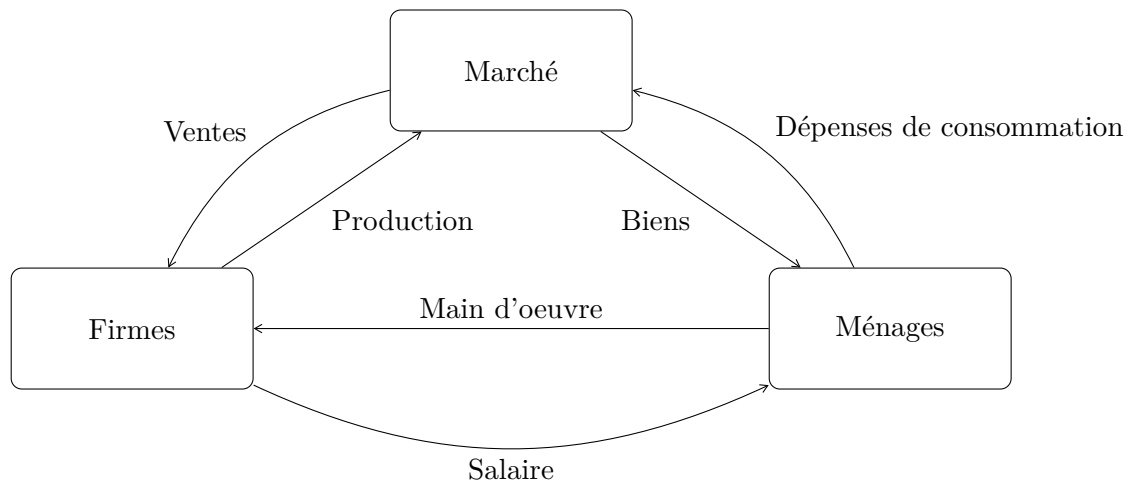


Figure 3.7: Structure du modèle du Chapitre 2

Dans un premier temps, nous testons un scénario sans automatisation. Le modèle reproduit certains faits stylisés, tels qu'une croissance endogène du PIB réel, une hétérogénéité dans le niveau de productivité des firmes ou encore une corrélation positive entre le taux de chômage et les inégalités de revenu. Ce premier scénario aboutit à une polarisation des salaires, liée notamment à un marché du travail compétitif qui offre un salaire plus élevé aux compétences les plus rares; mais la distribution de la demande de compétences, elle, ne se polarise pas. Empiriquement, la polarisation du marché du travail est caractérisée par une déformation de la distribution des salaires et du niveau de qualification des emplois ; ce premier scénario n'est donc pas à même de reproduire pleinement ce phénomène.

Dans un second temps, nous testons un scénario avec automatisation. Contrairement à l'explication fondée sur le RBTC, nous ne posons aucune hypothèse quant aux types de tâches ciblées par l'automatisation ; nous postulons simplement que les firmes automatisent certaines tâches pour réduire leurs coûts de production. Ce scénario reproduit les faits stylisés listés précédemment, ainsi que le caractère erratique de l'automatisation, qui se matérialise par une automatisation irrégulière, qui se caractérise par des pics. En conservant les valeurs des paramètres identiques à celles du premier scénario, l'introduction de l'automatisation aboutit à une polarisation des salaires et de la distribution de la demande de compétences. Ce résultat s'explique par l'aspect déqualifiant de l'automatisation : quand une tâche est automatisée, les compétences nécessaires à sa réalisation ne sont plus requises, et par conséquent les travailleurs ne développent plus cette compétence. Le temps de travail économisé est réalloué sur

les autres tâches non automatisées, permettant d'accélérer le développement des compétences liées à ces tâches, car le mécanisme d'apprentissage se fonde sur des effets d'expérience. Il en résulte que les travailleurs sont de plus en plus spécialisés mais, en moyenne, moins qualifiés car leurs compétences qui ne sont plus mobilisées se sont dépréciées, aboutissant à une perte de polyvalence. Dans ce modèle, l'automatisation génère du chômage technologique, avec une hausse de 13.6 points du taux de chômage entre les deux scénarios. Enfin, nous testons un scénario intermédiaire dans lequel l'automatisation est plus lente, et montrons que le degré de polarisation du marché du travail et l'ampleur du chômage technologique sont proportionnels au rythme de l'automatisation.

Après avoir étudié le lien empirique entre robotique, intelligence artificielle et chômage (Chapitre 1), et après avoir explicité les mécanismes théoriques pouvant expliquer l'émergence du chômage technologique et de la polarisation du marché du travail (Chapitre 2), nous nous intéressons dans le Chapitre 3 aux politiques permettant de contrecarrer les effets négatifs de l'automatisation. Le progrès technologique est, au regard de l'histoire, vecteur de développement économique ; il ne s'agit donc nullement de chercher à le freiner, mais plutôt à déployer des moyens permettant d'accompagner les mutations qu'il engendre sur le marché du travail. Pour ce faire, nous étendons le modèle développé dans le Chapitre 2 afin de tester quatre politiques : une régulation des licenciements et des démissions (FRR) visant à stabiliser le marché du travail, un salaire minimum (MW), un système d'assurance chômage (UB) et un système de formation (TP).

Les résultats indiquent que ces politiques ont des effets hétérogènes qui sont parfois complémentaires, à l'image de la régulation des contrats de travail et de l'instauration d'un salaire minimum qui permettent de réduire la polarisation des salaires, tandis que le système de formation atténue la polarisation des compétences. À l'inverse, les effets sont parfois opposés, à l'image de la régulation des contrats de travail qui impacte négativement la productivité, alors que le SMIC et l'assurance chômage l'augmentent. Le tableau ci-dessous synthétise les résultats obtenus :

	FRR	MW	UB	TP
Indice de polarisation des compétences	/	/	/	-
Indice de polarisation de salaires	-	-	/	/
Indice de Gini	-	-	/	-
Indice de compétences médian	+	/	/	+
Salaire réel médian	+	+	/	/
Taux de chômage	-	+	+	-
Productivité du travail	-	+	+	/
Part des tâches automatisées	/	-	-	-
PIB réel	+	/	/	/

Table 3.23: *Synthèse des résultats du Chapitre 3 - "/" = Pas d'effet statistiquement significatif*

Certaines politiques testées semblent efficaces pour réduire la polarisation de la distribution des salaires ou de la demande de compétences, mais aucune ne parvient à agir sur les deux de manière simultanée. Dans un dernier scénario, nous testons conjointement les quatre politiques, et les résultats sont très positifs avec une réduction de la polarisation des salaires et des compétences, une hausse du salaire réel médian, du niveau général de qualifications, de la productivité du travail et du PIB réel qui engendre une baisse du taux de chômage. Ces résultats montrent qu’il est possible de lutter efficacement contre le chômage technologique et la polarisation du marché du travail en instaurant des politiques appropriées ; et que les décideurs politiques doivent penser les politiques publiques dans leur ensemble et non de manière séparée, au risque de passer à côté de potentielles synergies positives.

Nous avons montré dans le chapitre 1 qu’il y avait un lien statistiquement significatif entre robots, IA et chômage. Dans le chapitre 2, nous avons explicité les mécanismes théoriques pouvant expliquer l’émergence du chômage technologique et de la polarisation du marché du travail. Dans le chapitre 3, nous avons testé différentes politiques visant à lutter contre ces deux phénomènes, et montré qu’il était possible de les atténuer grâce à l’intervention publique. Ces trois chapitres contribuent à la littérature sur le lien entre progrès technologique, emploi et inégalités en couvrant à la fois les aspects empiriques, théoriques et politiques du sujet ; mais comportent également plusieurs limites.

Dans le chapitre 1, les effets des robots sur le chômage sont aisément interprétables, tandis que les effets de l’intelligence artificielle sont moins robustes et plus difficilement interprétables. Alors que pour la variable « robots » nous utilisons des données sur le stock effectivement déployé dans les entreprises, pour la variable « IA » nous construisons les données à partir de la collecte de brevets en lien avec l’intelligence artificielle. Dès lors, nos données sur l’IA sont davantage une mesure de l’innovation dans ce domaine qu’une mesure du stock de logiciels et algorithmes effectivement utilisés dans les processus de production, une différence pouvant expliquer le manque de robustesse des résultats liés à cette variable. La meilleure solution serait de disposer, à l’image des robots industriels, de données sur le stock de capital lié à l’IA ; mais de telles données ne sont pas, à notre connaissance, disponibles à l’échelle macroéconomique.

Dans le chapitre 2, les entreprises se font concurrence sur un marché unique en produisant un bien homogène. Bien que l’hétérogénéité des salaires et des compétences requises pour un même type de métier émerge déjà des propriétés du modèle, il serait intéressant d’observer si la dynamique de ces deux variables diffère dans le cadre d’un modèle multisectoriel. En effet, la consultation des données du Bureau Of Labor Statistics nous indique qu’il y a de grandes disparités de salaires pour un même emploi entre différents secteurs ; et il serait intéressant de pouvoir étudier si l’hétérogénéité des trajectoires technologiques entre les secteurs peut contribuer à expliquer ces différences. Enfin, une approche sectorielle permettrait, grâce à une matrice input-output, d’inclure les consommations intermédiaires dans la fonction de production afin de pouvoir étudier les mécanismes de compensation intersectoriels. En effet, l’automatisation dans un secteur d’activité A permet de réaliser des gains de productivité et donc de faire baisser le coût

unitaire de production et donc les prix. Cette baisse de prix se répercute également sur les autres secteurs qui utilisent l'output du secteur A comme input dans leur fonction de production, permettant aux entreprises de ces secteurs d'être plus compétitives, de gagner des parts de marché, d'augmenter leur production et, in fine, de créer des emplois. Enfin, l'ouverture du modèle aux échanges internationaux pourrait atténuer l'ampleur du chômage technologique et de la polarisation du marché du travail générés par le modèle. En faisant baisser les coûts, l'automatisation permet aux entreprises nationales de réduire l'écart de compétitivité-prix face aux pays où la main d'œuvre est peu chère, permettant in fine de préserver des emplois.

La dernière limite soulevée pour le chapitre 2 peut également s'appliquer au chapitre 3. En effet, la combinaison des quatre politiques testées a pour effet de faire augmenter le salaire réel médian de 35%, ce qui ne pose pas de problème dans une économie fermée, mais pourrait s'avérer pénalisante pour la compétitivité des entreprises nationales dans le cas d'un modèle en économie ouverte. Dès lors, l'effet net sur l'emploi, légèrement positif dans le cas d'une économie fermée, serait beaucoup plus incertain et pourrait même devenir négatif.

Ces limites sont autant de pistes de recherche futures. La combinaison de données plus précises et de modèles plus développés constitue un terrain de recherche fertile pour analyser plus en détail les transformations du marché du travail induites par le progrès technologique.

Robotics and artificial intelligence: what impacts on employment and inequalities?

Résumé

Cette thèse étudie l'impact des robots et de l'intelligence artificielle (IA) sur l'emploi et les inégalités sous ses aspects empiriques, théoriques et politiques. Dans le premier chapitre, nous menons une étude économétrique sur 33 pays de l'OCDE, et observons que les robots et l'IA augmentent le taux de chômage, tout en soulignant la faible magnitude et l'hétérogénéité des impacts en fonction de l'âge et du niveau d'étude de la force de travail. Dans le second chapitre, nous construisons un modèle théorique visant à expliquer l'émergence du chômage technologique et de la polarisation du marché du travail. Nos résultats confirment que l'automatisation est un facteur clef de ces deux phénomènes, dont l'ampleur est proportionnelle à la vélocité du progrès technologique. Enfin, le troisième chapitre teste quatre politiques, et les résultats indiquent qu'il est possible de lutter efficacement contre le chômage technologique et la polarisation du marché du travail.

Mots-clés : Robotique, intelligence artificielle, chômage technologique.

Résumé en anglais

This thesis studies the impact of robots and artificial intelligence (AI) on employment and inequality from its empirical, theoretical and policy aspects. In the first chapter, we conduct an econometric study on 33 OECD countries, and observe that robots and AI increase the unemployment rate, while highlighting the low magnitude and heterogeneity of the impacts according to the age and education level of the labor force. In the second chapter, we build a theoretical model to explain the emergence of technological unemployment and labor market polarization. Our results confirm that automation is a key factor in these two phenomena, the magnitude of which is proportional to the velocity of technological progress. Finally, the third chapter tests four policies, and the results indicate that technological unemployment and labor market polarization can be effectively addressed.

Keywords : Robotics, artificial intelligence, technological unemployment.