

GROWINPRO

Growth Welfare Innovation Productivity

Working Paper

Effects of Technological Change and Automation on Industry Structure and (Wage-)Inequality: Insights from a Dynamic Task-Based Model

Herbert Dawid

Department of Business Administration and Economics and Center for Mathematical Economics, Bielefeld University

Michael Neugart

Technische Universität Darmstadt, Department of Law and Economics

10/2022 May



This project has received funding from the European Union Horizon 2020 Research and Innovation action under grant agreement No 822781

Effects of Technological Change and Automation on Industry Structure and (Wage-)Inequality: Insights from a Dynamic Task-Based Model*

Herbert Dawid[†] and Michael Neugart[‡]

October, 2021

Abstract

The advent of artificial intelligence is changing the task allocation of workers and machines in firms' production processes with potentially wide ranging effects on workers and firms. We develop an agent-based simulation framework to investigate the consequences of different types of automation for industry output, the wage distribution, the labor share, and industry dynamics. It is shown how the competitiveness of markets, in particular barriers to entry, changes the effects that automation has on various outcome variables, and to which extent heterogeneous workers with distinct general skill endowments and heterogeneous firms featuring distinct wage offer rules affect the channels via which automation changes market outcomes.

Keywords: automation, artificial intelligence, tasks, wage inequality, firm dynamics

JEL-codes: C63, L10, J31, O33

1 Introduction

Advances in artificial intelligence, data science, automation, or robotics have led to the notion that economies are undergoing a fundamental change with largely uncertain consequences. Radical transformation of the workplaces (Brynjolfsson and McAfee, 2011;

*We are grateful for helpful comments from participants of the 18th Joseph A. Schumpeter Conference and the 27th International Conference on Computing in Economics and Finance. This research has been supported by the European Unions Horizon 2020 grant No. 822781 - Project GROWINPRO.

[†]Corresponding Author: Department of Business Administration and Economics and Center for Mathematical Economics, Bielefeld University. Email: hdawid@uni-bielefeld.de

[‡]Technische Universität Darmstadt, Department of Law and Economics, Hochschulstraße 1, 64289 Darmstadt, Germany. E-mail: michael.neugart@tu-darmstadt.de

Ford, 2015) may lead to the displacement of large fractions of the workforce as tasks will increasingly be done by intelligent machines (Frey and Osborne, 2017). This is one projection of what to expect. Another perspective emphasizes that also this transformation will bring prosperity as it was the case with technological advancements in the past. Among the countervailing forces to the displacement of workers, it has been put forward that automation might augment tasks carried out by worker, thereby increasing their productivity, that total productivity increases may lead to higher aggregate demand thus increasing labor demand, and that it may create new jobs as tasks will emerge that did not exist in the past (Acemoglu and Restrepo, 2018b, 2019; Autor and Salomons, 2018; Vivarelli, 2014).

In this paper, we aim to shed more light on the consequences of automation for economies paying particular attention to the distributional effects of automation. We do this by building a dynamic industry model with task-based production on the firm level where firms decide on whether to allocate workers or machines to different tasks. Firms compete both on the product and the labor market. We are interested in how the allocation decisions of firms change in the presence of different forms of automation that may affect the productivity of machines, the productivity of workers, both, or shift the set of tasks that firms have to do for production. A central feature of our analysis is that we analyze the consequences of the various forms of automation on the distribution of wages between and within firms, and on the functional income distribution. An important aspect of our analysis is that we explore how the effects of automation depends on the competitiveness of the industry. In particular, we distinguish between industries in which high barriers to entry prevent the entry of additional competitors and industries without entry barriers where recurrent entry and exit, triggered by positive respectively negative firm profits, occurs. We examine to which extent and why the competitiveness of markets alters the distributional effects of automation. Moreover, we analyze the role of heterogeneous as opposed to homogeneous agents. We compare the outcomes of automation when workers have homogeneous skills to the outcomes we get with high skill and low skill workers. In addition, we explore a version of our model with two types of firms. One set of firms has high wage setting power, and another set of firms has low wage setting power.

Our analysis, paying particular attention to the distributional consequences in the light of potential market entry and exit of firms, extends existing analyses of the economic effects of automation. Our investigation borrows from task-based models studying the consequences of automation for labor markets and the macroeconomy (Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018a, 2019). Also the distinction between different types of automation addressed in our analysis builds on this stream of literature. However, we consider a dynamic industry framework, where the task allocation decisions are made on the firm level rather than on a macroeconomic level. This allows us to incorporate

competition effects and to more closely pay attention to industry dynamics and the role that heterogeneous workers and firms may play with respect to the economic consequences of automation.

We rely in our analysis on an agent-based modeling approach, which has become a widely used approach for studying issues related to innovation and industrial dynamics from an evolutionary perspective, see, e.g. Nelson and Winter (1982); Dosi et al. (1995); Landini et al. (2017); Dawid et al. (2019), or the survey in Dawid (2006), as well as in labor economics (Neugart and Richiardi, 2018), and macroeconomics (Dawid and Delli Gatti, 2018; Dosi and Roventini, 2019). Several recent contributions have studied different implications of automation using an agent-based approach. Vermeulen et al. (2020) investigate an agent-based model with multiple sectors and occupations. They investigate the effect of productivity increases and the creation of new sectors on employment and wages, and various policies to mitigate the negative consequences of automation on workers. They find that robotization causes the stagnation of wages, but that the emergence of new sectors pushes up wages. Fierro et al. (2021) also base their analysis on a multi-sector model. Workers have heterogeneous skills and firms post vacancies requiring different skills due to an endogenous skill-biased technological change. They show that automation leads to job polarization, and how a minimum wage policy may increase aggregate productivity. Dosi et al. (2021) extend the well established K+S agent-based macroeconomic framework to a multi-sector setting with endogenous occurrence of technological paradigms to study the employment effects of automation. Bertani et al. (2020, 2021) incorporate digital assets, which are produced in a separate sector and increase productivity of consumption good producers into an agent-based macroeconomic model, and show that in such a setting technological unemployment emerges in the long run. More generally, effects of (rapid) technological change on the wage distribution have been studied in numerous agent-based settings, see e.g. Dosi et al. (2017); Dawid et al. (2018); Caiani et al. (2019); Dawid and Hepp (2021).

In contrast to these agent-based studies, our approach features market entry and exit of firms, and various forms of automation that we compare. Moreover, using a task-based approach, we are able to speak to the effects of automation on the wage distribution disentangling the effects on between firm and within firm wage effects. Within our framework, these wage effects can be traced to the firms' decisions on how to allocate machines and workers on tasks, and the firms' decision to enter or leave the market as a consequence of automation that changes profits to be earned.

In this respect, our theoretical analysis relates to empirical studies that explore the consequences of automation on the firm level. Due to data limitations this is still a small literature, that, however, has already interesting insights to offer.¹ Barth et al. (2020)

¹For empirical studies on the aggregate employment effects of automation, see e.g. Acemoglu and Restrepo (2020a,b); Aghion et al. (2020); Graetz and Michaels (2018).

shows that robots have positive effects on average wages, and that wages of high skill workers increase relative to wages of the low skill workers in firms that robotize. Dixon et al. (2021) find that investment in robots increases total employment within firms, which, in their sample of Canadian firms, is a composite effect of higher employment of high skill workers and of low skill workers, and lower employment of middle skill workers. In contrast, Koch et al. (2021), based on Spanish data, find that employment in firms adopting robots increases for all skill groups. Moreover, they show that the labor cost share in firms adopting robots decreases. Domini et al. (2021b) document that the channel through which automation raises firms’ employment growth rate is through a higher hiring rate and a lower separation rate. Interestingly, they also show that automation spikes are not associated with changes in the occupational composition of firms or changes in jobs that can be characterized as routine or non-routine. In a follow-up paper Domini et al. (2021a) again use French firm-level data to explore the effect of automation on within and between firm wage inequality. They find that the largest part of the wage inequality within each sector is due to within firm inequality and that automation spikes have hardly any impact on this inequality.

Our task-based theoretical analysis with heterogeneous worker skills offers economic mechanisms that may help to better understand the empirical effects of automation on the firm level. In particular, we can demonstrate how automation changes the task allocation of firms. This is accompanied by changing wages paid to high skill and low skill workers driven by changes in the productivity of tasks and by changes in the composition of workers doing specific tasks. Moreover, we can show how high skill workers may crowd out low skill workers on lower tasks, or in some instances may withdraw from lower tasks and rather apply for higher tasks as new firms enter the market due to better profit opportunities caused by automation.

We proceed by introducing our task-based simulation framework in Section 2. In Section 3, we parameterize our model and introduce the experiments. Section 4 discusses the results, and Section 5 concludes.

2 The Model

2.1 General Setting and Task-based Production

We consider an industry in which m_t firms offer a homogeneous consumption good. Denoting by $Y_{i,t}$ the output of firm $i \in \{1, \dots, m_t\}$ at time t , the price of the good at t is given by the iso-elastic inverse demand curve

$$p_t = \frac{B}{\sum_i Y_{i,t}}. \quad (1)$$

Here B is the total consumption budget of all consumers in this industry. Since our focus is on the supply side effects of different types of automation, we assume that the consumption budget (in nominal terms) is constant over time. Following an approach similar to Acemoglu and Autor (2011); Acemoglu and Restrepo (2018a,b, 2019), we also assume that different tasks have to be carried out in order to produce the consumption good. More precisely, $\mathcal{X}_t = \{N_t - n + 1, \dots, N_t\}$ denotes the set of tasks at time t , and at each time period each firm has to allocate either a machine or a worker to each task $x \in \mathcal{X}_t$. There is a constant population of J workers who aim to work in this industry and the function $l_{i,t}(x) \in \{0, 1, \dots, J\}$ denotes whether a machine ($l_{i,t}(x) = 0$) or a worker ($l_{i,t}(x) = j \in \{1, \dots, J\}$) carries out task $x \in \mathcal{X}_t$ at firm i in period t . The output of task x at firm i is denoted by $y_{i,t}(x)$ and is given by

$$y_{i,t}(x) = \gamma_L(\min(x, s_{l_{i,t}(x),t}))\mathbf{1}_{[l_{i,t}(x)>0]} + \gamma_M(x)\mathbf{1}_{[l_{i,t}(x)=0]}, \quad (2)$$

where $\mathbf{1}_{[\cdot]}$ denotes the indicator function and $s_{j,t}$ the specific skill level of worker j at time t . Hence, the output of task x is given by $\gamma_L(\min(x, s_{j,t}))$ if task x is filled with worker j and by $\gamma_M(x)$ if task x is filled with a machine. The functions $\gamma_L(\cdot), \gamma_M(\cdot)$ are assumed to be increasing which means that workers and machines generate more output when allocated to a higher tasks. The fact that for tasks carried out by workers the argument of $\gamma_L(\cdot)$ is the minimum of the task level x and the specific skill of the worker ($s_{j,t}$), captures that workers lacking the necessary skills cannot fully fulfill a high level task. Finally, in accordance with Acemoglu and Autor (2011) we assume that the productivity of (sufficiently skilled) workers relative to machines is larger for high level compared to low level tasks. Formally, this means that the ratio $\gamma_L(x)/\gamma_M(x)$ increases with x .² Output of firm i is given by the production function

$$Y_{i,t} = \prod_{x=N_t-n+1}^{N_t} y_{i,t}(x), \quad (3)$$

which captures the complementarity between the different tasks. Workers increase their specific skills at the end of each period in which they have been allocated to a task above their current skills. Hence,

$$s_{j,t+1} = \max[s_{j,t}, x_{j,t}], \quad (4)$$

where $x_{j,t}$ denotes the tasks to which worker j was allocated in period t . In a scenario with two types of workers, characterized by low and high general skills, only workers with high general skills can be allocated to any task, whereas workers with low general skills can never carry out a task above a level x^{max} , and hence firms also cannot allocate low

²In our simulations, we use the functional forms $\gamma_k(x) = \bar{\gamma}_k + \tilde{\gamma}_k x$, $k = L, M$ with parameters chosen in a way that all assumptions are fulfilled.

skill workers to a task $x > x^{max}$.

2.2 Determination of Labor Demand and Labor Market Interactions

Each firm every period develops a plan for the allocation of the different tasks to workers respectively machines. More precisely, each firm calculates for each task $\tilde{x} \in \mathcal{X}_t$ the expected profit $\hat{\pi}_{\tilde{x},t}$ arising if all tasks below \tilde{x} are filled with machines and starting with task \tilde{x} all higher tasks are allocated to workers. This calculation is based on naive expectations about the consumption good and machine prices as well as about wages. Denoting by $p_{m,t}$ the machine price in period t and by $\bar{w}_{x,t}$ the economy wide average wage of workers allocated to task x in t , the estimated profit can be written as

$$\hat{\pi}_{\tilde{x},t} = p_{t-1} \left[\prod_{x=N_t-n+1}^{\tilde{x}-1} \gamma_M(x) \cdot \prod_{x=\tilde{x}}^{N_t} \gamma_L(x) \right] - (\tilde{x} - (N_t - n))p_{m,t-1} - \sum_{x=\tilde{x}}^{N_t} \bar{w}_{x,t-1}.$$

The firm then determines the lowest task it plans to allocate to a worker as $x_{i,t}^* = \arg \max_{N_t-n+1 \leq \tilde{x} \leq N_t} \hat{\pi}_{\tilde{x},t}$. If the firm in its current workforce has workers allocated to tasks below $x_{i,t}^*$ it dismisses these workers. Afterwards, it posts vacancies for all tasks $x \geq x_{i,t}^*$ which are currently not filled with workers. Vacancies are posted stating the level of the task for the advertised position.

The labor market follows a simple search and matching protocol without on-the-job search. Once all vacancies have been posted, every unemployed worker j applies to ξ posted vacancies. If for at least ξ posted vacancies the worker's specific skills exceed the skill requirement of the task (i.e. $s_{j,t} \geq x$ if the posted vacancy is for task x) the worker applies to ξ randomly chosen vacancies among those. Otherwise, the worker applies to all posted vacancies for which she has sufficient skills and chooses the additional positions to apply to randomly among the opening for which she has in-sufficient skills. Workers with low general skills, however, never apply to openings for tasks above x^{max} .

Firms enter negotiations with all applicants for its posted vacancies and for each applicant elicits information about her specific skill level $s_{j,t}$ and her reservation wage $w_{j,t}^{res}$. Based on this, the firm develops a plan for making job offers. As a first step, for each task to be filled the firm calculates the maximal wage it can pay for a worker with sufficient skills for this task such that the resulting firm profit is still higher compared to a scenario where this task is filled by a machine. This calculation is carried out under naive expectations about the machine price and the assumption that apart from this considered task all tasks below $x_{i,t}^*$ are allocated to machines and all above $x_{i,t}^*$ to workers. This gives rise to a maximal wage $w_{i,t}^{max}(x)$ the firm is willing to pay for task x . Then, starting from the highest task to be filled, for each open task x the firm determines the set of yet not allocated applicants who have sufficient specific skills for the considered task

($s_{j,t} \geq x$) and at the same time a reservation wage below the firm's maximal willingness to pay ($w_{j,t}^{res} \leq w_{i,t}^{max}(x)$). From this set of applicants the firm chooses the worker with the highest reservation wage and allocates it to task x . The rationale for this first step in the procedure is that the firm tries to make sure that given the pool of applicants it can fill as many positions as possible. Since the maximal wage the firm is willing to pay increases with x , potential frictions arising from a mismatch between $w_{i,t}^{max}(x)$ and $w_{j,t}^{res}$ are minimized if workers with high reservation wages are allocated to high tasks.

After each open task has been treated in this way, the firm determines the set of all open tasks to which so far no applicant has been allocated. If this set is non-empty then, starting with the highest among these tasks, for each task x the firm selects among the currently unassigned applicants the one with the highest specific skills and allocates her to task x . After a temporary allocation of applicants has been generated in this way, as a final planning step, the firm checks for each open task whether there are non-allocated applicants whose specific skills are above the level of the task and whose reservation wage is lower than that of the currently allocated applicant. If such a case occurs, the applicant with the higher reservation wage is replaced by that with the lower reservation wage in the planned allocation. This final step reduces the expected wage bill of the firm without increasing the frictions. This procedure results in a planned allocation $\hat{l}_{i,t} = \{\hat{l}_{i,t}(x)\}_{x=N_t-n+1}^{N_t}$. Note that for $x < x_{i,t}^*$ always $\hat{l}_{i,t}(x) = 0$ holds.

For every task x with $\hat{l}_{i,t}(x) > 0$ the firm makes a job offer with wage

$$w_{i,t}(x) = \zeta_i w_{j,t}^{res} + (1 - \zeta_i) w_{i,t}^{max}(x) \quad (5)$$

to the applicant $j = \hat{l}_{i,t}(x) \in \{1, \dots, J\}$. The parameter ζ_i indicates the fraction of the jointly generated surplus the firm intends to keep rather than to transfer to the worker. If a worker receives several job offers she accepts the one with the highest wage offer and rejects all other offers. If the job offer of firm i for task x has been accepted we have $l_{i,t} = \hat{l}_{i,t}(x)$, otherwise $l_{i,t} = 0$. Once all applicants have accepted/rejected their offers the labor market closes. All applicants who did not receive job offers remain unemployed in period t and receive unemployment benefits b^{ue} .³ Workers update their reservation wage for $t + 1$ as a weighted average between their reservation wage and their wage (or unemployment) income in t :

$$w_{j,t+1}^{res} = \begin{cases} (1 - \rho) w_{j,t}^{res} + \rho w_{i,j}(x) & \text{if } l_{i,t}(x) = j \\ (1 - \rho) w_{j,t}^{res} + \rho b^{ue} & \text{if } j \in J_t^{UE}, \end{cases}$$

³Although in what follows we refer to such workers as unemployed, they could also accept a position in a different industry. In this case b^{UE} should be interpreted as the expected income of a worker who is not able to find a job in the considered industry. Since we assume that all J workers aim to work in this industry, in any case all such workers apply for open positions in the industry.

where $J_t^{UE} = \{j : l_{i,t}(x) \neq j \ \forall i = 1, \dots, m_t, \forall x \in \mathcal{X}_t\}$ is the set of unemployed workers at t .

2.3 Capital and Consumption Goods Markets, Firm Profits

All tasks which firm i has not filled with workers after the closure of the labor market, either because $\hat{l}_{i,t}(x) = 0$ or because the firm's job offer has been rejected, are allocated to machines. Hence, the demand for machines by firm i in period t is given by $q_{i,t}^m = \sum_{x=N_t-n+1}^{N_t} \mathbf{1}_{[l_{i,t}(x)=0]}$. To abstract from issues of inter-temporal investment planning, we assume that firms rent the required machines every period. There is a competitive rental market for machines with an upward sloping supply curve such that the rental price is

$$p_t^m = \eta \left(\sum_{i=1}^{m_t} q_{i,t}^m \right) + \kappa,$$

where $\eta, \kappa > 0$ are parameters. Since the firm is never rationed on the capital market either a worker or a machine is allocated to each task. Based on the resulting allocation $l_{i,t}$ the firm's output $Y_{i,t}$ is determined by (2) and (3) and the price of the consumption good p_t is given by the inverse demand curve (1). The actual profit of firm i in period t then reads

$$\pi_{i,t} = p_t Y_{i,t} - q_{i,t}^m p_t^m - \sum_{x \notin v_{i,t}^m} w_{i,t}(x).$$

At the end of period t there is a random separation of matches between workers and firms. More precisely, every worker leaves her current firm and becomes unemployed with probability $\alpha > 0$. The worker then applies for new jobs in $t + 1$ and the firm enters the planning for $t + 1$ with the corresponding task not filled.

2.4 Firm Exit and Entry

In what follows, we will consider the effects of different types of automation both in scenarios where entry barriers to the industry are so high that even in cases of positive profits of firms in the industry there is no entry. For this scenario we consider parametrizations where firm profits are always positive and, therefore, there is also no exit, i.e. $m_t = m_0$ for all t . Additionally, we consider scenarios with entry and exit, where we use a reduced form representation of the entry and exit process. More precisely, denoting by $\bar{\pi}_t = \frac{\sum_{i=1}^{m_t} \pi_{i,t}}{m_t}$ the average profit of all firms in period t , we assume that in periods with positive average profits, i.e. $\bar{\pi}_t > 0$, an additional firm enters the industry at the beginning of period $t + 1$ with probability $P^{en} = \frac{\bar{\pi}_t}{\beta^{en} + \bar{\pi}_t}$, where β^{en} is a parameter, which is positively related to the barriers of entry in the industry. In periods with $\bar{\pi}_t < 0$ there is no entry, but with probability $P^{ex} = \frac{|\bar{\pi}_t|}{\beta^{ex} + |\bar{\pi}_t|}$ a randomly selected firm among those with $\pi_{i,t} < 0$ exits the industry.

2.5 Sequence of Events

The following list provides an overview over the different steps and their sequence carried out in every period:

1. Exit and entry of firms (in the exit/entry scenario)
2. Firms carry out their production planning and determine labor demand
3. Firms dismiss excess workers
4. Firms post vacancies
5. Unemployed workers apply to vacancies
6. Firms generate allocation of applicants to open tasks and make wage offers to selected applicants
7. Applicants accept highest wage offer
8. Firms determine their demand for machines and acquire machines at the market clearing rental price
9. Firms produce and sell consumption goods
10. Firm profits are determined
11. Workers update their specific skills and reservation wages
12. Random quitting of workers from their employers

3 Parameters and Experimental Set-up

3.1 Parameters

To calibrate our model, we follow a mixed approach where we, when possible, recur to empirical evidence for specific parameter choices, or choose parameters to replicate macroeconomic indicators. See Table 1 for an overview of the parameters and initial values.

We populate our model with 300 workers. The task range encompasses 11 tasks. We interpret an iteration as one quarter of a year. Choosing a quitting rate of 2.5% then would lead to an average job tenure of about 10 years (if firms were not to dismiss workers for economic reasons), which matches what Auer and Cazes (2000) report for a sample of OECD countries. Studies on earnings losses of displaced workers after reemployment suggest that these workers earn 10% to 25% less (cf. Bartel and Borjas, 1981; Ruhm,

Table 1: Parameters and initializations for baseline model

Parameters and variables		Value
<i>Parameters</i>		
Number of workers	J	300
Applications per worker	ξ	2
Quitting probability	α	0.025
Reservation wage update	ρ	0.08
Wage setting	ζ_i	0.5
Lower bound of tasks	$N_t - n + 1$	10
Upper bound of tasks	N_t	20
Labor productivity	$\bar{\gamma}_l$	1.22
	$\tilde{\gamma}_l$	0.024
Machine productivity	$\bar{\gamma}_m$	1.44
	$\tilde{\gamma}_m$	0.005
Supply curve for machines	η	1
	κ	3
Entry of firms	β^{ent}	50
Exit of firms	β^{ex}	50
Consumption Budget	B	
Unemployment benefit	b^{UE}	5
<i>Initialization</i>		
Number of firms	m_0	25
Reservation wage	$w_{j,0}^{res}$	5
Machine price	p_0^m	0.8

1987; Burda and Mertens, 2001). We use these estimates to set our parameter on the reservation wage update by calculating how high it would have to be so that a displaced worker would earn approximately 20% after an unemployment spell of two iterations when being reemployed. Parameters for labor productivity and machine productivity and for the supply function of machines are chosen to get a task allocation with workers and machines that brings us to an empirically plausible labor share of about 60%. The parameters guiding entry and exit of firms set on a level so that it roughly takes 10 years to erase market profits through entry of new firms - starting from a number of firms (25) that implies mark-ups, calculated as the ratio of total profits in the economy to sales, in the range of 20%.⁴ Mark-ups of this size have been reported, for example, by Bresnahan (1989).

Variables that have to be initialized but are determined endogenously by the model afterwards are the machine price ($p_0^m = 0.8$), the reservation wage ($w_{j,0}^{res} = 5$), and the number of firms ($m_0 = 25$).

⁴Our parametrization implies that at least for the initial number of firms the supply of workers always exceeds labor demand such that rationing of firms on the labor market is entirely due to frictions. We choose this setup to rule out that task reallocation effects of changes in labor respectively machine productivity are crucially affected by labor shortage.

Table 2: Experiments

No.	Description	Change in parameters to
1	Deepening of automation	$\bar{\gamma}_m = 1.4703$
2	Labor augmenting technological advances	$\bar{\gamma}_l = 1.2516$
3	Deepening of automation and labor augmenting technological advances	$\bar{\gamma}_m = 1.4703$ and $\bar{\gamma}_l = 1.2516$
4	Shift in tasks	$N_t = 21$ and $N_t - n + 1 = 11$

3.2 Experiments

Within our framework, we conduct four different experiments related to different types of automation, see Table 2 for a summary. These stylized experiments are meant to capture distinct dimensions of automation in a task-based framework (see also Acemoglu and Restrepo, 2019). Experiment (1) we call “deepening of automation”. It relates to an increase in the productivity of machines with which tasks can be operated. Historically old machines have been replaced with newer vintages that are more productive. In agriculture plows have been replaced by more modern technology where modern technology makes use of global positioning systems, or hand-operated milling machines have been replaced by CNC (computer numerical control) milling machines. We operationalize the experiment by increasing parameter $\bar{\gamma}_m$ so that the productivity of tasks increases by 2% on average. Experiment (2) we call “labor augmenting technological advances”. These are technologies that increase a worker’s output on a task. Examples are augmented reality glasses for the design of new products, exo-skeletons that support workers in executing manual tasks, or software that makes teaching to heterogeneous students more effective. These labor augmenting technological advances correspond to an increase in the productivity of tasks that are operated with workers. This change in productivity relates to $\bar{\gamma}_l$. Again the size is chosen so that the productivity of tasks increases by 2% on average. Experiment (3) is a combination of both changes in automation conducted in Experiments (1) and (2). Finally, Experiment (4) analyzes a “shift in tasks”. This allows us to capture technological advances leading to the creation of new tasks. We have been observing that new tasks or occupations emerged as artificial intelligence has become more powerful. News tasks arose related, for example, to the development or maintenance of these systems. In our experiment, we add one more task, that has a higher productivity than the tasks that existed so far. To keep the number of tasks constant, the hitherto lowest task, i.e. the one with the lowest machine or labor productivity disappears.

4 Results

4.1 Dynamics of the Burn In Phase

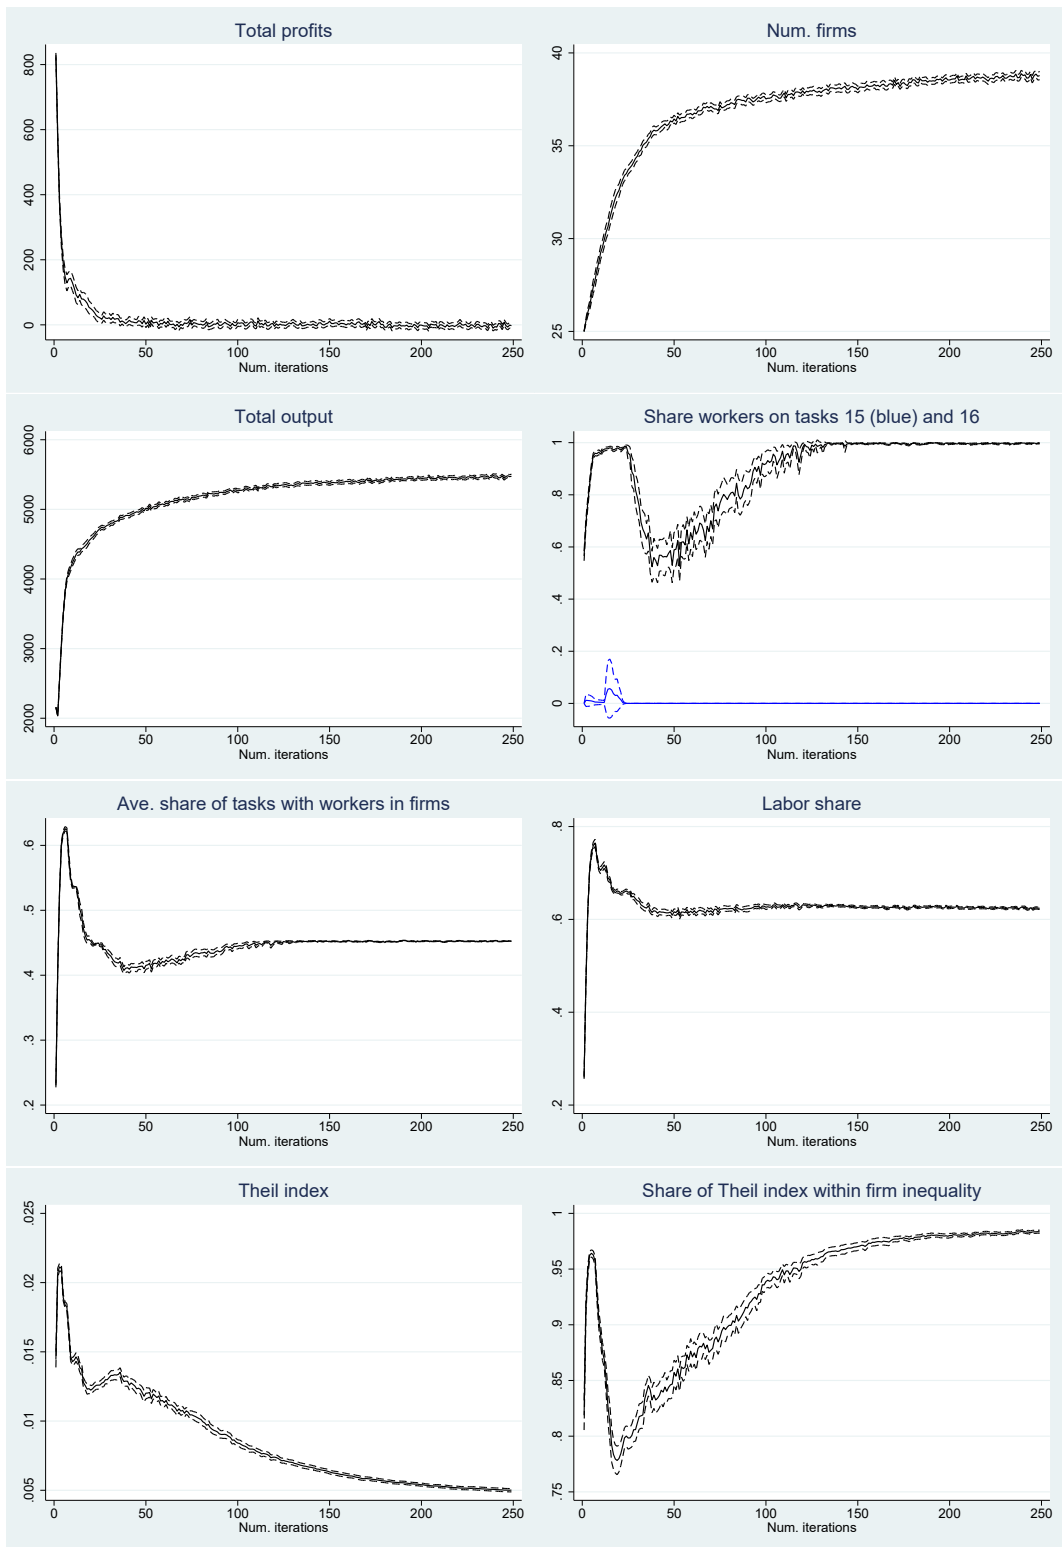
Choosing initial values for the variables of our model and parameters, as summarized in Table 1, without any technological change results in stationary outcomes, as exemplified in Figure 1. The time series shown for various variables relate to the case where we allow for entry and exit of firms and initialize the model with 25 firms (which is the number of firms we keep constant for the case of no exit and entry of firms.) Initially, profits summarized over all firms in the market are positive, triggering firm entry. As the total number of firms settles at 39, total profits in the market become zero. The entry of firms initially drives up real output, which however eventually converges to a stationary level as the number of firms ceases to grow further. Firms partly produce with labor and partly with machines. Lower tasks are operated with machines as exemplified with the share of firms putting workers on Task 15, which becomes zero after a short burn in phase. Task 16, in contrast, is operated by almost all firms with workers. Hence, our model endogenously, based on independent decisions by the different firms, generates a rather sharp distinction between tasks that are carried out by machines and by workers. In the long run, firms allocate a fraction 45% of tasks to workers and the labor share in the industry settles at approximately 62%. Throughout our analysis, we use the Theil index as a measure of wage inequality, as it will allow us to easily decompose the overall wage inequality into the components of a within firm wage inequality and a between firm wage inequality. As can be seen in Figure 1, the Theil index also converges to a stationary level, and most of the wage inequality can be attributed to within firm wage inequality.

4.2 Long-run Effects

The main focus of our analysis is on the longer run effects of automation on various market outcomes. To analyze these effects, we shock the economy at iteration 250 and record outcomes briefly before the shock and 100 iterations after the shock. The results are summarized in box plots, where the coding on the vertical axis relates to the baseline (0), i.e. outcomes before the shock, and the experiments “deepening of automation” (1), “labor augmenting technological advances” (2), the combination of the two (3), and “shifts in tasks” (4). The box plots summarize the outcomes of 50 replications showing the 25th, 50th, and 75th percentile, and the upper and lower adjacent values. Dots refer to outside values.

In particular, we analyze the consequences arising from the various forms of automation for the labor share, real wages, employment, and the wage distribution. We compare the distributional effects when barriers to entry are so high that there is no entry and exit of firms in the market and when entry barriers are low such that there is entry and exit

Figure 1: Time series



Notes: Panels show time series for various outcomes (solid lines) with 95% confidence intervals (dashed lines). Standard errors are calculated on the basis of 50 replications.


of firms in the market. These are two extreme forms of the competitiveness of industries that may, and as we will see, do matter for how automation affects market outcomes. Moreover, our model allows us to analyze the effects of automation on market outcomes along two further dimensions. First, in the model without entry and exit of firms, we introduce firm heterogeneity with respect to their wage setting behavior, cf. parameter ζ_i in our model description. Rather than assuming that all firms have the same wage offer rule, we split firms randomly into one half that wants to keep a larger fraction of the joined surplus rather than transferring it to the workers, and another half that transfers a larger fraction of the joined surplus to the workers via the wage offer. Second, we analyze the outcome of our experiments for scenarios where workers have different general skills. In this scenario, as opposed to our baseline where all workers can acquire necessary specific skills for all tasks, there are high skill workers who can be employed on all tasks of a firm, and low skill workers who can only be assigned to tasks up to x^{max} and do not have the skills to operate higher tasks.

4.2.1 No Market Entry or Exit

Homogeneous Firms We start discussing our results by considering the scenario with no entry and exit of firms and homogeneous agents. Figure 2 summarizes the main findings. With all four forms of automation real output in the economy increases. However, the distributional consequences differ substantially. With respect to the baseline of no automation, we observe that a deepening of automation (Experiment (1)) decreases the labor share and labor augmenting technological advances (Experiment (2)) increase the labor share. If both types of automation occur simultaneously (Experiment (3)) the effects on the labor share more or less compensate each other, whereas a shift in tasks (Experiment (4)) raises the labor share. The signs of the effects of the four experiments on the fraction of tasks allocated to workers are equal to the signs of the effects on the labor share. We can also see that automation has distinct effects on the total profits in the industry. Compared to the baseline, automation deepening, i.e. higher machine productivity, increases total profits, whereas labor augmenting technological change, i.e. higher labor productivity, decreases total profits. When productivity rises for both production factors, total profits stay approximately the same, while they decrease with a shift in tasks. The wage distribution, measured with the Theil index, becomes more equal for a deepening of automation and becomes more unequal for higher labor productivities. The effects of automation on the wage distribution are almost exclusively related to a change in the within firm wage inequality and not to a change in the between firm wage inequality as the decomposition of the Theil index reveals. Average real wages slightly fall with automation deepening and increase with respect to the baseline for all other types of automation.

The observation that real output increases in all four experiments is a straightforward

Figure 2: Treatment effects on market outcomes without market entry and exit of firms



Notes: Panels show the long run effects of the four automation experiments, cf. Table 2, compared to a baseline before the automation shock. Box plots summarize the simulation results of 50 replications. There is no market entry or exit, and firms and workers are homogeneous.

consequence of the higher task productivities. Less obvious are the observations that the positive output effect is much stronger for labor augmenting technological advances compared to automation deepening (although our experiments are set up in a way that the average productivity increase on each task is the same in Experiments (1) and (2)), and that only automation deepening leads to an increase in firm profits, whereas an increase in labor productivity results in lower profits compared to the baseline. To gain an intuitive understanding of these observations, it should be noted that the decision whether a task is (planned to be) allocated to a machine or a worker is essentially based on the ratio between productivity and costs (i.e. machine price respectively expected wage) for the two options. For tasks close to the marginal task⁵, where the allocation switches from machines to workers, this ratio is approximately identical under both allocations, but both productivity and costs are lower for a machine compared to a worker allocation. If machine productivity increases and hence firms re-allocated one (or several) tasks from workers to machines this reallocation implies a decrease in output (due to the lower productivity of machines) and in firm costs. In contrast, in the case of an increase of labor productivity the induced reallocation of tasks to workers induces a further increase of output (in addition to the increases triggered by the pure effect of a productivity increase) and at the same time an increase in costs. Hence, total output increase is larger under labor augmenting technological advances compared to automation deepening. The implications for firm profits are then driven by the resulting price effect. The stronger increase in output under labor augmenting technological change induces a stronger decrease in the price of the consumption good. Since at the same time the costs of the firms increase this induces a reduction in profits. In contrast, under automation deepening the reallocation of tasks to machines has a positive impact on the consumption good price and therefore induces an increase of firm profits.⁵

The effects on the average real wages in the industry are the result of countervailing forces. As described above, a higher machine productivity reduces the average real wage. As firms substitute labor for machine, labor demand declines which exerts downward pressure on wages (labor demand effect). At the same time average real wages (net of a demand effect) increase because only those workers operating better paid higher tasks remain employed (composition effect). As workers are spread over fewer tasks also wage inequality declines. For a higher labor productivity the composition effect and the labor demand effect have opposite signs when compared to the experiment with a higher machine productivity. Labor demand increases pushing up wages. At the same time workers are spread over a larger set of tasks including less well paid tasks, which reduces average wages (net of the labor demand effect). As the labor demand effect is far stronger

⁵It should be noted that under our iso-elastic demand curve with a price-elasticity of 1 a uniform increase in productivity across all firms without any induced task reallocation would not have any effects on revenues and profits since the price would decrease with the same factor with which output increases.

than the composition effect, real wages increase. Moreover, wage inequality increases as workers operate less well paying tasks that formerly have been done by machines.

Heterogeneous Wage Bargaining So far we have shown the consequences of various forms of automation when firms are homogeneous with respect to how wages are determined. We proceed by altering our assumption on wage setting to investigate whether heterogeneity of firms in this respect changes the obtained results. We randomly divide the population of firms into two groups. Half of the firms have a wage setting parameter of $\zeta_i = 0.25$, whereas the remaining firms have $\zeta_i = 0.75$. In the former case, firms on average offer higher wages compared to the latter case, see eq. (5).

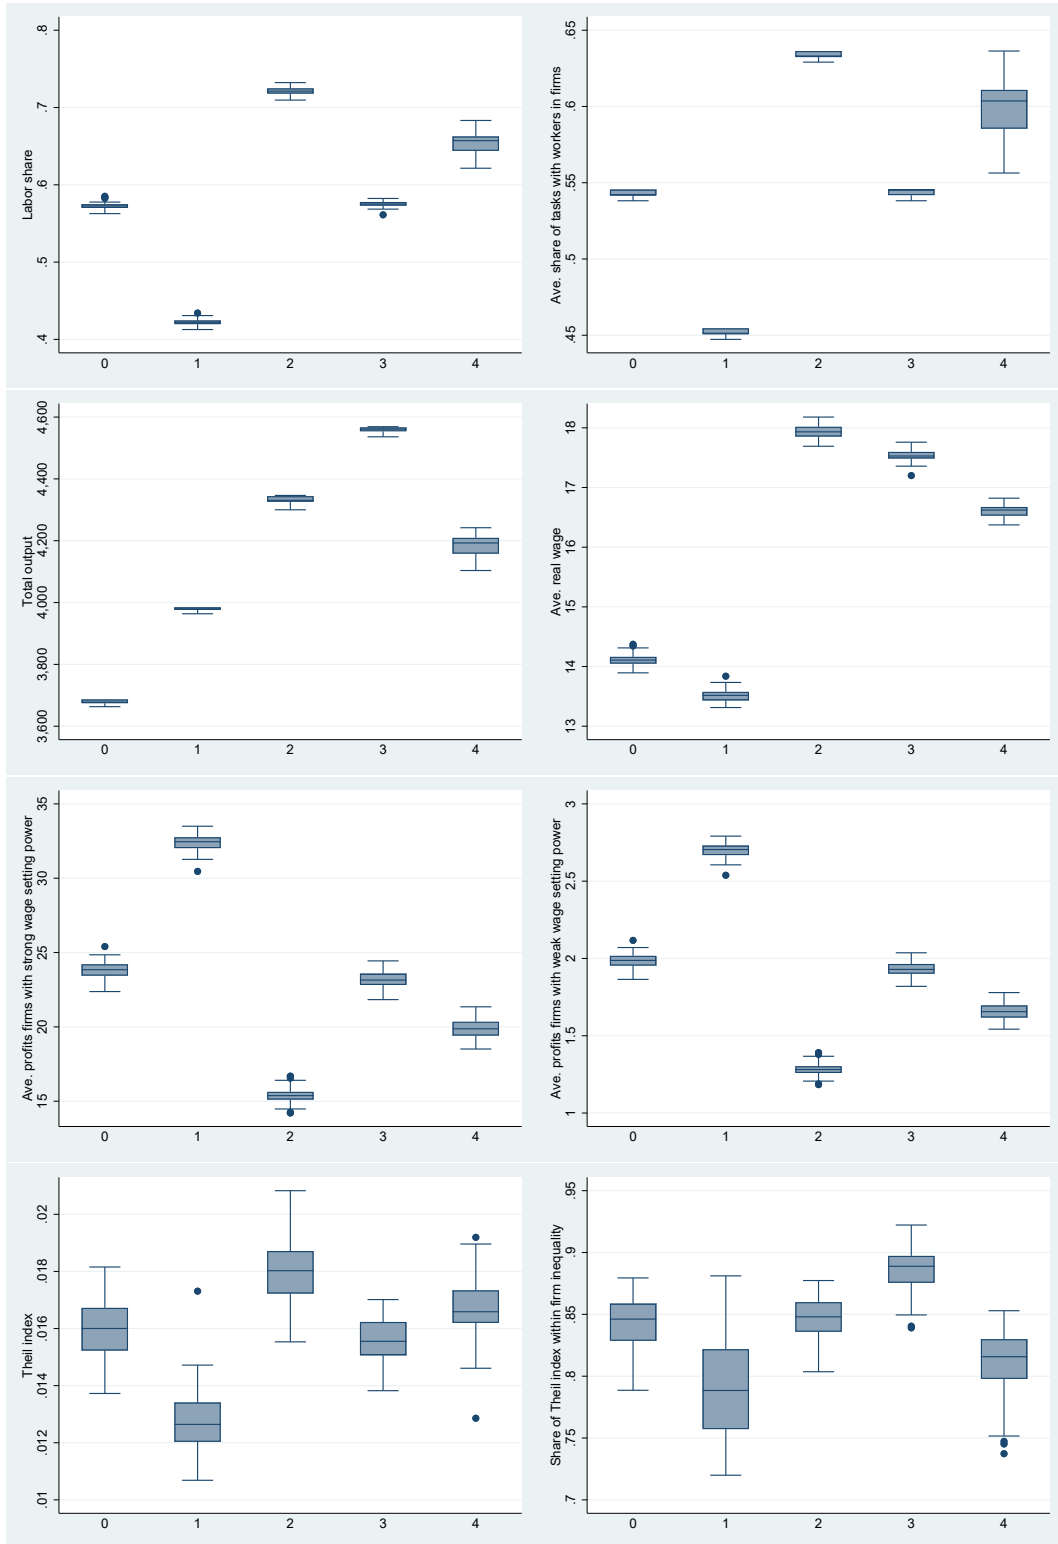
Compared to the case with homogeneous firms, we observe that total real output, employment, and the labor share are hardly affected by heterogeneous wage setting rules across firms (see Figure 3). Since half of the firms pay lower wages than the rest of the firms, these firms with lower wage costs make significantly higher profits. Average real wages are only slightly affected by the heterogeneous population of firms compared to a homogeneous population. As to be expected, the share of wage inequality (measured again by the Theil index) attributable to the between firm wage inequality increases substantially. How automation affects size of wage inequality and the share of within firm wage inequality in relation to total inequality is, however, unaffected by introducing heterogeneous firms.

Heterogeneous general skills of workers Next, we consider the effect of heterogeneity on the worker side. More precisely, we assume that workers have heterogeneous general skills (and firms are homogeneous again). As explained in Section 2.1, workers with high general skills can carry out and acquire the necessary specific skills for all task, whereas workers with low skills can only carry out tasks up to level x^{max} . In our experiment, the share of high skill workers is 30% and we set $x^{max} = 18$, where the highest possible task is $N_t = 20$ in Experiments (1) - (3) and $N_t = 21$ for $t \geq 250$ in Experiment (4).

The effects of automation with heterogeneous general skills on output, tasks assigned to workers, labor share, total profits, average real wages, and wage inequality are almost unchanged (figures are not shown) when compared to the case of homogeneous skills. However, the two skill groups are differently affected by automation. This becomes evident as we look into wages and employment by skill group (Figure 4). As to be expected the wages of the high skill workers relative to the low skill workers are always higher and so are relative employment rates. Different types of automation, however, affect these ratios differently.

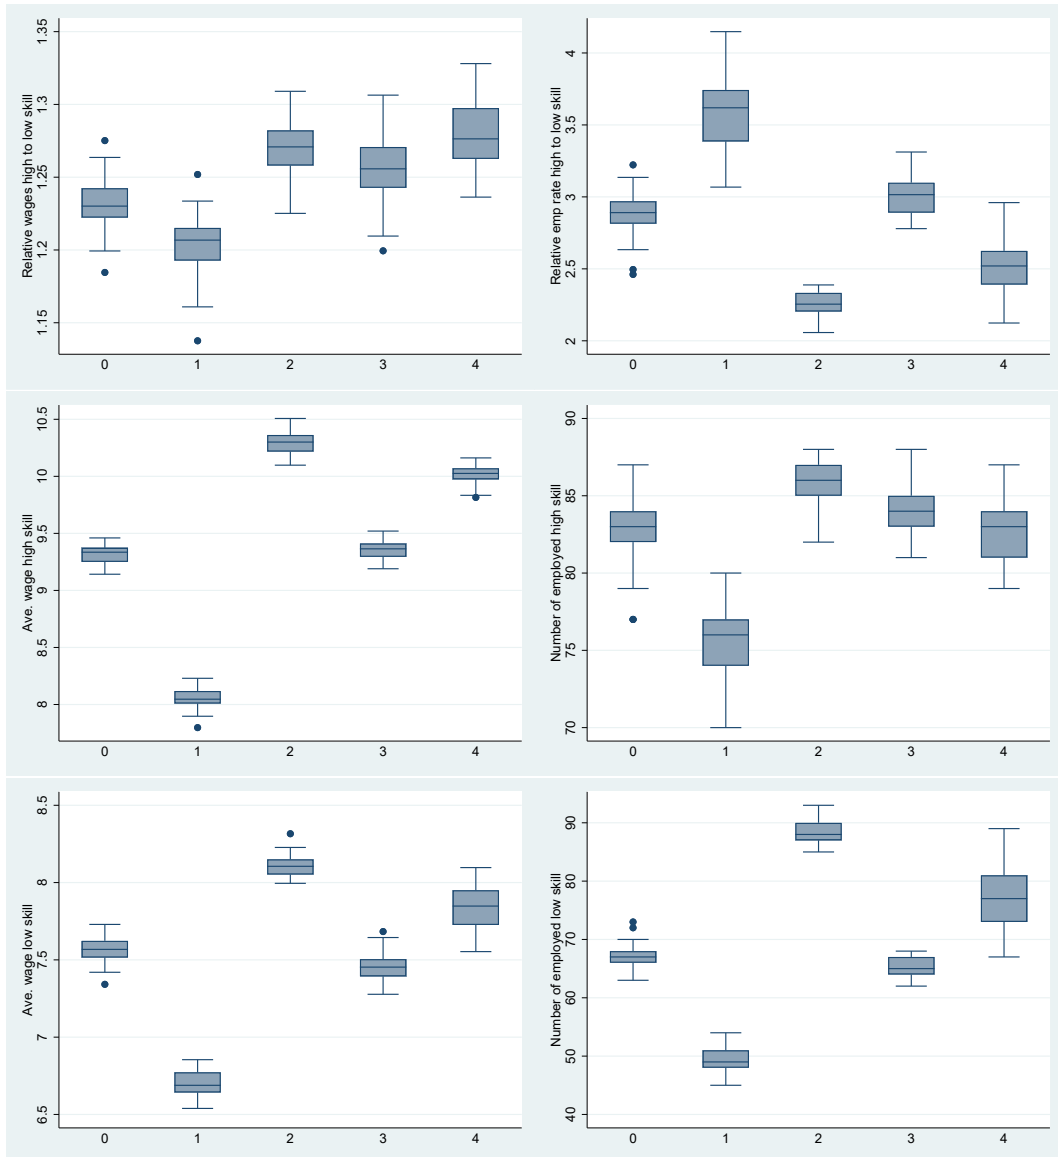
Automation deepening, i.e. an increase in machine productivity decreases the wage ratio of the high skill workers relative to the low skill workers, while increasing the relative employment rate. This is because the reallocation of tasks to machines triggered by

Figure 3: Treatment effects on market outcomes without entry and exit of firms and with heterogeneous wage bargaining



Notes: Panels show the long run effects of the four automation experiments, cf. Table 2, compared to a baseline before the automation shock. Box plots summarize the simulation results of 50 replications. There is no market entry or exit, firms are heterogeneous with respect to their wage offer rule, and workers are homogeneous.

Figure 4: Treatment effects on market outcomes without market entry and exit of firms and heterogeneous general skills of workers



Notes: Panels show the long run effects of the four automation experiments, cf. Table 2, compared to a baseline before the automation shock. Box plots summarize the simulation results of 50 replications. There is no market entry or exit, firms are homogeneous, and workers differ with respect to general skills.

automation deepening mainly affects relatively low tasks. While employment decreases for both skill groups, the employment effect is stronger for the low skill workers for whom fewer tasks exist to apply for. Workers with high general skills are also less employed, but the effect on their employment level is substantially smaller since only a fraction of high skill workers is allocated to low level tasks. Due to a composition effect, the average wage of the employed low skill workers increases relative to that of the high skill workers, where, like in the baseline, all wages decrease in real terms.

Higher labor productivity due to labor augmenting technological advances triggers exactly the opposite effects, since here tasks are reallocated from machines to workers. Both types of workers gain from higher labor productivity which translates into higher wages. For the low skill workers, however, as more tasks are done by workers, average wages do not increase by so much. This composition effect is almost absent for the high skill workers as they are still running the higher tasks. Hence the ratio of average wages of high skill and low skill workers increases. For a better understanding of the effect on the relative employment rate, we can disentangle the high skill workers' employment into employment on high tasks and low tasks (not shown in figures). What we observe is that high skill workers' employment on high tasks stays fairly constant. High skill workers' and low skill workers' employment on lower tasks increases with higher labor productivity. So both groups gain in terms of employment, but low skill workers gain more so that the relative employment rate decreases.

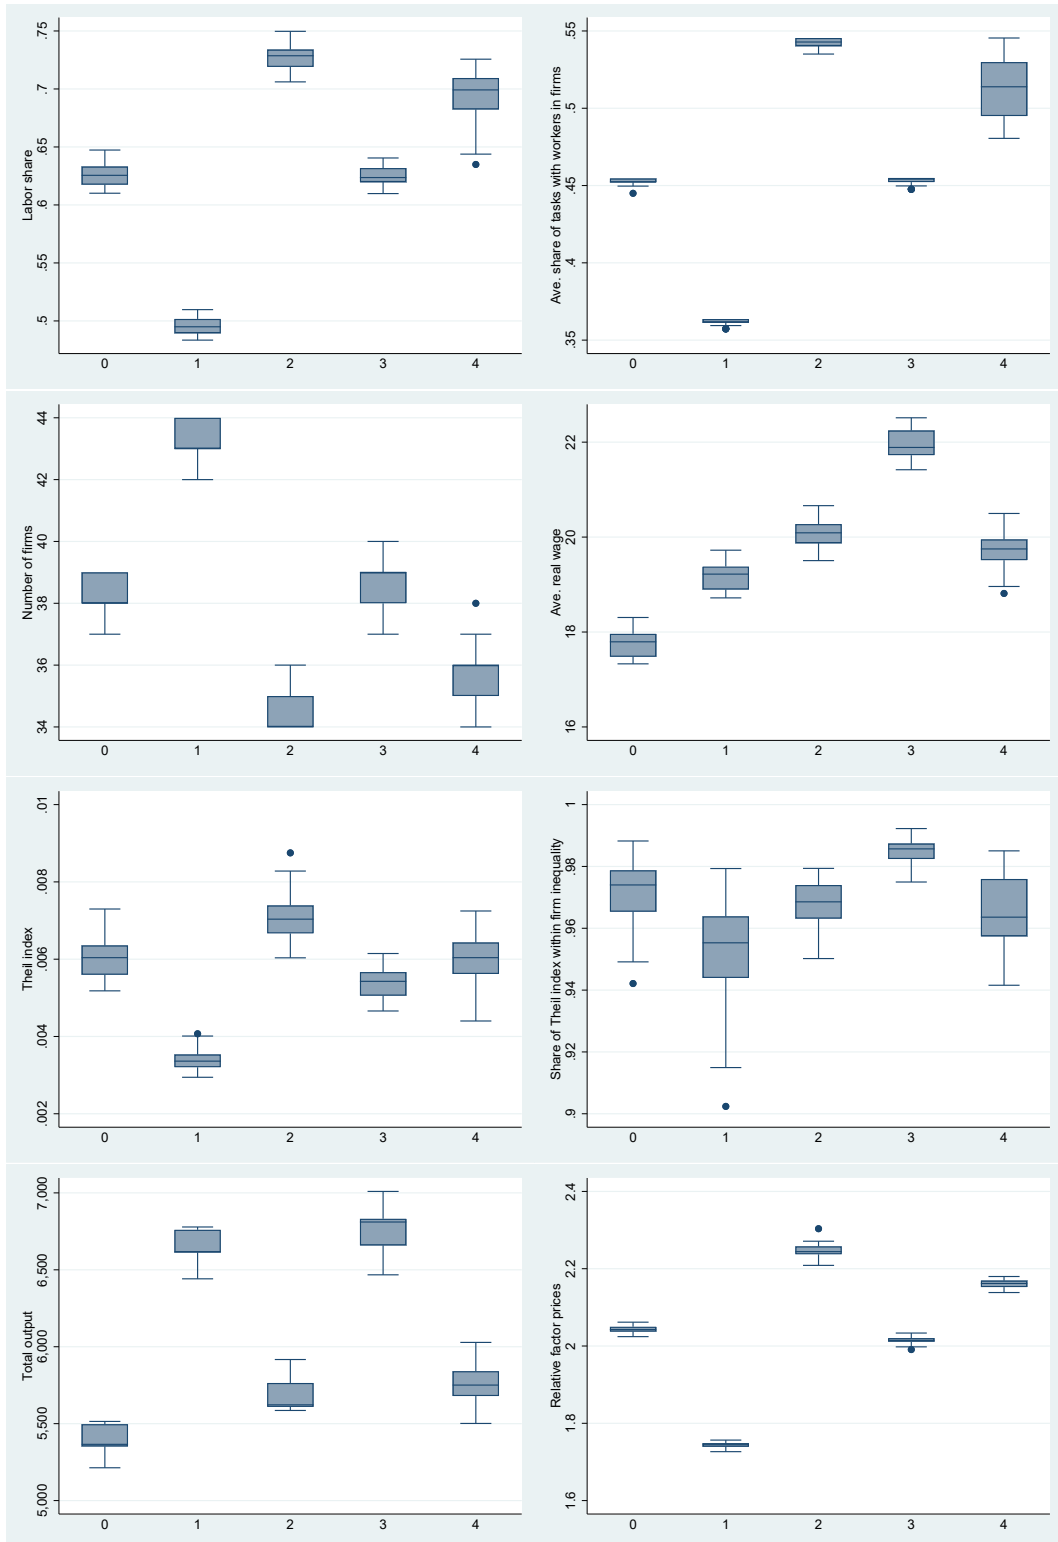
For automation relating to higher labor and machine productivity, i.e. Experiment (3), the two effects almost cancel out. For a shift in tasks, the mechanisms described in relation to the experiment on a higher labor productivity are prevailing. There are more tasks at the upper part of the task range which are done by high skill workers. Some of these high skill workers were assigned to lower tasks before. Thus employment of the high skill workers hardly changes, and employment of the low skill workers increases as their is less competition from high skill workers for the lower tasks.

4.2.2 Market Entry and Exit

Homogeneous workers We turn to a discussion of our results for the scenario with entry and exit of firms and homogeneous agents now. These results are summarized in Figure 5. Compared to an industry with high barriers to entry for new firms, see the results in Figure 2, we observe that in the baseline (without technology shocks) the labor share is higher, fewer tasks are allocated to workers, total output is higher, real wages are higher, and wage inequality is reduced.

Without entry barriers the positive firm profits, which arise under the initial (and fixed) number of firms (this is the same number of firms present in the runs underlying Figure 2) triggers entry by additional firms. This increases the production capacity and total output in the industry compared to the scenario without entry and exit. The new

Figure 5: Treatment effects on market outcomes with entry and exit of firms



Notes: Panels show the long run effects of the four automation experiments, cf. Table 2, compared to a baseline before the automation shock. Box plots summarize the simulation results of 50 replications. There is market entry and exit, firms and workers are homogeneous.

firms demand labor and thereby push up total employment and at the same time the consumption good price decreases due to higher total output. Together this drives up real wages. Since labor becomes relatively more expensive the average number of tasks in a firm carried out by workers goes down. The wage effect dominates leading to a higher labor share in the scenario with exit and entry. As expected, firms' profits go down to zero (not shown in Figure 5). Since firms operate with fewer workers on their task range, wage inequality declines because of the resulting composition effect.

Considering the effects of our four automation experiments, we observe that interestingly the pattern of changes of total output and average real wages across experiments is different compared to the no entry and exit case. Under automation deepening average real wages are higher compared to the baseline now, whereas the opposite effect occurs without entry and exit. Furthermore, again in contrast to the case without exit and entry, total real output is higher with an increase in machine productivity compared to the total real output after an increase in labor productivity.

Both of these observations are driven by the different changes in the number of firms induced under automation deepening compared to labor augmenting technological change. Whereas the former leads to the entry of additional firms, the latter results in a smaller long-run number of firms compared to the baseline without a technology shock. The mechanism underlying this difference is essentially the same that induces an increase (decrease) of profits due to high capital (labor) productivity in the case without exit and entry. The fact that there is a larger long-run number of firms in the industry with automation deepening than with labor augmenting technological change is the dominant factor in explaining why total output in the industry now is larger in the scenario with a capital productivity increase compared to a labor productivity increase. The higher output results in a lower price of the consumption good and this price decrease is sufficiently strong to induce higher real wages under automation deepening compared to the baseline although nominal wages (not shown here) are actually lower. The change in nominal wages is the result of the interplay of two effects: a composition effect that increases wages because fewer (and mainly high level) tasks are done by workers and a demand effect. This second effect is negative, the effect of substitution of labor with capital on the firm level outweighs the additional labor demand due to firm entry. Under labor augmenting technological change both the composition effect and the demand effect have opposite signs compared to automation deepening. Again the demand effect dominates such that already nominal wages are larger compared to the baseline. The fact that the consumption good price is lower, due to higher output, compared to the baseline, induces that the positive effect of an increase in labor productivity on real wages is even stronger.

As mentioned above, the level of wage inequality in this scenario with entry and exit is substantially lower compared to the case with high barriers to entry. This is mainly driven by a composition effect since in the scenario with entry and exit firms on average

fill a smaller set of tasks with workers. The fact that the reduction in inequality is mainly due to a reduction in within firm inequality becomes apparent from the observation that the fraction of the Theil index stemming from within firm inequality is lower in Figure 5 compared to Figure 2. Considering the effect of the different types of automation on wage inequality and its distribution on within and between firm components, the qualitative findings obtained without exit and entry, however, carry over also to the scenario with no entry barriers.

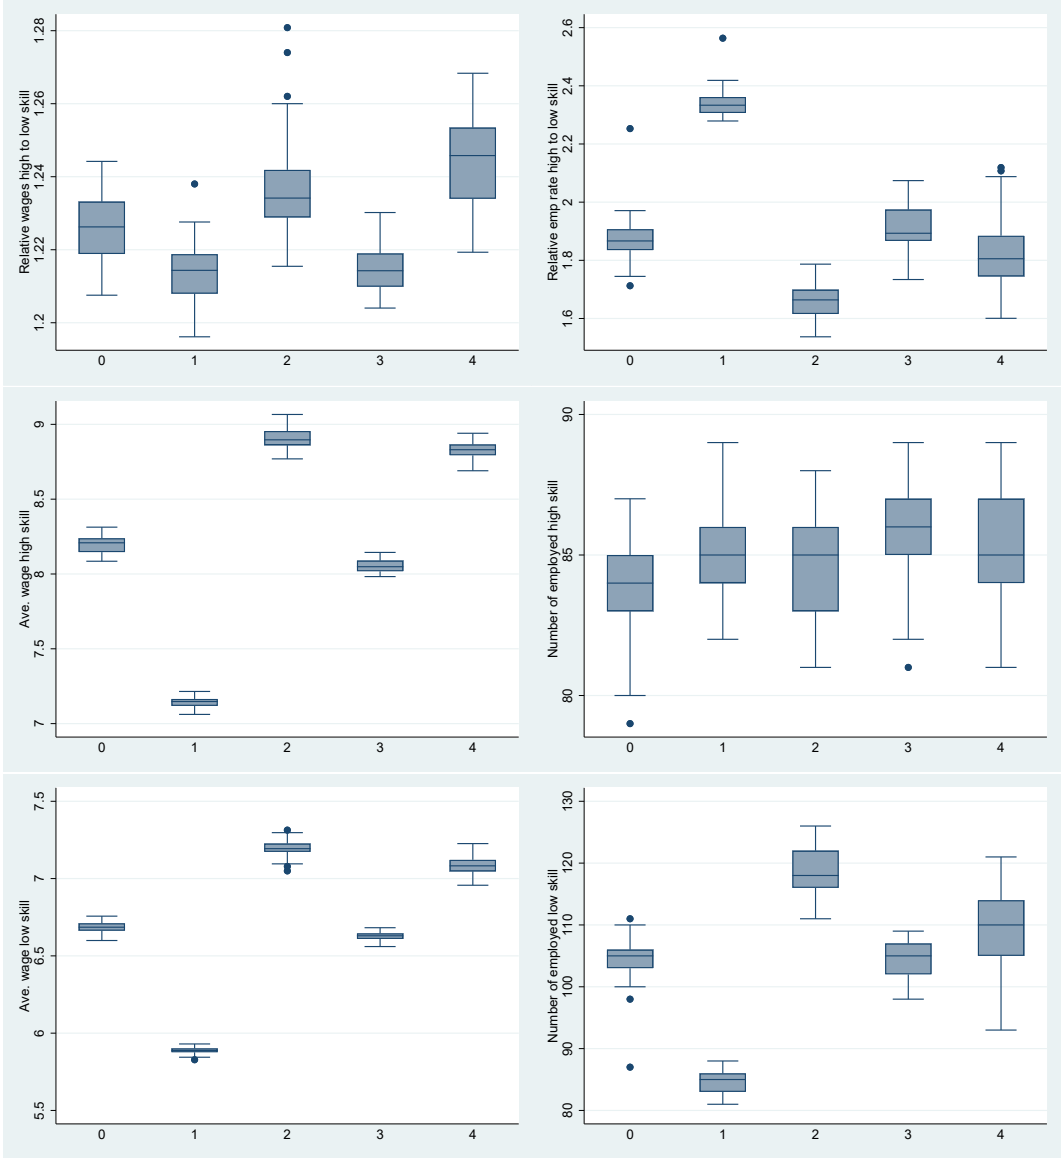
Heterogeneous General Skills of Workers Finally, we turn to the case of a competitive economy where firm can enter and exit the market, and workers' general skills are heterogeneous.⁶ Again, we compare the effects of the automation shocks with those observed in the no entry and exit scenario with heterogeneous skills analyzed earlier in Section 4.2.1.

As we have already observed for the case with free entry and homogeneous general skills, also in the case with heterogeneous skills total real output and the labor share increase as firms enter the market. Wage inequality is lower for all types of automation. Furthermore, allowing for entry and exit of firms changes the effects that automation deepening or labor augmenting technological change have on the average real wages and total output. The qualitative pattern, however, is the same as for the case of homogeneous general skills of workers, and so are the underlying mechanisms. To save space we do not illustrate these observations with separate figures.

What qualitatively differs is the effect of a higher labor or machine productivity on employment of the high skill workers. Figure 6 summarizes these findings. Without entry and exit of firms, high skill employment decreases with more productive machines. With entry of firms, high skill employment is slightly increasing with more productive machines. The underlying mechanisms explaining the observation is related to the entry of new firms. These new firms employ additional high skill workers on their high tasks. This overcompensates the fewer employment opportunities of the high skill workers on lower tasks, which is associated with a higher machine productivity. In fact, we can check (not shown in the figures) that the employment rate of the high skill workers and the low skill workers on lower tasks falls. While employment of the high skill workers increases with higher labor productivity in the scenario without entry and exit, it hardly changes in the scenario with an endogenous number of firms. Here firms exit the market after an increase in labor productivity, which worsens the employment opportunities of the workers. This counteracts the positive effect on employment chances that comes with the firms' decision to run more tasks with workers (who have become more productive).

⁶We do not present experiments with heterogeneous firms in the scenario with entry and exit, since in our setting due to the profit differences between the types of firms in the long run only one type of firm prevails.

Figure 6: Treatment effects on market outcomes with market entry and exit and heterogeneous general skills for workers



Notes: Panels show the long run effects of the four automation experiments, cf. Table 2, compared to a baseline before the automation shock. Box plots summarize the simulation results of 50 replications. There is market entry and exit, firms are homogeneous, and workers differ with respect to general skills.

5 Conclusion

In this paper, we study the effects of different types of automation on output, the wage distribution and the labor share in a dynamic industry model with a task-based production technology on the firm level. We explicitly model the (adaptive) process through which firms operating in an uncertain environment with labor market frictions allocate tasks to workers respectively machines and also consider the implications of high respectively low barriers to entry for the effect of automation. Furthermore, we capture complementarity between task levels and required worker skills.

Our analysis yields several key insights. First, we demonstrate that in spite of the considered frictions and myopic decision rules of the firms a large degree of coordination evolves in the industry with workers being allocated to tasks above a threshold, which is almost uniform across firms. Second, in the presence of high barriers to entry into the industry, labor augmenting technological change, i.e. an increase in labor productivity, has stronger output expanding effects compared to a comparable increase in machine productivity, which we refer to as automation deepening. In the absence of entry barriers this relationship is reversed. Third, whether automation deepening results in a decrease or an increase of real wages depends on the size of entry barriers. Finally, we show that the different types of automation have qualitatively different effects on wage level and inequality based on the interplay of composition effects and labor demand effects. These qualitative findings on the industry level are shown to be robust with respect to heterogeneities of firms and workers, where the impact of automation is, however, substantially different for workers with different general skills.

More generally, our analysis highlights in an evolutionary setting, in which firms and workers adaptively react to technological change, how the qualitative effects of different types of automation depend on industry characteristics, like barriers to entry and agent heterogeneity. Our analysis has empirically testable implications, in particular with respect to the importance of composition effects as drivers of changes in wage inequality. Our results suggest that the evolution of within firm wage inequality plays an important role in determining distributional effects of automation. From a policy perspective these insights have important implications for the type of automation a policy maker who is concerned about distributional effects should foster. Also, they show that an active industrial policy aiming to reduce barriers to entry helps to avoid negative (real) wage effects of automation.

Whereas in this paper we have focused on the implications of a single exogenous technology shock, the developed framework also is suitable for expanding our analysis to repeated automation steps and also for considering firm actions, like training measures, which potentially accompany technological change and automation. Furthermore, due to our focus on task allocation decisions of firms, in this paper, we have abstracted from

strategic quantity decisions of firms, assuming that the scale of the production process of each firm is fixed. Incorporating strategic quantity decisions by firms into our framework could shed additional light on the implications of automation on industry dynamics. Considering these additional aspects is beyond the scope of the current paper, but opens many avenues for future work.

References

- ACEMOGLU, D. AND D. AUTOR (2011): “Skills, tasks and technologies: Implications for employment and earnings,” in *Handbook of Labor Economics*, ed. by D. Card and O. Ashenfelter, Elsevier, vol. 4, 1043–1171.
- ACEMOGLU, D. AND P. RESTREPO (2018a): “Low-skill and high-skill automation,” *Journal of Human Capital*, 12, 204–232.
- (2018b): “The race between man and machine: Implications of technology for growth, factor shares, and employment,” *American Economic Review*, 108, 1488–1542.
- (2019): “Artificial intelligence, automation, and work,” in *The Economics of Artificial Intelligence*, ed. by A. Agrawal, J. Gans, and A. Goldfarb, University of Chicago Press, chap. 8.
- (2020a): “Competing with robots: Firm-level evidence from France,” *AEA Papers and Proceedings*, 110, 383–388.
- (2020b): “Robots and jobs: Evidence from US labor markets,” *Journal of Political Economy*, 128, 2188–2244.
- AGHION, P., C. ANTONIN, S. BUNEL, AND X. JARAVEL (2020): “What are the labor and product market effects of automation? New evidence from France,” CEPR Discussion Paper No. DP14443.
- AUER, P. AND S. CAZES (2000): “Resilience of the long-term employment relationship: Evidence from the industrialized countries,” *International Labour Review*, 139, 379.
- AUTOR, D. AND A. SALOMONS (2018): “Is automation labor share-displacing? Productivity growth, employment, and the labor share,” *Brookings Papers on Economic Activity*, 1–63.
- BARTEL, A. P. AND G. J. BORJAS (1981): “Wage growth and job turnover: An empirical analysis,” in *Studies in Labor Markets*, University of Chicago Press, Chicago, 65–90.
- BARTH, E., M. ROED, P. SCHØNE, AND J. UMBLIJS (2020): “How robots change within-firm wage inequality,” IZA Discussion Paper No. 13605.

- BERTANI, F., L. PONTA, M. RABERTO, A. TEGLIO, AND S. CINCOTTI (2020): “The productivity and unemployment effects of the digital transformation: an empirical and modelling assessment,” *Review of Evolutionary Political Economy*, 1, 329–355.
- (2021): “The complexity of the intangible digital economy: an agent-based model,” *Journal of Business Research*, 129, 527–540.
- BRESNAHAN, T. F. (1989): “Empirical studies of industries with market power,” in *Handbook of Industrial Organization*, Elsevier, vol. 2, 1011–1057.
- BRYNJOLFSSON, E. AND A. MCAFEE (2011): *Race against the machine: How the digital revolution is accelerating innovation, driving productivity, and irreversibly transforming employment and the economy*, Lexington, Massachusetts: Digital Frontier Press.
- BURDA, M. C. AND A. MERTENS (2001): “Estimating wage losses of displaced workers in Germany,” *Labour Economics*, 8, 15–41.
- CAIANI, A., A. RUSSO, AND M. GALLEGATI (2019): “Does inequality hamper innovation and growth? An AB-SFC analysis,” *Journal of Evolutionary Economics*, 29, 177–228.
- DAWID, H. (2006): “Agent-based models of innovation and technological change,” in *Handbook of Computational Economics*, ed. by L. Tesfatsion and K. Judd, Elsevier, vol. 2, 1235–1272.
- DAWID, H. AND D. DELLI GATTI (2018): “Agent-Based Macroeconomics,” in *Handbook of Computational Economics*, ed. by C. Hommes and B. LeBaron, Elsevier, vol. 4 of *Handbook of Computational Economics*, 63–156.
- DAWID, H., P. HARTING, AND M. NEUGART (2018): “Cohesion policy and inequality dynamics: Insights from a heterogeneous agents macroeconomic model,” *Journal of Economic Behavior and Organization*, 150, 220–255.
- DAWID, H., P. HARTING, AND S. VAN DER HOOG (2019): “Manager remuneration, share buybacks and firm performance,” *Industrial and Corporate Change*, 28, 681–706.
- DAWID, H. AND J. HEPP (2021): “Distributional effects of technological regime changes: hysteresis, concentration and inequality dynamics,” *Review of Evolutionary Political Economy*, forthcoming.
- DIXON, J., B. HONG, AND L. WU (2021): “The robot revolution: Managerial and employment consequences for firms,” *Management Science*, 67, 5586–5605.
- DOMINI, G., M. GRAZZI, D. MOSCHELLA, AND T. TREIBICH (2021a): “For whom the bell tolls: the firm-level effects of automation on wage and gender inequality,” GROWINPRO Working Paper 28/2021.

- (2021b): “Threats and opportunities in the digital era: automation spikes and employment dynamics,” *Research Policy*, 50, 104137.
- DOSI, G., O. MARSILI, L. ORSENIGO, AND R. SALVATORE (1995): “Learning, market selection and the evolution of industrial structures,” *Small Business Economics*, 7, 411–436.
- DOSI, G., M. NAPOLETANO, A. ROVENTINI, AND T. TREIBICH (2017): “Micro and macro policies in the Keynes+Schumpeter evolutionary models,” *Journal of Evolutionary Economics*, 27, 63–90.
- DOSI, G., M. C. PEREIRA, A. ROVENTINI, AND M. E. VIRGILLITO (2021): “Technological paradigms, labour creation and destruction in a multi-sector agent-based model,” *GROWINPRO Working Paper 18/2021*.
- DOSI, G. AND A. ROVENTINI (2019): “More is different... and complex! The case for agent-based macroeconomics,” *Journal of Evolutionary Economics*, 29, 1–37.
- FIERRO, L. E., A. CAIANI, A. RUSSO, ET AL. (2021): “Automation, job polarisation, and structural change,” Working Paper 2021/09, Universitat Jaume I.
- FORD, M. (2015): *Robots: technology and the threat of a jobless future*, New York: Basic Books.
- FREY, C. B. AND M. A. OSBORNE (2017): “The future of employment: How susceptible are jobs to computerisation?” *Technological Forecasting and Social Change*, 114, 254–280.
- GRAETZ, G. AND G. MICHAELS (2018): “Robots at work,” *The Review of Economics and Statistics*, 100, 753–768.
- KOCH, M., I. MANUYLOV, AND M. SMOLKA (2021): “Robots and firms,” *The Economic Journal*, 131, 2553–2584.
- LANDINI, F., K. LEEC, AND F. MALERBA (2017): “A history-friendly model of the successive changes in industrial leadership and the catch-up by latecomers,” *Research Policy*, 46, 431–446.
- NELSON, R. AND S. WINTER (1982): *An evolutionary theory of economic change*, Harvard University Press, Cambridge, MA.
- NEUGART, M. AND M. RICHIARDI (2018): “Agent-based models of the labor market,” in *The Oxford Handbook of Computational Economics and Finance*, Oxford University Press, 667–687.

- RUHM, C. J. (1987): “The economic consequences of labor mobility,” *Industrial and Labor Relations Review*, 41, 30–42.
- VERMEULEN, B., A. PYKA, AND P. P. SAVIOTTI (2020): “Robots, structural change, and employment: Future scenarios,” in *Handbook of Labor, Human Resources and Population Economics*, ed. by K. F. Zimmermann, Springer International Publishing.
- VIVARELLI, M. (2014): “Innovation, employment and skills in advanced and developing countries: A survey of economic literature,” *Journal of Economic Issues*, 48, 123–154.