

« Automation and labor market polarization in an evolutionary model with heterogeneous workers »

Auteurs


Florent Bordot, Andre Lorentz

Document de Travail n° 2021 – 39

Septembre 2021

**Bureau d'Économie
Théorique et Appliquée
BETA**

www.beta-umr7522.fr

 @beta_economics

Contact :
jaoulgrammare@beta-cnrs.unistra.fr

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

Florent Bordot* André Lorentz†

September 24, 2021

Abstract

The purpose of this paper 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 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.

Keywords: Automation; Wage Polarization ; Technical Change ; Employment ; Agent-Based Model

JEL:C63,E14,J21,J31,033

*LEM, Sant'Anna School for Advanced Studies, Pisa, Italy ; Université de Strasbourg, Université de Lorraine, CNRS, BETA, Strasbourg, France ; f.bordot@santannapisa.it

†Université de Strasbourg, Université de Lorraine, CNRS, BETA, Strasbourg, France; alorentz@unistra.fr

1 Introduction

Advances in robotics and artificial intelligence are increasingly raising questions about the impact of technology on 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 has grown leading to an enlarge 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, the evidence supporting a negative effect of replacement investment on labor demand is rather weak. 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 UK (Salvatori 2018), Japan (Furukawa and Toyoda 2018) and Europe (Goos et al. (2014), Náplava (2019)); and is induced by many factors: immigration, population aging, female labor force entry, offshoring and technical change. Autor and Dorn (2013) analyze the effect of these factors, and 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 employment inequality. 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 drop in "degree of automation" of 25% over this period. These data are based on expert estimates and are, as the authors point out, sensitive to the subjectivity of the expert 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 statically robust correlation between the complexity index and the evolution of wages between 1980 and 2005, while the routine index is not statically significant. With regard to 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 decomposed into two types of tasks: routine and non-routine tasks. Therefore, the production function is composed of routine and non-routine labor input and computer capital. Workers have a heterogeneous productivity endowment and may reallocate to only one of the two tasks, or to 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 a 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 have 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 (2013). 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 zero and the employer would have no interest in hiring him. To illustrate, we make the reasonable assumption that a cashier cannot perform the job of a neurosurgeon. In order to simplify and keep a tractable model, we use 6 different occupations in our model (data are presented in the Appendix and the methodology used to aggregate the data is discussed in Section 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 will 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 is able to 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 use and extend the original K+S model developed by Dosi et al. (2006): Dosi, Pereira, Roven-tini and Virgillito (2021) extend the Dosi et al. (2018) model, itself based on Dosi et al. (2006), 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, Caiani, Russo et al. (2021) developed their own ABM to show that automation can triggers 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 will indeed 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 will always be 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 for example 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 will 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 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 patient, it is very likely that he or she will be much less efficient than the hairdresser. For practical reasons, we keep an aggregate skill index so that we can calculate the skill 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 will 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 kind of technical progress: one which improve capital productivity and is labor-augmenting, and a second one which is labor-displacing and aim at automating some tasks. Regarding the labor-enhancing technical progress, the firms targets 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, he will increase more rapidly the skill required to perform task A than an agent spending an equal amount of time between the two tasks, but will improve more slowly the skill required to perform task B. 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-augmenting technical progress balances the allocation of time between tasks: workers will 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 required skills associated with them. As a result, agents will be more

specialized and therefore more efficient on the remaining tasks.

The structure of the model reflect 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 will be different, and in a sense more complex, than a set of tasks requiring the same type of skill but with 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 pre-suppositions 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 intertemporal maximization calculation impossible. Third, there is path dependence in both innovation and skill dynamics. Difficulties in recruiting engineers and/or poor innovation output will have a persistent effect on the firm's productivity. Periods of underemployment or unemployment will have a permanent effect on an agent's skill levels. Fourth, distributions matter: a broad polarization of the skill structure will 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 the exception. Finally, the model also used Schumpeterian dynamics: technical progress is endogenous and generates a process of 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.

1.1 Research questions

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 The Model

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




Figure 1: The structure of the model

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.

2.1 Production

The model is composed by a population of F heterogeneous firms indexed $f \in [1; F]$ producing an homogeneous consumption good. The firms uses labor to produce this good. The production capacities of the firms are assumed to reflect the evolution of capital, and the productivity it embodies and the organizational changes at the firms 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:

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

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 given job m requires a combination of tasks to be realized. We assume here that there exist a vector of I tasks indexed $i \in [1; I]$. Each tasks 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:

$$\frac{L_{m,f}(t)}{A_{m,f}(t)} = \min \left\{ \frac{L_{m,f}(t)}{B_{1,m,f}(t)}; \dots; \frac{L_{m,f}(t)}{B_{i,m,f}(t)}; \dots; \frac{L_{m,f}(t)}{B_{I,m,f}(t)} \right\} \quad (2)$$

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 the task i on the job m .

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

$$\frac{1}{A_{m,f}(t)} = \min \left\{ \frac{1}{B_{1,m,f}(t)}; \dots; \frac{1}{B_{i,m,f}(t)}; \dots; \frac{1}{B_{I,m,f}(t)} \right\} \quad (3)$$

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 workers 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 workers 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 on a specific task, the more it benefits from the productivity embodied in the technology (see section 2.4), the higher their individual productivity. More formally, $B_{i,m,f}(t)$ can be formalized as follows:

$$B_{i,m,f}(t) = \begin{cases} \varepsilon & \text{if the task is automated} \\ \left(a_{i,m,f}(t-1) \sum_{j=1}^{L_{m,f}(t)} \frac{b_{j,i,m,f}(t)}{L_{m,f}(t-1)} \right)^{-1} & \text{if the task is performed by workers} \end{cases} \quad (4)$$

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 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 current stocks ($Y_f(t-1) - D_f(t-1)$):

$$Y_f^D(t) = \alpha Y_f^D(t-1) + (1 - \alpha) D_f(t-1) - (Y_f(t-1) - D_f(t-1)) \quad (5)$$

In this respect, we assume that firms have no information on the level on the current state of demand when planning their production and build 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 exact same quantity than 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 Demand, Market Dynamics and Pricing Behavior

The level of 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 an homogeneous product. The aggregate demand corresponds to the sum of the consumption level desired by the agents. Formally, the demand addressed to each firms, in nominal terms can be expressed as:

$$D_f(t) = z_f(t) \left[\sum_{j=1}^J C_j^D(t) + \sum_{\tau=1}^{t-1} \left(\sum_{j=1}^J C_j^D(\tau) - \sum_{f=1}^F p_f(\tau) Q_f(\tau) \right) \right] \quad (6)$$

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

$$C_j^D(t) = \psi * C_j(t-1) + (1 - \psi) * W_j(t) \quad (7)$$

We assume here a non-Walrasian market dynamics (Gaffard (2018)) on this homogeneous good market. There is no market clearing, but at each time period, the total amount of expenditures is shared among firms according to a replicator dynamics mechanism representing an imperfect competition process (Metcalf 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:

$$z_f(t) = z_f(t-1) \left[1 + \phi \left(\frac{E_f(t)}{\bar{E}(t)} - 1 \right) \right] \quad (8)$$

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 backlogs or 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 turns, 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:

$$E_f(t) = \frac{1}{\omega p_f(t) + (1 - \omega) \left(\frac{D_f(t-1)}{p_f(t-1)Q_f(t-1)} - 1 \right)} \quad (9)$$

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:

$$\bar{E}(t) = \sum_{f=1}^F z_f(t-1) E_f(t) \quad (10)$$

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

$$p_f(t) = (1 + \bar{\mu}) \sum_{m=1}^M w_{m,f}(t) A_{m,f}(t) \quad (11)$$

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

$$Q_f(t) = \min \left\{ Y_f(t); \frac{D_f(t)}{p_f(t)} \right\} \quad (12)$$

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

$$\pi_f(t) = p_f(t) Q_f(t) - \sum_{m=1}^M w_{m,f}(t) L_{m,f}(t) \quad (13)$$

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.

2.3 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 jobs, corresponds to the 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 in 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, Pereira, Roventini and Virgillito 2017). Formally, the level of labor demand per jobs can be expressed as follows:

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

This expression can be simplified as follows:

$$L_{m,f}^D(t) = A_{m,f}(t-1)Y_f^D(t) + (1-\lambda)(T_{m,f}(t-1) - L_{m,f}(t-1)) \quad (14)$$

$L_{m,f}^D(t)$ is interpreted in our model as the number of slots available a 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 in the past period to attract enough workers to satisfy its needs on a specific job, the wages raises. More formally, wages per jobs are set as follows:

$$w_{m,f}(t) = w_{m,f}(t-1) \left[1 + \xi_1 \left(\frac{A_f(t-1)}{A_f(t-2)} - 1 \right) + \xi_2 \max \left\{ 0; 1 - \frac{L_{m,f}^S(t-1)}{L_{m,f}^D(t-1)} \right\} \right] \quad (15)$$

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 evidences (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 type of job opening to satisfy the firm's labor demand $L_{m,f}^D(t)$ is characterized by a vector of tasks each requires a minimum level of skill for each tasks to be performed. For a given 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:

$$\bar{S}_{m,f}(t) = \begin{pmatrix} \bar{s}_{1,m,f}(t) \\ \vdots \\ \bar{s}_{i,m,f}(t) \\ \vdots \\ \bar{s}_{I,m,f}(t) \end{pmatrix} \quad (16)$$

We assume, here that the firms are conservative in the skill levels of the required in the job openings, so that they expect a minimum skill level corresponding to the 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 the 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:

$$\bar{s}_{i,m,f}(t) = \begin{cases} 0 & \text{if the task is automated} \\ \frac{\sum_{j=1}^{L_{m,f}^{(t-1)}} s_{j,i,m,f}(t-1)}{L_{m,f}^{(t-1)}} & \text{if the task is performed by labor} \end{cases} \quad (17)$$

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 tasks i :

$$S_j(t) = \begin{pmatrix} s_{j,1}(t) \\ \vdots \\ s_{j,i}(t) \\ \vdots \\ s_{j,I}(t) \end{pmatrix} \quad (18)$$

An agent can apply for a job if and only if his/her skills fit the minimum required for that job. Formally, each elements of the vector of skills of a potential

candidate have to be greater or equal to the elements of the vector of minimum requirements of that job:

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

At each time 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 match the minimum requirements in terms of skill levels 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 opening for which they match the minimum requirements.

We assume here that each agent can only apply to one job at the time, which is randomly drawn from the set of jobs the agent can apply for. The probability to candidate 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 job 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 agent, applying for jobs, 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; second agents are then ranked accordingly; third the firms hire according to the ranking until the openings for the job m are filled.

We assume that the firms only have an 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. More formally this estimate can be described as follows:

$$s_{j,i,m,f}^E(t) \sim N(s_{j,i}(t); (s_{j,i}(t) - \bar{s}_{i,m,f}(t))) \quad (19)$$

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 :

$$S_{j,m,f}^E(t) = \sum_{i=1}^I s_{j,i,m,f}^E(t) \frac{B_{i,m,f}(t-1)}{\sum_{i=1}^I B_{i,m,f}(t-1)} \quad (20)$$

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 remain unemployed. A 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 quantity 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.

$$L_{m,f}(t) = (1 - \lambda)(L_{m,f}(t-1) - T_{m,f}(t)) + \min\{L_{m,f}^D(t); L_{m,f}^S(t)\} \quad (21)$$

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 are required for the firm to increase its competitiveness in the long-run and survive the market selection mechanisms.

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. 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 benefits from the productivity gains resulting from the technical changes resulting by their R&D activity.

We assume that individual skills evolves 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 micro-economic models of production (See among others Llerena, Lorentz, Marengo and Valente 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 workers therefore 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 learning pace is correlated to the distance between its skill and the frontier. Say it in a different way, 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:

$$s_{j,i}(t) = s_{j,i}(t-1) \left[1 - \delta_1 \left(1 - \frac{B_{i,m,f}(t-1)}{\sum_{i=1}^I B_{i,m,f}(t-1)} \right) + \delta_2 \frac{B_{i,m,f}(t-1)}{\sum_{i=1}^I B_{i,m,f}(t-1)} (s_i^{MAX} - s_{j,i}(t-1)) \right] \quad (22)$$

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 the 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.

$$b_{j,i,m,f}(t) = e^{\kappa(s_{j,i,m,f}(t-1) - s_i^{MAX})} \quad (23)$$

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

This technical change is rooted in the firms R&D activity. Each period, the firm devotes part of its resources in R&D expenditures ($R_f(t)$). These investments are constrained by the resources accumulated by the firm in the past. These assumption are both in line with empirical evidences (Coad and Rao 2010) and standard in the evolutionary literature (See among others Llerena and Lorentz 2004, Lorentz and Savona 2008, Ciarli, Lorentz, Savona and Valente 2010, Dosi et al. 2006, Dosi et al. 2018). These highlight the financial constraints impose 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:

$$R_f(t) = \min \left\{ \eta p_f(t-1) * Q_f(t-1); \sum_{\tau=1}^{t-1} (\pi_f(\tau) - R_f(\tau)) \right\} \quad (24)$$

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 $w_{E,f}(t)$:

$$L_{E,f}^D(t) = \frac{R_f(t)}{w_{E,f}(t)} \quad (25)$$

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 jobs are set as follows:

$$w_{E,f}(t) = w_{E,f}(t-1) \left[1 + \xi_1 \left(\frac{A(t-1)}{A(t-2)} - 1 \right) + \xi_2 \max \left\{ 0; 1 - \frac{L_{E,f}^S(t-1)}{L_{E,f}^D(t-1)} \right\} \right] \quad (26)$$

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:

$$L_{E,f}(t) = \min \{ L_{E,f}^D(t); L_{E,f}^S(t) \} \quad (27)$$

The matching mechanism is symmetric to the one use at the production level described in paragraph 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 improve the task level labor productivity $a_{i,m,f}$, on the other. For a given time 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 the task i with respect to the total wage bill of the firm:

$$\hat{i}(t) = i \in [1; I] \mid x_{\hat{i},f}(t) > x_{i,f}(t) \forall i \neq \hat{i} \quad (28)$$

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

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 $\frac{1}{A_{E,f}(t)}$:

$$P[\text{Innovation} = 1] = 1 - e^{-\rho \frac{1}{A_{E,f}(t)} L_{E,f}(t)} \quad (29)$$

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 $\varepsilon_{i,f}^\sigma(t)$ is randomly drawn from a Gaussian Law with centered around zero and an endogenous variance depending of the distance between the maximum skill level of the task to automate (s_i^{MAX}) and the technological frontier of the firm $\sigma_{i,f}(t-1)$:

$$\sigma_{i,f}(t) = \sigma_{i,f}(t-1)(1 + \max\{0; \varepsilon_{i,f}^\sigma(t)\}) \quad (30)$$

$$\text{with } \varepsilon_{i,f}^\sigma(t) \sim N(0; \beta(s_i^{MAX} - \sigma_{i,f}(t-1)))$$

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 around zero and an exogenous

variance γ_i :

$$a_{i,m,f}(t) = a_{i,m,f}(t-1)(1 + \min\{0; \varepsilon_{i,m,f}^a(t)\}) \quad (31)$$

with $\varepsilon_{i,f}^a(t) \sim N(0; \gamma_i) \forall m$

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 put the emphasis on the importance of the qualitative transformations of job activities in order to explain the employment dynamics.

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.

2.6 Timeline of the model

1. Based on their production and demand during the previous period, firms fixed their desired level of production [equation 5]
2. Firms fixed 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 24]
3. Based on $Y_f^D(t)$ and $RD_f^D(t)$, firms compute the number of people they want to hired or fired (equation 14 for production jobs, and equation 25 for engineers).
4. For each job, firms made a wage proposal [equation 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 candidate 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 19 and 20], rank them and hire only the best ones. If, for a given job, the firm want 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 11]
9. Workers set their desired level of consumption based on their consumption during the previous period and their current income. [equation 7]
10. The skills levels of each agent are updated. [equation 22] Consequently, the agent's productivity on each type of tasks is updated [equation 23].
11. The R&D outputs are revealed [equations 30 and 31].
12. Firms update their minimum skills levels requirement based on the average skills levels of their workers [equation 17].
13. Given the evolution of workers' productivity and R&D outputs, firms' productivity is updated. [equations 3 and 4].
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.

3 Simulation Results

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. In order 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 let 26.5% of them unemployed. We deliberately chose a high initial unemployment rate in order to avoid a labor shortage right during the first simulation steps of the model. Workers' skills are initialized on the basis of the skills required to perform their job, plus a stochastic component to generate some heterogeneity among workers having the same occupation :

$$s_{j,i} = \bar{s}_{i,m,f}(t)(1 + \max\{0; N(0, k)\}) \quad (32)$$

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

$$s_{j,i} = \bar{s}^u_i(t)(1 + \max\{0; N(0, k^u)\}) \quad (33)$$

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 on 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 continues over the following periods, and adding more periods would not have added more information. 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 4 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%.

3.2 Polarization measurement

Labor market polarization is characterized by two phenomena: 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 two others type of workers. We do not use this dichotomy, which could potentially influence the results via a threshold effect linked to the skill values used for ranking, but we do aggregate the skills in order to have a metric for computing the degree of labor market polarization. To do this, we calculate a skill index \hat{s} for each job that takes into account both the average skill level required and the degree of specialization:

$$\hat{s}_{m,f}(t) = \frac{\sum_{i=1}^I \bar{s}_{i,m,f}(t)}{I} \left[1 + \sum_{i=1}^I \left(\frac{\bar{s}_{i,m,f}(t)}{\sum_{i=1}^I \bar{s}_{i,m,f}(t)} \right)^2 \right] \quad (34)$$

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

$$PI_s(t) = \frac{1}{L_t} \left[\sum_{f=1}^F \sum_{m=1}^M \left[L_{m,f}(t) \frac{(\hat{s}_{m,f}(t) - m(\hat{s}(t)))}{m(\hat{s}(t))} \right]^r \right] \quad (35)$$

Where \hat{s} is the required skill index, and $m(\hat{s})$ the median skill index. Wang and Tsui (2000) recommend to choose 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. To illustrate that the choice of this value does not drive the results, the graphs of the wage and skill distributions will accompany the results of the polarization indexes presented for the baseline and automation scenario.

Similarly, we calculate a wage polarization index:

$$PI_w(t) = \frac{1}{L_t} \left[\sum_{f=1}^F \sum_{m=1}^M \left[L_{m,f}(t) \frac{(w_{m,f}(t) - m(\hat{w}(t)))}{m(\hat{w}(t))} \right]^r \right] \quad (36)$$

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

3.3 Empirical validation

In the baseline scenario without automation, the model is able to 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 the heterogeneity of agents in term of skills and the pro-cyclical accumulation of skills for workers are guaranteed by Equation 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 17).

Other stylized facts generated by the model are the results of complex micro-dynamics. At the micro level, we observe in figure 2 a strong heterogeneity in the dynamics of firms' productivity (Bartelsman and Doms (2000) and Dosi (2008)).

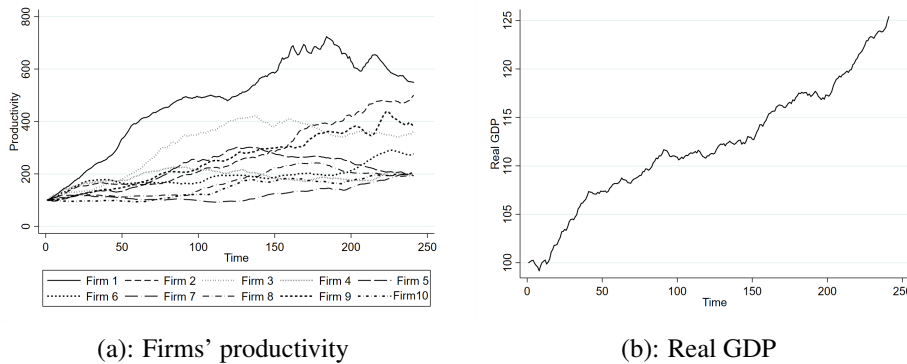




Figure 2: Evolution of firms' productivity (panel (a)) and Real GDP (panel (b)) - Base 100 in step 0 - 10-periods moving average - Baseline scenario

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, real GDP is growing and fluctuating endogenously. Labor productivity and the average real wage follows a similar pattern, which translates into a wage share that remains roughly constant, fluctuating between 82% and 85% (Figure 3).

Once we introduce automation, the model successfully generates automation spikes at the firm level (figure 5). Logically, labor productivity is stronger in the scenario with automation, leading also to a higher average wage. However, we notice that the labor gap between productivity and wages widens over time (figure 6). In period 50, the gap is about 6.5 points, and peaks at 17.3 points at the end of the simulation. This dynamic is reflected in the trend decline in the wage share, which falls by 10 points between the beginning and the end of the simulation. The downward trend in the wage share is a well-documented stylized fact (Karabarbounis and Neiman (2014)). One explanation provided in the literature (Dao et al. (2017), Acemoglu and Restrepo (2019), and Cheng et al. (2021)) is related to automation, and the results of our model seems to corroborate this hypothesis.



(a): Labor productivity and average real wage



(b): Wage share

Figure 3: Evolution of Labor productivity and average real wage (panel (a) - Base 100 in step 0) and Wage share (panel (b)) - both 10-periods moving average - Baseline Scenario

Finally, figure 4 shows a positive correlation between the change in unemployment and inequality, consistent with empirical observations (Mocan (1999), Pontusson and Weisstanner (2016), Deysappriya (2017)).

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 1.1.




Figure 4: The unemployment-inequality nexus





Figure 5: Automation spikes



(a): Labor productivity and average real wage

(b): Wage share

Figure 6: Evolution of Labor productivity and average real wage (panel (a) - Base 100 in step 0) and Wage share (panel (b)) - both 10-periods moving average - Automation Scenario

3.4 Baseline scenario

In this baseline scenario, there is no automation. The only form of technological progress is labor-augmenting. To generate this scenario, we set $\beta = 0$ in equation 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 7, and show that the labor market exhibits a polarization pattern in the wage

structure, but not in the skill structure.




Figure 7: Polarization indexes (equations 35 and 36) - Baseline scenario

As regards to wages, this result is explained by the dynamics of wage setting (equation 15), which depends on two factors: the evolution of the firm's productivity (ξ_1) and the difficulty of recruitment (ξ_2). With respect to 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 22) and skill demand (equation 17), heterogeneity in labor supply and labor demand increase across time. As a result, tensions can appear on the labor market, especially for jobs requiring high and varied skills, thus accelerating the growth of high wages.

The presence of wage polarization is confirmed by the evolution of the wage distribution between the first and last period of the simulation (Figure 8). This polarization is mainly driven by high wages: the share of wages above 1.5 times the median has strongly increased, while that at the other end of the distribution, the share of wages below 0.5 times the median, which were almost non-existent in step 1, have also increased but to a lesser extent. Regarding skills, we observe in Figure 7 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 9).

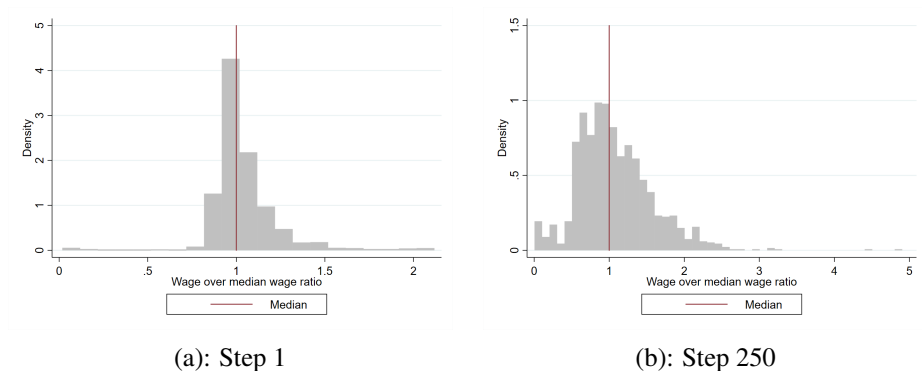


Figure 8: Wage distribution Step 1 (panel (a)) vs. Step 250 (panel (b) - Baseline Scenario

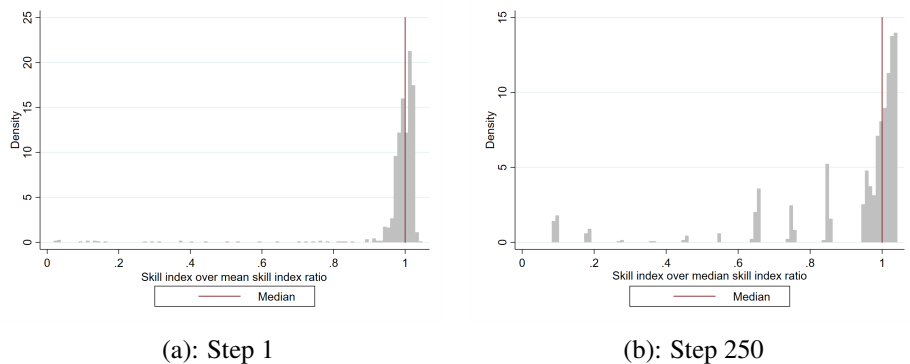


Figure 9: Skill index distribution Step 1 (panel (a)) vs. Step 250 (panel (b) - Baseline Scenario

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 10 shows the evolution of the real median wage and the median skill. Over the whole simulation, the real median wage remains relatively constant and fluctuates around its initial value. Interestingly, this dynamic is similar to the evolution of the median hourly wage in the US during last 4 decades, with an increase of only 15.1% between 1979 and 2019 (Gould (2020)).

The evolution of skills shows that there is a general rise in workers' qualifications. Consequently, the agents on the left of the median (Figure 9) 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




Figure 10: Median skill and median real wage evolution - Baseline - Base 100 on step 0

(equation 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 allow us to answer the first research question (RQ1). Without 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. Nevertheless, this polarization differs from the one traditionally observed in the literature: it is an asymmetric type of polarization, which is driven by the increase of high wages only, without really affecting wages around the median. Finally, the distribution of skills does not seem to be affected, and the polarization observed in many countries are 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.

3.5 Automation scenario

In this second scenario, we introduce automation by setting $\beta = 0.1$ in equation 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 certain tasks. Firms follow a simple rule of cost reduction by trying to automate the tasks that are most costly. Figure 11

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 wage polarization index at the end of the simulation very close to the one we obtained in the baseline scenario.



Figure 11: Polarization indexes - Automation scenario




Figure 12: Dynamics of the Gini indices

A look at the dynamics of the Gini indices, presented in Figure 12, computed only on workers to avoid 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.

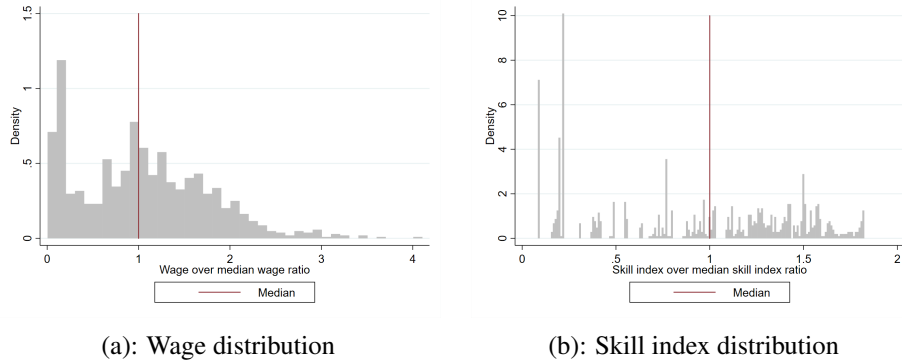


Figure 13: Wage distribution (panel (a)) and Skill index distribution (panel (b)) - Automation scenario - Step 250

The upward trend in polarization index for skills in Figure 11 and the atomized skill distribution presented in Figure 13 highlight the effect of automation of the skill structure. The strong increase in the density of skills indexes lower than half of the median skill index testifies to the impact of automation on the qualitative aspect of labor demand. We observe a strong increase in the density of jobs having skills requirement corresponding to less than half of the median skill index, and similarly for the right-hand side of the distribution requiring a skill level 1.5 times higher than the median level.

In contrast to the baseline scenario, the median skill decreases over time (Figure 14). In fact, 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, around step 50, corresponds to the moment when firms begin to automate tasks. As agents improve their skills through a process of learning by doing, they will tend to forget how to perform some tasks that they no longer need to do. In some extreme cases where almost all tasks are automated in a job, a worker will lost most of his skills and only retain a kind of "supervisory skill" that allows him to perform only one task: 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




Figure 14: Median skill and median real wage evolution - Automation scenario

(2019) use a panel of more than 160 countries show that automation in the 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.

In order to confirm that these findings are robust and not simply a matter of randomness, we have carried out a t-test presented in Table 1. We find a statistically significant difference between the skill polarization index of the baseline scenario and the one of the automation scenario, but no difference in the wage 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 employment rate half 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 13) compared to the baseline scenario (Figure 9). 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.

We have seen that automation has an impact on the polarization of the skills structure and on wage inequality. However, the question arises whether slower

	Automation	Baseline	Difference	P-value
Skills polarization index	0.189	0.031	0.158	1.14e-36
Wages polarization index	0.824	0.871	-0.047	0.1720
Gini index	0.736	0.589	0.147	2.67e-12
Median skill index	3.446	4.714	-1.260	7.64e-52
Real median wage	943.6	747.1	196.5	0.000703
Employment rate	0.402	0.618	-0.216	3.69e-22
Labor productivity	2181.3	1332.6	848.7	1.45e-14

Table 1: *Automation and baseline scenarios*

automation, expressed in the model by a lower value of the parameter Beta, leads to similar conclusions. In order to answer this question, we consider a third scenario in which the automation of production proceeds more slowly.

3.6 Slow automation scenario

In this scenario, we set the value of the Beta parameter in equation (30) to 0.06. The choice of this value is motivated by the rate of automation generated by such a value. Figure 15 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, 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 have a negative impact on the employment rate, but not on 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. We can observe a relation between the pace of automation and the deformation of the skill structure, with an increase in the value of the skill 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 (equation 15). We observed that an acceleration of automation lead to a decrease of the employment rate, which is consistent with fragmentation of the skill structure leading to an increase of the mismatch between labor supply and labor demand.




Figure 15: Share of automated tasks - Automation and slow automation scenarios

Finally, we observe in figure 16 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 workers will be on tasks that have not been automated yet. But this increasing specialization is a double-edged sword: given the loss of versatility, it will be more complicated for a worker to find another job if he is dismissed. Indeed, a task automated in firm A will not be necessarily automated in firm B, and so the second firm will require candidates to have sufficient skills to perform this task.

The faster the automation, the higher the specialization but the lower the median skill index (Table 2). To resolve this apparent contradiction, we need to refer to the formula used to compute the skill index, described in equation 34. In order 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 have 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 does not, on average, compensate for the loss of skills resulting from the automation of the others 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.




Figure 16: Dynamics of specialization indexes

	Slow	Auto.	Baseline	Difference	P-value	Difference	P-value
	(1)	(2)	(3)	(1) - (2)	(1) - (2)	(1) - (3)	(1) - (3)
Skills polarization index	0.074	0.189	0.031	-0.115	1.22e-24	0.044	2.14e-25
Wages polarization index	0.883	0.824	0.871	0.059	0.0608	0.012	0.757
Gini index	0.685	0.736	0.589	-0.051	0.00970	0.096	0.000002
Median skill index	4.199	3.446	4.714	0.752	3.81e-30	-0.516	4.22e-37
Real median wage	752.1	943.6	747.1	-191.5	0.00151	5.034	0.914
Employment rate	0.469	0.402	0.618	0.067	0.00104	-0.149	3.83e-15
Labor productivity	1731.2	181.3	1332.6	-450.1	0.00006	398.5	1.20e-09

Table 2: *Slow automation scenario versus Automation and baseline scenarios*

3.7 Results summary

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 ?
- 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 accentuate 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 (30)) to automate medium-skilled routine tasks. There is also a feedback 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 will be relatively low and therefore firms will 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 wages structure does not vary much between the two scenarios, the introduction of automation in the model leads to a rising wages inequality illustrated by the increase of almost 25% of the Gini index. (table 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 classical economic factors, as for example labor costs, in the explanation of the polarization process. Further theoretical and empirical works should focus more on these factors.

In addition to the main results answering the two research questions, we observe two secondary results that were not in our initial focus. A comparison between the baseline scenario and the automation scenario shows that in this model, automation is a driving force in the decline of the labor share (Figure 6). Even though the average real wage almost doubled between the first and the last stage of the simulation, wages moved more slowly than the evolution of labor productivity, leading to a value added distribution favorable to profits over wages. This theoretical result supports the evidence provided by Dao et al. (2017), Acemoglu and Restrepo (2019) and Cheng et al. (2021) on the role of automation on the decline of the labor share.

Finally, in this model, automation appears to be a "deskilling" process: while the degree of specialization increases in the scenario with automation (16), the median skill index decreases as the number of automated tasks increases (Figure 14). 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)).

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 is able to 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, companies do not manage to automate them but do manage to automate slightly less expensive but much less complex tasks. On the other hand, companies will not necessarily

devote R&D efforts to a task that is not very complex but also not very costly due to an abundance of available labor will not necessarily be automated. This result provides some support to the complex-task biased technical change hypothesis developed by Caines et al. (2017).

We have also highlighted a perverse 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 a decrease in the employment rate. Finally, automation seems to have a negative impact on the labor income share: in the scenario without automation the wage share is roughly constant, oscillating between 82% and 85% (Figure 3), while it exhibits a downward trend in the scenario with automation (Figure 6), dropping by 10 points between the beginning and the end of the simulation. These results provide support to the literature that studies the role of automation on the wage share (Dao et al. (2017), Acemoglu and Restrepo (2019) and Cheng et al. (2021)).

A severe polarization of the labor market would lead to two problems: increasing difficulties in entering the labor market for people without the adequate skills, and increasing 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 induce by technology do not lead to a "winners take all" scenario that would significantly increase inequalities.

References

- 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. and Restrepo, P.: 2019, Automation and new tasks: How technology displaces and reinstates labor, *Journal of Economic Perspectives* **33**(2), 3–30.
- Acemoglu, D. and Restrepo, P.: 2020, Robots and jobs: Evidence from us labor markets, *Journal of Political Economy* **128**(6), 2188–2244.
- Amendola, M. and Gaffard, J. L.: 1998, *Out of Equilibrium*, Clarendon Press, Oxford.
- Arntz, M., Gregory, T. and Zierahn, U.: 2016, The risk of automation for jobs in oecd countries: A comparative analysis.
- 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, 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.
- Bartelsman, E. J. and Doms, M.: 2000, Understanding productivity: Lessons from longitudinal microdata, *Journal of Economic literature* **38**(3), 569–594.
- 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.
- 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.
- Cheng, H., Drozd, L. A., Giri, R., Taschereau-Dumouchel, M. and Xia, J.: 2021, The future of labor: Automation and the labor share in the second machine age, *Technical report*, Federal Reserve Bank of Philadelphia.
- 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.
- 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.
- Coad, A. and Rao, R.: 2010, Firm growth and r&d expenditure, *Economics of Innovation and New Technology* **19**(2), 127–145.
- 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.
- Dao, M. C., Das, M. M., Koczan, Z. and Lian, W.: 2017, *Why is labor receiving a smaller share of global income? Theory and empirical evidence*, International Monetary Fund.
- Dauth, W., Findeisen, S., Suedekum, J., Woessner, N. et al.: 2018, Adjusting to robots: Worker-level evidence, *Opportunity and Inclusive Growth Institute Working Papers* **13**.
- Dawid, H. and Hepp, J.: 2021, Distributional effects of technological regime changes: Hysteresis, concentration and inequality dynamics.
- Deyshappriya, N.: 2017, Impact of macroeconomic factors on income inequality and income distribution in asian countries.
- 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., 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. and Virgillito, M. E.: 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.
- 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.
- Fierro, L. E., Caiani, A., Russo, A. et al.: 2021, Automation, job polarisation, and structural change, *Technical report*.
- 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.
- 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.

- Gatti, D. D., Desiderio, S., Gaffeo, E., Cirillo, P. and Gallegati, M.: 2011, *Macroeconomics from the Bottom-up*, Vol. 1, Springer Science & Business Media.
- Goos, M., Manning, A. and Salomons, A.: 2014, Explaining job polarization: Routine-biased technological change and offshoring, *American economic review* **104**(8), 2509–26.
- Gould, E.: 2020, State of working america wages 2019: A story of slow, uneven, and unequal wage growth over the last 40 years, *Economic Policy Institute* .
- Graetz, G. and Michaels, G.: 2018, Robots at work, *Review of Economics and Statistics* **100**(5), 753–768.
- Hein, E. and Lavoie, M.: 2019, Post-keynesian economics, *The Elgar Companion to John Maynard Keynes*, Edward Elgar Publishing.
- Jo, Y. J.: 2019, Downward nominal wage rigidity in the united states, *Job Market Paper* .
- Kaldor, N.: 1957, A model of economic growth, *Economic Journal* **December**, 65–94.
- Karabarbounis, L. and Neiman, B.: 2014, The global decline of the labor share, *The Quarterly journal of economics* **129**(1), 61–103.
- Knight, F. H.: 1921, *Risk, uncertainty and profit*, Vol. 31, Houghton Mifflin.
- Kunst, D.: 2019, Deskillling among manufacturing production workers.
- 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. and Savona, M.: 2008, Evolutionary micro-dynamics and changes in the economic structure, *Journal of Evolutionary Economics* **18**(3-4), 389–412.
- 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.

- 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.
- Metcalfe, J. S.: 1994, Competition, fisher’s principle and increasing returns in the selection process, *Journal of Evolutionary Economics* **4**, 327–346.
- Mocan, H. N.: 1999, Structural unemployment, cyclical unemployment, and income inequality, *Review of Economics and Statistics* **81**(1), 122–134.
- Náplava, R.: 2019, Changing structure of employment in europe: Polarization issue, *Review of Economic Perspectives* **19**(4), 307–318.
- Nelson, R. and Winter, S.: 1982, *An Evolutionary Theory of Economic Change*, Harvard University Press.
- 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.
- Pontusson, J. and Weisstanner, D.: 2016, The political economy of compensatory redistribution: Unemployment, inequality and policy choice, *Technical report*, LIS Working Paper Series.
- Salvatori, A.: 2018, The anatomy of job polarisation in the uk, *Journal for labour market research* **52**(1), 8.
- Scholl, K. and Hanson, R.: 2020, Testing the automation revolution hypothesis, *Economics Letters* **193**, 109287.
- 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.
- 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.

A Appendix




Figure 17: Real GDP - Baseline scenario - 1000 steps and 100 replications




Figure 18: Labor productivity - Baseline scenario - 1000 steps and 100 replications

Agent	Component	Equation
Firms	Production	$Y_f(t) = \min \left\{ \frac{L_{1,f}(t)}{A_{1,f}(t)}, \dots, \frac{L_{m,f}(t)}{A_{m,f}(t)}, \dots, \frac{L_{M,f}(t)}{A_{M,f}(t)} \right\} \quad (1)$
	Desired production	$Y_f^D(t) = \alpha Y_f^D(t-1) + (1-\alpha) D_f(t-1) - (Y_f(t-1) - D_f(t-1)) \quad (5)$
	Job productivity	$\frac{1}{A_{m,f}(t)} = \min \left\{ \frac{1}{B_{1,m,f}(t)}, \dots, \frac{1}{B_{i,m,f}(t)}, \dots, \frac{1}{B_{I,m,f}(t)} \right\} \quad (3)$
	Task intensity	$B_{i,m,f}(t) = \begin{cases} 0 & \text{if the task is automated} \\ \left(a_{i,m,f}(t-1) \sum_{j=1}^{L_{m,f}(t)} \frac{b_{j,i,m,f}(t)}{L_{m,f}(t-1)} \right)^{-1} & \text{if the task is performed by workers} \end{cases} \quad (4)$
	Labor demand	$L_{m,f}^D(t) = A_{m,f}(t) Y_f^D(t) + (1-\lambda)(T_{m,f}(t-1) - L_{m,f}(t-1)) \quad (14)$
	Wage proposal	$w_{m,f}(t) = w_{m,f}(t-1) \left[1 + \xi_1 \left(\frac{A_f(t-1)}{A_f(t-2)} - 1 \right) + \xi_2 \max \left\{ 0; 1 - \frac{L_{m,f}^S(t-1)}{L_{m,f}^D(t-1)} \right\} \right] \quad (15)$
	R&D	$R_f(t) = \min \left\{ \eta p_f(t-1) * Q_f(t-1); \sum_{\tau=1}^{t-1} (\pi_f(\tau) - R_f(\tau)) \right\} \quad (24)$
	Technological frontier	$\sigma_{i,f}(t) = \sigma_{i,f}(t-1) (1 + \max\{0; \varepsilon_{i,f}^\sigma(t)\}) \quad (30)$ with $\varepsilon_{i,f}^\sigma(t) \sim N(0; \beta (s_i^{MAX} - \sigma_{i,f}(t-1)))$
	Capital productivity	$a_{i,m,f}(t) = a_{i,m,f}(t-1) (1 + \min\{0; \varepsilon_{i,m,f}^a(t)\}) \quad (31)$ with $\varepsilon_{i,f}^a(t) \sim N(0; \gamma_i) \forall m$
	Workers	Skills
Worker's efficiency		$b_{j,i,m,f}(t) = e^{\kappa (s_{j,i,m,f}(t-1) - s_i^{MAX})} \quad (23)$
Desired level of consumption		$C_j^D(t) = \pi * C_j(t-1) + (1-\pi) * W_j(t) \quad (7)$

Table 3: 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''	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 4: 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 5: Data used to initialize the model.