

Training Time, Robots and Technological Unemployment*

Fenicia Cossu

Alessio Moro

University of Cagliari

University of Cagliari, CEPR and IZA

Michelle Petersen Rendall

Monash University and CEPR

April 29, 2026

Abstract

We document that labor training requirements for high-skilled occupations increased in the U.S. from 2006 to 2019. These greater training requirements reduce the extent to which workers displaced from shrinking occupations can relocate to expanding (high-skilled) occupations, thus potentially affecting both the equilibrium occupational structure and the unemployment level. We build a quantitative model in which labor is displaced by task-replacing technological change embodied in robots (“tasks shock”) and the extent of occupational switching depends on the destination occupations’ training requirements. We find that: (i) task-displacing technological change increases steady-state unemployment, but it *reduces* unemployment along the transition; (ii) in contrast, a comparable shock to capital embodied technological change produces larger unemployment rates with respect to the tasks shock, both in the transition and the steady state; and (iii) greater training requirements in high-skilled occupations increase steady-state unemployment and affect the occupational structure along the transition, but the shape of their effect depends on the size of the technological shock.

JEL Classification: J23, J24, O33, J63, C68.

Keywords: Occupational Sorting, Robots, Technological Unemployment.

*We thank Fabio Cerina, Georg Duernecker, Elisa Keller, Ramon Marimon, Alexander Monge-Naranjo, Pontus Rendahl, Andrew Rendall, Florian Trouvain, Gianluca Violante and seminar participants at Auckland, Princeton, UT Sydney, Victoria University Wellington, St. Gallen, XX RIDGE forum, PSE Macro Days 2024, Spanish Macro Network 2024, “The Economics of AI and Automation” (Bicocca), BSE Summer Forum, Fabulous Macro Conference (St. Gallen), Sailing the Macro Workshop (Siracusa), NUS Macroeconomics Workshop, X Workshop on Structural Transformation and Macroeconomic Dynamics (Cagliari), 48th SAEe (Salamanca) and SASCA Conference (Sassari) for the useful comments. We acknowledge financial support under the Italian National Recovery and Resilience Plan funded by NextGenerationEU, “The Role of Skill-Biased Technological Change in Immiserizing Growth, Spatial Polarization, and Product Quality Upgrading” – CUP F53D23003240006. The usual disclaimers apply.

1 Introduction

In recent decades, the U.S. occupational structure changed substantially, with high-skilled occupations increasing their employment shares relative to middle- and low-skilled roles. These changes have been largely attributed to technological change, and more specifically to automation patterns (Acemoglu and Autor, 2011, Autor and Dorn, 2013, Acemoglu and Restrepo, 2020). However, while a large part of the literature investigates these long-run patterns, less attention has been placed on the short-run dynamics of this restructuring and the barriers that might hinder it.

In this paper, we focus on a specific barrier to accessing an occupation, namely the broad *training requirements* needed to perform it. We document that, while the demand for high-skilled labor has surged since the 2000s, training requirements in these occupations also increased, rising on average from 10.05 to 11.33 years over the sample period. In contrast, training requirements for low- and middle-skilled occupations remained relatively constant. Concurrently, the pace of automation had also been accelerating. Robots can already perform a wide range of tasks in various occupations with limited human intervention, including welding, painting and packaging.¹ AI promises to substitute a large fraction of tasks in white collar occupations in the next few years.² Taken together, these facts suggest that while labor-replacing technologies are likely to be rapidly and extensively adopted, barriers to occupational mobility increase. Motivated by these facts, we investigate how the interaction between fast automation and rising training barriers can lead to the emergence of *technological unemployment*, for instance by raising the aggregate unemployment rate or increasing the length of unemployment spells of displaced workers.

Investigating the potential emergence of technological unemployment matters beyond its effects on labor market aggregates. A growing body of evidence documents that prolonged unemployment has severe effects on individual health and well-being. Sullivan and Von Wachter (2009) show that job displacement leads to a 50–100% increase in mortality rates in the year following the event, with effects persisting even twenty years later. More broadly, the literature on the so-called *deaths of despair* has shown that the decline in economic opportunities for less-educated workers in the U.S. is associated with rising mortality

¹Morris (2023) in a BBC news article reports that the Ocado Technology warehouse in Luton (UK) has 44 robotic arms, which currently account for 15% of the products handled in the facility, about 400,000 items per week. The rest are handled by workers. In two or three years, Ocado expected robots will handle 70% of the products.

²While conceptually the introduction of a robot or AI is no different from the introduction of textile machinery in the 19th century, their speed of improvement and adoption is substantially faster, as they can be easily re-programmed to perform new sets of tasks. See Cossu (2025) for the effects that AI and Robots have on the skills required to perform the various occupations.

from suicide, drug overdose, and alcohol-related causes (Case and Deaton, 2015, 2017). More recently, and closely related to our focus, O’Brien, Bair, and Venkataramani (2022) provide causal evidence that the adoption of industrial robots in U.S. local labor markets has increased working-age mortality, operating both through direct material effects on displaced workers - such as reduced employment and wages - and through a diminished sense of future opportunity in affected communities, which fosters risky health behaviors. In this context, understanding how the interaction between automation and training barriers shapes unemployment spells, unemployment rates and the future prospect of workers is of first-order importance, not only for labor market outcomes, but also for population health.

To study the emergence of technological unemployment in an environment in which (i) the arrival of technology can displace workers in certain occupations and (ii) occupational mobility is hindered by training barriers, we build and calibrate a dynamic multi-occupation growth model. We model occupations as productive units which comprise a set of tasks. In a similar spirit to the growing literature on task-based technological change, we assume that current technology allows broadly defined capital (i.e., robots) to perform a subset of existing tasks and, in those tasks, labor and capital are perfect substitutes. As robots are cheaper than labor in equilibrium, they perform all the tasks for which a technology exists, while labor is left performing the remaining tasks. Technological change in the model occurs as an exogenous increase in the tasks that can be performed by robots, and such an increase can be heterogeneous across occupations. When technological change occurs within an occupation, *ceteris paribus*, the marginal product of labor declines and workers tend to flow out of that occupation.³

On the workers side, we assume a household is composed of a unit measure of agents. Each agent has instantaneous utility in consumption and labor time that they provide to the household, where all agents’ utilities are equally weighted. In each period, individual agents work in one of the J market occupations or they are unemployed, while the household owns the stock of robots in the economy and makes investment decisions. There is a unique good produced in the economy that can be used for consumption or investment (in robots) purposes.

We assume that workers laid off from an occupation must be trained before moving to a new occupation. Thus, given the sector k in which an agent works at time t , which is determined at $t - 1$, the amount of labor time that the agent provides is optimally split into two components: one which is sold by the household in the market in exchange for a wage and another which is used by the household to train new workers in occupation k for the

³We use the terms capital, machine and robots interchangeably in the paper. However, when we measure technological change, we specifically refer to robots’ capabilities in substituting labor in specific tasks.

next period $t + 1$.

To model the training process, and the fact that this can take time, we consider that the labor force in one occupation at time t is given by workers already in that occupation at $t - 1$, less some exogenous separations, plus newly trained workers. The latter are drawn from the unemployment pool through an occupation specific *training function* that uses, as inputs, the unemployed and the amount of time that workers in the destination occupation k use to train new workers. Thus, the household chooses in each period t : (1) consumption, which is equal for all agents in equilibrium; (2) the work effort of each type of agent in occupation k ; (3) how much of that effort goes into training instead of production; (4) the stock of robots at $t + 1$; and (5) the size of occupations at $t + 1$.

The chain of events is as follows. There is a one time, zero probability exogenous (unanticipated or MIT) shock that increases the number of tasks that robots can perform in a specific occupation. This induces the household to reallocate part of the labor force from that occupation. However, reallocating labor is costly, as it requires that part of employed workers' labor time is devoted to training new workers for that occupation.

We use variation across U.S. Metropolitan Statistical Areas (MSA) and time to estimate the parameter governing the training barriers in our training function. The estimation suggests that no training is required to transition to a low-skilled occupation, while the required training for high-skilled occupations is substantially larger than for middle-skilled occupations. We then calibrate the rest of the parameters such that the model replicates the occupational structure of the U.S. in 2006.

To identify the technological shock we use two methods. Using a natural language processing (NLP) algorithm, we classify which intermediate work activities (IWA) in O*NET can be potentially performed by robots within each occupation by using data from the International Federation of Robotics (IFR).⁴ Alternatively, we use the mapping of robots to tasks from [Teubert, Rendall, and Dowe \(2024\)](#) who rely on Large Language Models (LLMs) to create a match based on IFR applications. These maps allow us to determine the number of IWAs that can be performed by robots.

In both cases we aggregate IWAs to a three occupation classification level (high-, middle- and low-skilled) and we compute the ratio of these over the total number of IWAs by occupation. Given the small relative stock of robots with respect to other types of capital in the economy, we interpret this mapping as providing the set of IWAs for which a technology *already exists* to perform certain tasks, but that has not yet been adopted on a large scale.

⁴IWA are broad occupational tasks defined using a common standard across occupations. There are a total of 332 unique IWAs in O*NET. Unlike IWAs, tasks in O*NET are mostly unique to each occupation with virtually no overlap across occupations.

This gives us the size of the shock and the *future* state of the economy.⁵

We report three main sets of results. First, the shocks to the number of tasks performed by robots increase steady state unemployment, but reduce unemployment in the initial part of the transition. The main intuition for this result is that for shrinking occupations, the household has no incentive to speed up the transition. This is because workers attached to these occupations are valuable as they are productive assets. On the other hand, to speed-up the transition, the household increases training in the expanding occupations, which leads to a fall in unemployment as the unemployed are trained and move into the growing occupation. However, training drains resources from productive activities. As a result, flows out of shrinking occupations and into expanding ones are optimally balanced by the household and unemployment then rises only slowly towards the new steady state.

Second, we compare the tasks shock to a more standard shock to robots' productivity (capital embodied technological change), keeping the final steady state output fixed across the two experiments. We find that, while the tasks shock decreases unemployment in the first part of the transition, the productivity shock does not. With the productivity shock, the stock of robots increases less than with the task shock and labor reallocation across occupations is smaller. Real wages, instead, increase more with the productivity shock. Consequently, time spent in training increases less in the high-skilled occupation, which is the one receiving unemployed workers along the transition. As a consequence, there is an increase in the unemployment rate along the transition and in the new steady state. This result highlights that, when training time is an input in the reallocation of workers, labor reallocation can be hindered as the technological shock can raise real wages.

Third, we experiment with what happens when the task shock occurs simultaneously to a shock in the training requirements of high-skilled occupations of the same size as the one we estimate in the data between 2006 and 2019. Greater training requirements make unemployment increase along the transition, rather than decrease as in the benchmark experiment. In addition, steady state unemployment is larger than in the benchmark case. This is because increasing training requirements in high-skilled occupations create a larger barrier to entry into this type of occupation, reducing its size along the transition and in the steady state. Thus, in this case, additional unemployment is not created because workers flow from low- and middle-skilled occupations to high-skilled ones, but because the increase in training time

⁵One could argue that, given that these technologies already exist, the workers could forecast their adoption and react to this expectation. In the context of our model, we assume neither the household nor workers have such knowledge. Alternatively, one can assume that these agents cannot forecast when the new technology will be adopted and so they still take no action. Our experiment can also be viewed as the arrival of a technological revolution as in Caselli (1999). See the related literature section below for a discussion on this point.

to high-skilled occupations is not enough to offset the flow out of low- and middle-skilled occupations. This result highlights the importance of training for the occupational structure. However, if the shock to tasks is very large, the dynamics of unemployment are similar to the benchmark case without a shock to training. This is because the gains of moving agents to the high-skilled occupations are large despite the increase in training barriers, and it is optimal to have a substantial increase in training time in high-skilled occupations.

Related Literature

This paper is linked to three main strands of the literature related to (i) task based technological change and the race between technology and human capital, (ii) occupational switching and the costs involved in the process and iii) the health consequences of labor markets restructurings due to technological change.

We follow the recent literature that exploits a task based approach to technological change. This approach has gained popularity in the last two decades due to its ability to shed light on the evolution of the occupational structure (Autor, Levy, and Murnane, 2003, Autor and Dorn, 2013), something that the canonical model of skill-biased technological change in its basic version cannot do (Acemoglu and Autor, 2011).⁶ Recently, this approach has been used in macroeconomic settings to model technological change due to the arrival of robots. Acemoglu and Restrepo (2020) estimate the equilibrium impact of introducing robots in local labor market, finding negative effects in terms of employment and wages. Acemoglu and Restrepo (2022b) focus on wage inequality, finding that task displacement, due to the introduction of new technologies that can perform those tasks, can explain much of the changes in education wage differentials between 1980 and 2016. Acemoglu and Restrepo (2022a) find that automation is associated with aging of the working population, because the latter creates a shortage of middle-aged workers who perform manual tasks. Moll, Rachel, and Restrepo (2022) discuss how automation can increase inequality. The benefits of technological change accrue not only to high-skilled workers, who see an increase in their productivity, but also to owners of capital in the form of higher capital incomes. Concurrently, automation can lead to stagnant wages, and so stagnant incomes at the bottom of the distribution of assets, thus widening inequality in both incomes and assets over the population. These contributions suggest that the task approach to technological change displays factual predictions on the allocation of labor across occupations on one side, and on the wage distribution on the other side. In this paper, we exploit this type of technological change to study how the reallocation of labor across occupations after a technology shock

⁶See Cerina, Moro, and Rendall (2021) for a model of skill-biased technological change that can account for the evolution of the occupational structure when a gender dimension and home production are introduced.

can generate technological unemployment. In terms of the quantitative exercise, we study the response of the economy to a zero-probability unexpected (MIT) exogenous shock to the number of tasks that robots can perform in an occupation, which can be interpreted as a technological revolution in the spirit of [Caselli \(1999\)](#). In his model, new technologies that arrive exogenously require workers to acquire more skills to operate them with respect to previous technologies. In our model the technological shock in one type of occupation (e.g., middle-skilled) can trigger a reallocation of workers to another type of occupation (e.g., high-skilled), for which they require more training.⁷ Directly related to our research question is [Beraja and Zorzi \(forthcoming\)](#). They show how the degree of automation in a decentralized equilibrium is suboptimal, which opens a role for policy to improve welfare. As in this paper, they also study the effect of automation in a setting in which there are reallocation frictions across occupations. Compared to their approach, we assume that automation arrives as an exogenous shock and focus on the role of barriers to labor mobility created by training barriers. In addition, we calibrate the differential arrival of automation in high- middle- and low-skilled occupations, to quantitatively assess the effect of automation on unemployment and the occupational structure. We thus see the two papers as complementary.

Our paper also relates to the literature on the race between technology and human capital. Since the seminal contribution of [Tinbergen \(1974\)](#), a large body of work has emphasized that the evolution of the wage structure depends on the relative pace of technological progress and the accumulation of human capital. [Goldin and Katz \(2008\)](#) provide extensive evidence that the U.S. skill premium declined during most of the twentieth century, when the expansion of education outpaced the demand for skilled labor, and then rose sharply after 1980, when the growth in educational attainment slowed while skill-biased technological change continued unabated. In a complementary approach, [Krusell, Ohanian, Ríos-Rull, and Violante \(2000\)](#) show that capital-skill complementarity can account for the dynamics of the skill premium through observable changes in capital quantities and relative labor supplies. Our empirical finding that training requirements in high-skilled occupations have increased over time can be interpreted through this lens: as technology becomes more complex, the human capital needed to work alongside it rises, increasing the barriers to entry into high-skilled occupations. In our model, this mechanism operates through the training technology, which governs how quickly the labor force can be reallocated in response to technological shocks. Thus, the interaction between rising training barriers and task-displacing technological change can be viewed as a new dimension of the race between technology and human capital.

⁷A similar strategy is also adopted in [Marimon and Zilibotti \(1999\)](#), who study the effects of an MIT skill-biased technological innovation on search and matching labor markets that differ in the size of the unemployment benefits.

The second strand of the literature we relate to is that of cross-occupational mobility and the costs associated to it. [Phelan and Trejos \(2000\)](#) document that, when workers are forced to switch sector due to demand or supply shocks, the unemployment rate increases and the transition to the new steady state can be long. A recent paper by [Carrillo-Tudela and Visschers \(2023\)](#) documents that gross occupational mobility at reemployment is high and increases with unemployment duration. Occupations with mobility rates above 40% represent more than 80% of employment-unemployment-employment spells. Occupational movers have, on average, longer spells than stayers, thus, contributing substantially to the increase in aggregate unemployment. We build on these findings by focusing on occupational switchers and the switching costs involved, measuring training requirements of each occupation. Our theory rests on the assumption that longer unemployment spells of switchers are due to different training requirements between the origin and the destination occupation. We provide evidence that the length of unemployment spells is indeed positively correlated with the training index of the destination occupation for workers who switch occupation. Given this observation, our assumption on mobility frictions is different from another modeling strategy in the literature that assumes that workers lose part of their human capital when switching occupations ([Dvorkin and Monge-Naranjo, 2019](#)). In our theory, the requirements needed to perform an occupation must be acquired before entering the occupation, such that the average length of unemployment spells depends on the training requirements of the destination occupation.

Third, our paper connects to the literature on the health consequences of labor market disruptions driven by technological change. [Cutler, Deaton, and Lleras-Muney \(2006\)](#) review the economic determinants of mortality, identifying employment and income as key channels through which economic conditions affect health. [Case and Deaton \(2015, 2017\)](#) document and subsequently interpret the "deaths of despair" phenomenon, linking rising mortality among less-educated Americans to the erosion of labor market opportunities. [Sullivan and Von Wachter \(2009\)](#) establish a causal link between job displacement and long-run mortality using administrative data. Most directly related to our analysis, [O'Brien, Bair, and Venkataramani \(2022\)](#) show that automation affects health through two distinct channels: a direct material effect on displaced workers, and an indirect effect operating through the erosion of future economic prospects in affected communities. These two channels broadly correspond to the two horizons of our model's predictions. The material cost of displacement is reminiscent of the transitional dynamics generated by an automation shock. The erosion of future prospects, in turn, is reminiscent of the steady-state unemployment rate that emerges when training requirements rise, as a structurally larger barrier to expanding occupations permanently reduces the expected path back into employment. Our paper thus

complements this empirical literature by providing a structural mechanism that rationalizes why automation shocks can produce the persistent labor market disruptions linked to deteriorating health outcomes.

The rest of the paper is organized as follows: Section 2 discusses the evolution of the training index in the U.S.; Section 3 describes the model; Section 4 details the calibration and estimation; Section 5 presents the results; and Section 6 concludes.

2 Training Index

In this section we report the evidence on the training index and its time evolution in the U.S. From the O*NET database we use the *Education, Training and Experience* file, which contains data associated with the level of (i) education, (ii) work experience, (iii) on the job training and (iv) on site or in plant training for each occupation. The levels of each category are already classified in terms of months for work experience, job training and on site training. For education, we have educational attainment, which we convert into months of education above 12 years (a typical high-school degree). For each level within each of the four categories, we have the fraction of workers associated to it. We then start by estimating the time associated with each category in each occupation by using percentage frequency data. Take education for instance. For each year t and for each O*NET occupation code classification j we have that a fraction $p_{i,j,t}$ of workers performing the occupation has $m_{i,j,t}$ months of education (above a high-school degree), where i is the education level. Thus, we compute the education training index of occupation j at time t as,

$$Education_{j,t} = \sum_{i=1}^n (m_{i,j,t} \cdot p_{i,j,t}). \quad (1)$$

We then sum the indices obtained for education, work experience, on the job training and on site training to obtain the training index for a specific occupation in a given year. We use information on hours usually worked in a week and weeks usually worked in a year for US Census occupations available from 2006 to 2019 to compute employment shares for each year. The aggregate training index is a weighted average in which, for each occupation, the training index is weighted by the employment share. In this way, the aggregate training index is calculated for each reference year from 2006 to 2019.⁸

Figure 1 shows the results. The average time required to acquire the skills needed to

⁸We start in 2006, the first year in O*NET with a large enough occupation sample (i.e., larger than 300 occupations) after merging with Occ1990 Census classification and a consistent *Education, Training and Experience* file. We stop in 2019 to avoid compositional effects due to the Covid pandemic.

perform the representative occupation in the U.S. economy is increasing over time, from 6.29 years in 2006 to 7.04 years in 2019. As such, this increase implies that the representative occupation in the economy requires, on average, approximately three-quarters of a year more time to access over the period considered. We interpret this rise in light of the race between technology and education (Tinbergen (1974), Goldin and Katz (2008), Krusell, Ohanian, Ríos-Rull, and Violante (2000)): as technology becomes more complex, the human capital needed to work alongside it rises.⁹



Figure 1: The Aggregate Training Index

To investigate the sources of the rise of the aggregate training index, we calculate the index for a standard three occupation classification: high-skill, middle-skill, and low-skill.¹⁰ Figure 2 shows the results. There is a clear increase in the training index of high-skilled occupations,

⁹Modestino, Shoag, and Ballance (2020) look at job-posting during a recession, finding that education and experience requirements rose substantially during the Great Recession. This phenomenon, however, appear to vanish after the recession. Modestino, Shoag, and Ballance (2020) do not provide evidence that education and experience of workers *actually* hired during the recession was indeed higher than before. Thus, their findings are more in line with a strategic behavior of firms considering that “the recession is a wonderful opportunity to acquire top talent when workers are more plentiful”, rather than an actual increase in requirements due to technological reasons as we discuss in the text. O*NET data, instead, document that workers *currently employed* in an occupation display, on average, those levels of education, work experience, on the job training and on site or in plant training for each occupation.

¹⁰We use the definition in Cerina, Dienesch, Moro, and Rendall (2023) which builds on the occupational classification in Autor and Dorn (2013), in turn based on harmonized U.S. Census occupation codes (occ1990). High Skilled occupations are: all managerial and professional occupations (codes 004–199); middle skilled occupations are technicians, sales, admin services, mechanics and transport, precision worker, and machine operators (codes 203-889 except 405-471); and low skilled occupations are the remaining occupations (codes 405-471). We drop agriculture and mining.

rising from 10.05 to 11.33 years over the sample period. In contrast, the corresponding indices for middle- and low-skilled occupations are notably stable over time, increasing only marginally from 5.07 to 5.16 and from 2.65 to 2.67, respectively.¹¹ Since the index is expressed in years, this implies that accessing high-skilled occupations requires, on average, more than one additional year of training, while requirements for middle- and low-skilled jobs have remained broadly unchanged. Thus, the aggregate training index is increasing *both* because high-skill/high-training occupations are gaining employment shares in the economy, *and* because the training time required to perform these occupations is increasing over time.¹²

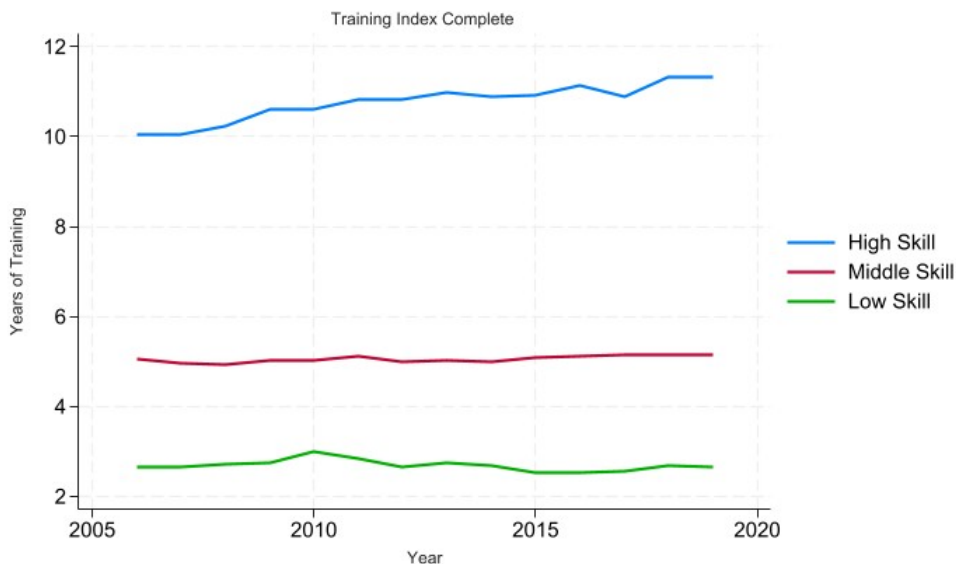


Figure 2: Training Index by occupation

2.1 Training time and Unemployment Spells

In general, switching occupational category at re-employment is associated with a larger unemployment spell compared to finding a new occupation within the same category. Thus, the duration of unemployment tends to be longer for switchers, as documented in [Carrillo-Tudela and Visschers \(2023\)](#). Here we aim to explore if the length of unemployment is positively related to the training time needed to perform the destination occupation. The idea is that, the higher the training time at the destination occupation, the higher the probability that the worker needs to acquire additional training before being able to perform the occupation.

¹¹We also computed the training index by broad sectors of economic activity: manufacturing and services. The patterns across high-, middle- and low-skilled occupations are similar to those observed in Figure 2.

¹²In Appendix A we discuss the rationale for using all types of training to construct our index.

To investigate this issue, we use data from CPS IPUMS to measure the unemployment spell of agents switching job from 1976 to 2019.¹³ We regress the average length of unemployment on the training index of the destination occupation and a binary variable that takes the value one if the destination occupation differs from the origin occupation. We also include an interaction term between the training index and the switching indicator which is our main variable of interest. We run three specifications. First, we exploit data from 1976 to 2019 and use a static version of the training index, where its value is fixed at the 2019 level. Second, we consider only the years 2006-2019, for which we can compute a measure of the training index at the three digit occupation level that changes over time.¹⁴ Third, we add time fixed effects to the previous specification. We thus estimate:

$$Y_{i,j} = \beta_0 + \beta_1 X_{1,j} + \beta_2 X_{2,j} + \beta_3 (X_{1,j} \times X_{2,j}) + \lambda_j + \xi_t + \varepsilon_i,$$

where $Y_{i,j}$ is the mean of weeks unemployed last year of agents laid off from occupation i and moving to occupation j ; $X_{1,j}$ is the training expressed in years associated with the j -th occupation of destination; $X_{2,j}$ is a dummy variable 'move' which assumes value 1 whenever the occupation of origin differs from the occupation of destination, $i \neq j$; λ_j are occupation fixed effects, and ξ_t are year fixed effects. Results are reported in Table 1.

As in Carrillo-Tudela and Visschers (2023), we find that workers switching occupation have longer unemployment spells, with the coefficient on Move being positive and statistically significant in the second and third specifications. The two variables Training and Training_Move allow us to distinguish the effect of training between workers who find a new job within the same occupational category and those that switch category at re-employment. The coefficient on the interaction term between training and occupational mobility is positive and statistically significant across all specifications, while the coefficient on training alone is not statistically different from zero. This implies that, only for workers who switch occupations, higher training requirements in the destination occupation are associated with longer unemployment spells. As expected, training requirements do not lengthen the unemployment spell of non-switchers, as those workers already have the training required to perform the destination occupation, which is the same as the origin occupation. Notably, this pattern holds for the four specifications. While these results do not claim causation, they show that longer unemployment spells are associated with higher training requirements specifically for workers who switch occupations, suggesting that an increase in training may

¹³We use the variable *weeks unemployed last year* (*wksunem1*), which provides information on the duration of unemployment in the previous year for both currently employed and unemployed respondents, as long as the respondents are in the universe for this variable.

¹⁴See Appendix A.

Table 1: Relationship between weeks unemployed last year, training and move

<i>Dependent Variable: Weeks unemployed last year</i>				
Variable	(1)	(2)	(3)	(4)
Move	-0.246 (0.243)	1.912*** (0.430)	1.440*** (0.418)	1.366*** (0.445)
Training	-0.038 (0.026)	-0.033 (0.046)	-0.038 (0.044)	-0.140 (0.173)
Training_Move	0.181*** (0.039)	0.180*** (0.067)	0.141** (0.065)	0.116* (0.068)
Intercept	15.70*** (0.175)	16.24*** (0.318)	12.75*** (0.611)	15.94*** (4.796)
Observations	30,958	10,455	10,455	10,455
Year FE	No	No	Yes	Yes
Occupation FE	No	No	No	Yes

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Standard errors are in parentheses below coefficients. Column (1) is based on observations from 1976 to 2019. Columns (2), (3) and (4) are based on observations from 2006 to 2019. The occupation classification is at the occ1990 level.

affect the time needed for displaced workers to transition into new occupations.¹⁵

2.2 Occupational flows, Robots and Training

Having established in the previous subsection that higher training requirements are associated with longer unemployment spells for workers who switch occupation, we now investigate the relationship between robots exposure and training requirements. In particular, the theory we present in Section 3 below builds on the insight that workers displaced from occupations exposed to automation find it more difficult to transition into new occupations when these require higher levels of training. To test this working hypothesis, we study whether i) *automation* exposure in the *origin* occupation and ii) *training* requirements in the *destination* occupation interact to determine observed labor flows across occupations.

For this purpose, we use CPS IPUMS data from 2006 to 2019 to construct measures of occupational mobility across occupations. We define occupational flows from origin occupation k to destination occupation j at time t as the share of total hours worked in occupation j that is performed by workers who were previously employed in occupation k and switched

¹⁵We have also estimated the same specification using an alternative version of the training index that excludes both work experience and on-the-job training components. The results are robust since the interaction term remains positive and statistically significant across all specifications.

occupation between $t - 1$ and t .

Formally:

$$Flow_{k \rightarrow j, t} = \frac{\sum_i w_{i,t} \cdot hours_{i,t} \cdot \mathbf{1}_{occ_{i,t} = j} \cdot \mathbf{1}_{occ_{i,t-1} = k}}{\sum_i w_{i,t} \cdot hours_{i,t} \cdot \mathbf{1}_{occ_{i,t} = j}}, \forall j \neq k \quad (2)$$

where $w_{i,t}$ denotes the CPS sampling weight and $hours_{i,t}$ is the number of hours worked by individual i at time t . This measure captures the turnover of occupation j generated by workers transitioning from occupation k .

We relate these flows to training requirements in destination occupations and to robots exposure in origin occupations. Training requirements are measured using the training index described in Section 2, constructed for the period 2006–2019. Robot exposure is measured as the share of IWAs in each occupation that can potentially be performed by robots, for the period 2006–2019.¹⁶ Our specification is:

$$Y_{k \rightarrow j, t} = \beta_0 + \beta_1 X_{1,j} + \beta_2 X_{2,k} + \beta_3 (X_{1,j} \times X_{2,k}) + \lambda_j + \xi_t + \varepsilon_i,$$

where $Y_{k \rightarrow j, t}$ is the occupational flows variable defined in (2); $X_{1,j}$ is the training expressed in years of the occupation of destination j ; $X_{2,k}$ captures robots exposure in the origin occupation k ; λ_j are occupation fixed effects, and ξ_t are time fixed effects. This specification allows us to test whether the relationship between technological displacement and reallocation depends on the training requirements of destination occupations.

Table 2 reports the results. Across all specifications, we find that flows into occupations requiring higher levels of training are lower, as indicated by the negative and statistically significant coefficient on the training variable. Robot exposure in the origin occupation k is positively associated with occupational flows towards occupation j , although the coefficient is significant only in column (3). The key result, however, is given by the interaction between training requirements and robot exposure, which is negative and statistically significant across all specifications. This implies that flows from occupations more exposed to robots are smaller when the training in the destination occupation is larger. As in the previous subsection, while not claiming causality, these findings support the idea that training requirements can act as a barrier to occupational mobility due to technological exposure. In the next section, we embed these insights into a dynamic multi-occupation growth model.

¹⁶See Appendix B for details.

Table 2: Occupational flows, Robots and Training

<i>Dependent Variable: Occupational Flow</i>			
Variable	(1)	(2)	(3)
Training	-0.0001*** (0.0000)	-0.0001*** (0.0000)	-0.0004** (0.0002)
Robots	0.0071 (0.0052)	0.0067 (0.0052)	0.0130** (0.0052)
Train_Rob	-0.0017*** (0.0006)	-0.0016*** (0.0006)	-0.0019*** (0.0006)
Intercept	0.0105*** (0.0003)	0.0078*** (0.0005)	0.0100** (0.0040)
Observations	56,227	56,227	56,227
Year FE	No	Yes	Yes
Occupation FE	No	No	Yes

Note: * p<0.1; ** p<0.05; *** p<0.01. Standard errors are in parentheses below coefficients. Observations are from 2006-2019. The occupation classification is at the occ1990 level.

3 Model

3.1 Household

There is a household composed of a unit measure of agents, $\iota \in [0, 1]$. Each agent has instantaneous utility,

$$\hat{u}_{\iota,t} = \frac{c_{\iota,t}^{1-\psi} - 1}{1-\psi} - B_k l_{\iota,t}^\mu,$$

where c_ι is consumption and l_ι is the amount of labor that the agents provides to the household. B_k is a parameter governing the disutility of working, which is occupation specific.¹⁷

Each agent has total lifetime utility given by:

$$u_\iota = \sum_{t=0}^{\infty} \beta^t \hat{u}_{\iota,t},$$

and utility of the household is:

$$W = \int_0^1 u_\iota d\iota.$$

At any time t , agents work in one of J market occupations, labeled $k = 1, 2, \dots, J$, or they

¹⁷Thus, if an agent switches occupation, her B_k changes.

are unemployed. Defining $N_{k,t}$ as the fraction of the population working in occupation k at time t , unemployment at t is defined by:

$$u_t = 1 - \left[\sum_{k=1}^J N_{k,t} \right]. \quad (3)$$

In each period, the household owns the stock of robots r_t in the economy and decides how much investment to make in robots,

$$r_{t+1} = I_t + r_t(1 - \delta),$$

where I_t is investment in robots and δ is the depreciation rate of the stock of robots. There is a unique good produced in the economy that can be used for consumption or investment purposes.

Given the sector k in which an agent works at time t , which is decided at $t - 1$, the amount of labor time the agent provides to the household, $l_{k,t}$, is optimally split by the latter into two components: $l_{k,t} - l_{k,t}^s$, which is sold in the market in exchange for a wage $w_{k,t}$ per unit of time, and $l_{k,t}^s$, which is used by the household to train new workers that will work in occupation k from $t + 1$ onward.

To model the fraction of agents available to work in each occupation at $t + 1$ we follow [Phelan and Trejos \(2000\)](#), and define the following law of motion:

$$N_{k,t+1} = (1 - \kappa) \left[(1 - \gamma_k) N_{k,t} + \pi_k(u_t, l_{k,t}^s) N_{k,t} \right], \quad (4)$$

where γ_k is an exogenous separation rate at which the labor force in occupation k becomes unemployed, and κ is a probability of dying (i.e., exiting the labor force).¹⁸ The function π_k is defined as:

$$\pi_k(u_t, l_{k,t}^s) = \Psi u_t^{\phi_k} (l_{k,t}^s)^{1-\phi_k}, \quad (5)$$

and represents a *training technology* available to the household, which denotes the amount of the new entrants into occupation k achieved per existing worker in occupation k .¹⁹ The inputs to this function are the current pool of unemployed workers, u_t , and the amount of time workers currently in occupation k spend training new workers, $l_{k,t}^s$. It follows that the

¹⁸Each period, a fraction κ of the population exits the labor force and is replaced by new agents who enter the economy as unemployed.

¹⁹[Phelan and Trejos \(2000\)](#) adopt a similar setting to model labor flows across different sectors of the economy. However, they consider the term $l_{k,t}^s$ as a search effort in recruiting new workers by firms, thus interpreting (5) as a matching function typical of the search and matching literature. For us, $l_{k,t}^s$ is time spent to train new workers, and so we interpret (5) as a training technology and estimate it accordingly in the quantitative part.

total amount of training time in occupation k at time t is given by $N_{k,t}l_{k,t}^s$. The parameter ϕ_k denotes how difficult/costly it is to train new workers for occupation k . A larger ϕ_k implies that unemployment is more important while a smaller ϕ_k implies that the training time of workers becomes more relevant.²⁰

Substituting the definition of unemployment and the training function into (4), we obtain that the law of motion of workers in occupation k is given by:

$$N_{k,t+1} = (1 - \kappa)N_{k,t} \left[(1 - \gamma_k) + \Psi(l_{k,t}^s)^{1-\phi_k} \left(1 - \left[\sum_{k=1}^J N_{k,t} \right]^{\phi_k} \right) \right]. \quad (6)$$

The utility function of the household weights all agents equally, so the marginal utility of consumption for each agent is the same at the maximized utility. We prove in Appendix C that this implies that all agents display a common level of consumption at any date t , which we denote by c_t . It follows that we can write the problem of the household as:

$$\max_{c_t, l_{k,t}, l_{k,t}^s, r_{t+1}, N_{k,t+1}} W = \sum_{t=0}^{\infty} \beta^t \left\{ \frac{c_t^{1-\psi} - 1}{1-\psi} - \left[\sum_{k=1}^J B_k l_{k,t}^{\mu} N_{k,t} \right] \right\}, \quad (7)$$

subject to the budget constraint,

$$c_t + r_{t+1} - r_t(1 - \delta) \leq \sum_{k=1}^J w_{k,t} N_{k,t} (l_{k,t} - l_{k,t}^s) + p_{r,t} r_t, \quad (8)$$

and condition (6) for each k . In (8), $p_{r,t}$ is the rental rate of robots, while the price of consumption and investment is taken as the numeraire and normalized to one. The household chooses in each period t : (1) the common consumption level of agents, c_t ; (2) the work effort of each type of agent in sector k , $l_{k,t}$; (3) how much of that effort goes to the training function instead of production, $l_{k,t}^s$; (4) the stock of robots at $t+1$; and (5) the size of each occupation at $t+1$.

²⁰The training function allows us to tractably model the flow of workers across broad occupational categories in such a way that they depend on the training time used in the training process. A potential concern with this specification is that it considers all unemployed workers alike. Thus, a worker previously employed in a high-skilled occupation is equal to one previously employed in a middle-skilled occupation from the standpoint of the training function. While this might appear a strong assumption, note that the training function in occupation k transforms only a fraction of unemployed workers into new labor force for occupation k . A potential interpretation here is that training time is the requirement to retrain the actual newly hired worker and not all unemployed workers. In addition, our three occupation classifications imply substantial heterogeneity within each group. Thus, an agent separated from an occupation like “management” might move to “science occupations.” While in our categorization both occupations belong to the high-skilled group, it is unlikely that no re-training is required for the switch.

3.2 Firm

3.2.1 Occupations

We follow the recent literature on task based technological change in defining an occupation j as a set of tasks $i \in [0, I]$. These produce an occupational output denoted by L_j , which can be interpreted as a quantity index of all the tasks performed within the occupation. The amount of labor used to perform a task i is denoted l_{ij} . We assume that there potentially exists a technology that allows robots to perform task i in occupation j , and denote r_{ij} as the stock of robots used to perform such task. When the technology exists for robots to perform a certain task, then labor and robots are perfect substitutes in performing the task (Autor, Levy, and Murnane, 2003). Thus, output of occupation j is defined as:

$$L_j = \left[\int_0^I \theta_i [\alpha_j r_{ij} + l_{ij}]^{\rho_j} di \right]^{\frac{1}{\rho_j}},$$

where ρ_j governs the elasticity of substitution between two different tasks, θ_i is the importance of task i in the occupation and α_j is the productivity of robots in occupation j , which can be interpreted as a measure of robots' embodied technology.

We assume that there is a firm that maximizes the output of a certain occupation j , given the amount of labor time, \bar{l}_j , and the stock of robots \bar{r}_j , allocated to the occupation.²¹ The problem is:

$$\max_{l_{ij}, r_{ij}} L_j = \left[\int_0^I \theta_i [\alpha_j r_{ij} + l_{ij}]^{\rho_j} di \right]^{\frac{1}{\rho_j}},$$

subject to $\sum_{i=1}^I l_{ij} = \bar{l}_j$ and $\sum_{i=1}^I r_{ij} = \bar{r}_j$. We make the following three assumptions:

- **Assumption 1:** We assume that technology is such that robots can perform tasks $[0, \hat{i}]$ while they cannot perform tasks $(\hat{i}, I]$.
- **Assumption 2:** All tasks have the same importance in the occupation, $\theta_i = \theta = 1$ for each task i .
- **Assumption 3:** Labor and robot prices in equilibrium are such that $\alpha_j w_j > p_r$.

Assumption 1 defines the exogenously set of tasks that robots can perform. In the quantitative part, we study the reaction of the economy to a sudden and unexpected (i.e., MIT) shock that increases this fraction. Assumption 2 is a symmetry assumption made for tractability.

²¹This is the same firm we introduce in the next section, that produces output by means of a technology that aggregates occupations. The problem discussed here is a sub-problem of the profit maximization problem of that firm.

Assumption 3 requires that, when a technology exists for robots to perform a certain task, these are cheaper than labor, so that only robots are used in performing that task.²² Under these assumptions, the maximized output of occupation j is a function of the fraction of tasks that can be performed by robots \hat{i} and it is given by:

$$L_j(\hat{i}_j) = \left[\alpha_j^{\rho_j} \hat{i}_j^{1-\rho_j} \bar{r}_j^{\rho_j} + (I - \hat{i}_j)^{1-\rho_j} \bar{l}_j^{\rho_j} \right]^{\frac{1}{\rho_j}}. \quad (9)$$

3.2.2 Production

There is one representative firm in the economy producing output at each time t with the following production function:

$$Y = A\bar{L}. \quad (10)$$

Here \bar{L} is a labor input given by an aggregation of the J occupations within the firm,

$$\bar{L} = \left[\sum_{k=1}^J \eta_k L_k(\hat{i}_k)^\sigma \right]^{\frac{1}{\sigma}}, \quad (11)$$

where $\sum_{k=1}^J \eta_k = 1$ and, in turn, $L_j(\hat{i}_j)$ can be expressed according to (9), as

$$L_j(\hat{i}_j) = \left[\alpha_j^{\rho_j} \hat{i}_j^{1-\rho_j} R_j^\rho + (I - \hat{i}_j)^{1-\rho} L_j^\rho \right]^{\frac{1}{\rho}}. \quad (12)$$

The variables R_j and L_j are the amounts of robots and labor that the firm decides to employ in occupation j , and \hat{i}_j defines the number of tasks in that occupation that can be performed by robots.²³

Under perfect competition, the firm maximizes profits,

$$\max_{L_k, R_k} \Pi = Y - \left[\sum_{k=1}^J w_k L_k \right] - p_r \left[\sum_{k=1}^J R_k \right], \quad (13)$$

subject to (10), (11) and (12). Note that in (13), the price of output is normalized to one, as this good is the numeraire of the economy.

²²In the general equilibrium of the model, we verify that this condition is always met.

²³Note that we differentiate notation for R_j and L_j with respect to the previous section, because in there robots and labor at the occupation level were taken as given, while here they represent choice variables for the firm.

3.3 Equilibrium

An equilibrium for the economy under study is a set of prices $\left\{ \{w_{k,t}\}_{k=1}^J, p_{r,t} \right\}_{t=0}^{\infty}$, allocations for the household $\left\{ c_t, r_{t+1}, \{l_{k,t}, l_{k,t}^s, N_{k,t+1}\}_{k=1}^J \right\}_{t=0}^{\infty}$ and allocations for the firm $\left\{ \{L_{k,t}, r_{k,t}\}_{k=1}^J \right\}_{t=0}^{\infty}$ such that:

- a) Given prices, $c_t, r_{t+1}, \{l_{k,t}, l_{k,t}^s, N_{k,t+1}\}_{k=1}^J$ solve the household's problem (7);
- b) Given prices, $\{L_{k,t}, r_{k,t}\}_{k=1}^J$ solve the firm's problem (13) at each t ;
- c) All markets clear at each t :

$$Y_t = c_t + I_t,$$

$$L_{k,t} = (l_{k,t} - l_{k,t}^s)N_{k,t+1}, \forall k = 1, \dots, J$$

$$r_{k,t} = \sum_{k=1}^J R_{k,t}.$$

4 Calibration and Estimation

To parameterize the model, we proceed in three steps. We first estimate the barriers to entry in each occupation (i.e., the importance of training) by exploiting the law of motion of workers by occupation, Equation (6). We then use two separate methods to identify the size of the technology shock by using data from the Occupational Information Network (O*NET) and the International Federation of Robotics. Among the remaining parameters, a set of them is based on previous literature estimates, while the rest is calibrated to match a number of data targets.

4.1 Estimating barriers

In this section we describe the methodology to estimate the size of barriers, which in the model are represented by the value of ϕ_k in each occupation. From the law of motion of occupational size (6), we can write:

$$\frac{N_{k,t+1} - (1 - \kappa)(1 - \gamma_k)N_{k,t}}{N_{k,t}} = (1 - \kappa)\Psi(u_t)^{\phi_k} (l_{k,t}^s)^{1-\phi_k}. \quad (14)$$

Taking logs on both sides, we obtain a log-linear equation suitable for estimation, in which the coefficient on unemployment provides the value of ϕ_k . Given the evidence in Figure 2, it is plausible that ϕ_k also varies over time. Thus, we use variation across U.S. Metropolitan Statistical Areas (MSA) and time to estimate a common training process. As we do not observe the value of $l_{k,t}^s$ in the data, we construct a proxy of it by exploiting training

requirements of unemployed workers by MSA in a given year based on the training index from Section 2. The training variable is constructed as $train_{k,t}^m = \sum_i (train_{k,t} - train_{i,t}^m)$, where $train_{i,t}^m$ is the training individual i from the unemployment pool in MSA m currently has and $train_{k,t}$ is the training required to work in occupation k at time t .²⁴ Thus, our proxy captures how much training unemployed workers in MSA m would need to be able to perform occupation k at time t . The larger the training needed by unemployed workers, the larger should be the time spent by workers already in occupation k to train new workers, $l_{k,t}^s$.

We estimate a pooled regression for all occupations, as the model implies a common training productivity Ψ across occupations. Defining the left hand side of equation (14) as $\Delta N_{k,t+1}$, we estimate,

$$\log(\Delta N_{k,t+1}^m) = a_0 + \sum_j \left((a_{1j} + b_{1j} \times t) \log(u_t^m) + a_{2j} \log(train_{j,t}^m) + \sum_h \gamma_{hj} X_{ht}^m \right) \times occ_j + \delta_t + \epsilon_{k,t}^m, \quad (15)$$

where occ_j is a dummy taking the value 1 if $j = k$ and zero otherwise, δ_t is a common time fixed effect, the superscript m refers to the MSA, the weight on unemployment is, $a_{1k} = \phi_k$, b_{1k} captures a time trend in ϕ_k and X_h are controls characteristic of the unemployment pool in an MSA m , namely the share of unemployed originally coming from occupation k and the share of agents not in the labor force (NILF) of the MSA.²⁵

We set the high-skilled occupation as the reference occupation. Table 3 shows a number of specifications with different subsets of X_h and with or without time trends in the ϕ_k 's. The bottom four rows compute the linear combinations and corresponding p-values for the parameter of interest, ϕ_k , using the Delta-method. Our preferred specifications are models (4) and (5). Consistent with the descriptive evidence in Section 2, these models estimate that the largest training barriers are in high-skilled occupations. The weight on the unemployed pool for the high-skilled lies between 0.376 and 0.592. For the simulations in the next section, we use the average between the two $\phi_h = 0.484$. We do the same for the middle-

²⁴To estimate barriers we consider our unemployment pool to include both temporary unemployed workers looking for a job (so the typical definition of unemployed workers) and those individuals of working age classified as "out of the labor force". We do this because we want to capture the whole pool of agents not in an occupation, but that can currently receive training to obtain a job in the future.

²⁵Note that while we include individuals not in the labor force in the unemployment pool, as described in note 24, we do exclude all disabled and retired workers. Also, note that although our proxy $train_{k,t}^m$ is expected to be correlated with training, it is not the training itself. For this reason, the coefficient a_{2j} should not be considered informative of the value of $(1 - \phi_k)$. For the same reason, we do not include a time trend similar to $b_{1j} \times t$ on a_{2j} , but instead allow the time fixed effects to absorb any broad changes over time not related to our coefficient of interest, ϕ_k .

Table 3: Estimation of Training Costs in Training Function

VARIABLES	(1)	(2)	(3)	(4)	(5)
Unemployment	0.139 (0.114)	0.300*** (0.116)	0.533*** (0.117)	0.376*** (0.118)	0.592*** (0.118)
Unemp. high-skilled x time				-0.022*** (0.005)	-0.018*** (0.005)
Unemp. x middle-skilled	0.530*** (0.130)	0.361*** (0.134)	0.215 (0.142)	0.315** (0.136)	0.181 (0.143)
Unemp. middle-skilled x time				-0.006 (0.005)	-0.005 (0.005)
Unemp. x low-skilled	0.661*** (0.114)	0.500*** (0.116)	0.267** (0.117)	0.424*** (0.118)	0.208* (0.118)
Train	-0.922*** (0.093)	-0.765*** (0.094)	-0.957*** (0.137)	-0.800*** (0.094)	-0.976*** (0.137)
Train x middle-skilled	0.123 (0.139)	-0.108 (0.148)	0.447** (0.193)	-0.074 (0.148)	0.451** (0.193)
Train x low-skilled	0.503*** (0.143)	0.271* (0.144)	0.677*** (0.184)	0.208 (0.147)	0.614*** (0.189)
a_0	-0.835*** (0.062)	-1.231*** (0.080)	0.649*** (0.198)	-1.123*** (0.087)	0.714*** (0.198)
Observations	5,249	5,249	5,243	5,249	5,243
Year FE	YES	YES	YES	YES	YES
Control unemployed pool	No	Yes x Occ	Yes x Occ	Yes x Occ	Yes x Occ
Control NILF	No	No	Yes x Occ	No	Yes x Occ
ϕ_h	.139	.3***	.533***	.376***	.592***
ϕ_m	.669***	.66***	.749***	.691***	.772***
ϕ_l	.8	.8	.8	.8	.8
$\phi_{h,2018}$.116.	.374***

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Standard errors are in parentheses below coefficients. The regression is based on observations from 2006 to 2018. The high-skilled occupation is the reference group.

skilled obtaining $\phi_m = 0.732$. The time trend on the training of the middle-skilled is close to zero and not statistically significant while the average counterfactual high-skilled parameter is $\phi_{h,new} = 0.245$, implying a $\Delta_{\phi_h} = 0.245 - 0.484 = -0.239$ over the period considered, which is used in quantitative Section 5.3 only.²⁶

We note here that in estimating the models in Table 3 we impose the restriction that $\phi_l = 0.8$. This is because running the unconstrained models gives a statistically insignificant ϕ_l in all specifications.²⁷ As we are mainly interested in estimating barriers into high-skilled occupations, we impose an exogenous and high value for ϕ_l , which implies that training barriers, consistent with the evidence in section 2, are low in low-skilled occupations.²⁸ We remark that we cannot use the natural choice/limit of $\phi_l = 1$, as our model requires a $\phi_k < 1$ to generate unemployment differences across steady states.²⁹

4.2 Shocks to \hat{i}_j

We use two separate methods to identify a lower and an upper bound to the size of the technology shock.³⁰ Both methods are based on the Occupational Information Network (O*NET) “Task Statement” module and the “Work Activities, Intermediate Work Activities (IWA) and Detailed Work Activities” module. To identify the shock(s) to task in the data, we proceed as follows. We first map tasks to Intermediate Work Activities (IWA) in O*NET. Then, by using the International Federation of Robotics (IFR) data, we either (1) create a match between the tasks in each occupation and the applications of robots through text mapping (*text shock*) or (2) make use of the mapping from Teubert, Rendall, and Dowe (2024), who use Large Language Models (LLM) to determine which tasks an IFR robot application can accomplish (*GPT shock*).³¹ Both these maps allow us to compute the number

²⁶The value $\phi_{h,new} = 0.245$ is obtained as the mean of 0.116 and 0.374 from the last row of Table 3.

²⁷Descriptive evidence, not reported here, about the relationship between changes in occupational employment and training requirements across MSAs appears to confirm this result. At the beginning of the sample period (2006), the relationship between changes in employment and training requirements is relatively flat across all skill groups. In contrast, by the end of the sample period (2018), a clearer and negative relationship emerges for high- and middle-skilled occupations, where changes in employment are more strongly associated with training requirements.

²⁸We also run robustness with $\phi_l \in [.9, .99]$. Results for ϕ_h and ϕ_m are very close to those in Table 3 and are available upon request. In general, the exogenous choice on ϕ_l does not have an impact on the estimated ϕ_h and ϕ_m .

²⁹To see this, consider the law of accumulation of occupation specific employment in (6). Imposing steady state and $\phi_k = 1$ leaves an equation which can be solved for the unemployment level. This depends only on the values of the parameters γ_k and κ , and so it is invariant to technological change.

³⁰In Appendix B, we also compute the shocks with other methods and discuss why these different methods provide different shock sizes.

³¹Details can be found in Appendix B.

of IWA that can be performed by robots in each occupation.³² Given the current small stock of robots with respect to other types of capital in the economy, we interpret these mappings as providing the set of IWA for which a technology *already exists* to perform certain tasks, but that has not yet been adopted on a large scale. This is particularly true as O*NET reports tasks performed by the worker, i.e., a task that appears in O*NET by the purpose of the survey is *currently* performed by a worker in that occupation, and not by a machine or robot. This gives us the size of the shock and the *future* state of the economy.

Both methods produce a similar pattern across occupations, with high-skilled occupations having the smallest proportion of tasks that can be replaced, and low-skilled occupations having the largest proportion. For example, LLMs suggest that robots are capable of doing up to 54 percent of human tasks in low-skilled occupations, while our most conservative text mapping algorithm suggests 8 percent. Given the empirical evidence on the development of robots, we take the LLMs estimates as upper bounds and the text mapping results as a lower bound of the shocks to \hat{i}_j .

Technically, we proceed as follows to introduce the shock in the model. We measure in the data the fraction x of labor tasks $I - \hat{i}_j$ that robots can potentially perform in occupation j with the two methodologies described above. It follows that $x(I - \hat{i}_j)$ gives us the shock $\Delta \hat{i}_k$ in occupation k .

4.3 Predetermined parameters

We start by normalizing to one total factor productivity $\{A\}$ and the number of tasks in each occupation $\{I\}$. Thus, \hat{i}_j is naturally interpreted as the share of tasks performed by robots in occupation j . We then set the parameters $\{\beta, \sigma, \delta, \psi, \mu\}$ to standard values in the literature. We set a subjective discount factor of $\beta = 0.96$, the inverse of the intertemporal elasticity of substitution to $\psi = 1.5$, and the inverse of wage elasticity of labor supply to $\mu = 3.0$. The elasticity of substitution between occupations is set to $\sigma = 0.254$ from estimates in [Caunedo, Jaume, and Keller \(2023\)](#). We also follow them in computing the depreciation rate $\delta = 0.113$ and in estimating the elasticity of substitution between robots (i.e., capital) and labor at the occupation level. Adopting their estimation to our 3-occupation definition delivers $\rho_l = 0.244$, $\rho_m = 0.247$, and $\rho_h = -0.153$. These numbers imply that capital and labor are substitutes in low- and middle-skilled occupations, while they are complements in high-skilled occupations. Finally, the separation rates $\{\gamma_j\}$ and the fraction of new workers $\{\kappa\}$ are computed using 2006-2018 CPS data and refer to the average separation rate by occupation during this period and the fraction of young new workers without prior history

³²We assume that an IWA can be performed by robots if there is at least one task pertaining to that IWA that is mapped to an IFR application.

in the labor market.

4.4 Calibrated Parameters

We first assume that $\eta_h = 1 - \eta_l - \eta_m$, $B_m = 1$, $\alpha_m = 1$. We are then left with ten parameters, $\{\eta_l, \eta_m, \Psi, B_l, B_h, \alpha_l, \alpha_h, \hat{i}_{l,0}, \hat{i}_{m,0}, \hat{i}_{h,0}\}$, where $\hat{i}_{k,0}$ refers to the current level of tasks that can be performed by capital/robots in occupation k . We use ten targets from 2006, employment shares of each occupation (two targets), worker shares of each occupation (two targets), unemployment rate (one target), wage premium (two targets), and labor expenditure shares by occupation (three targets) as follows:

1. The weight of each occupation in output $\{\eta_j\}$ determines the importance that each occupation has in producing output. We match CPS hours employment shares;
2. The productivity of the training functions $\{\Psi\}$ affects the rate at which agents can move from unemployment to employment. We set this to match an unemployment rate of 0.167 in the steady state;³³
3. The weight of labor in utility $\{B_k\}$ measures the disutility that the household has working in the market. We calibrate them by matching the CPS share of workers in each occupation;
4. The productivity of capital/robots $\{\alpha_k\}$. We calibrate them by matching the wage premium between high- and middle-skilled and between high- and low-skilled workers.
5. The initial tasks performed by robots $\hat{i}_{k,0}$. We match the expenditure share of labor by occupation. We compute labor expenditures using CPS data and [Caunedo, Jaume, and Keller \(2023\)](#) tool expenditure, imposing an aggregate labor share of 0.59.

All parameter values are found in Table 4, while data and model targets are in Table 5.

³³Note that our definition of unemployment includes agents not in the labor force.

Table 4: *Model Parameters*

Estimated	Type	Value
$\{\phi_l, \phi_m, \phi_h\}$	Weight on unemployed pool (low, middle, high)	$\{.8, .732, .484\}$
$\Delta\phi_h$	Shock to training importance (high-skilled)	-0.239
$\{\gamma_l, \gamma_m, \gamma_h\}$	Separation rates (low, middle, high)	$\{.055, .049, .018\}$
$\{\rho_l, \rho_m, \rho_h\}$	Elasticity of substitution between labor and robots	$\{.244, .247, -.153\}$
δ	Depreciation rate of robots	.113
κ	Probability of exiting the labor force	.008
$\left\{ \Delta\hat{i}_l, \Delta\hat{i}_m, \Delta\hat{i}_h \right\}_{lb}$	Lower-bound task shock (low, middle, high)	$\{.060, .042, .016\}$
$\left\{ \Delta\hat{i}_l, \Delta\hat{i}_m, \Delta\hat{i}_h \right\}_{ub}$	Upper-bound task shock (low, middle, high)	$\{.312, .240, .069\}$
Calibrated	Type	Value
$\{\eta_l, \eta_m, \eta_h\}$	Weight of occupations in output	$\{.164, .551, .285\}$
$\{\alpha_l, \alpha_m, \alpha_h\}$	Robot productivity by occupation	$\{2.545, 1, 2.982\}$
$\{\hat{i}_l, \hat{i}_m, \hat{i}_h\}$	Robot task share by occupation	$\{.308, .342, .423\}$
Ψ	Training function coefficient	.517
$\{B_l, B_m, B_h\}$	Disutility of work by occupation	$\{.231, 1, 2.275\}$
Predeter.	Type	Value
σ	Substitutability between occupations	0.254
A	Aggregate productivity	1
ψ	Income elasticity of labor supply	1.5
β	Discount factor	.96
μ	Inverse of wage elasticity of labor supply	3.0

Note: The first set of parameters is estimated from 2006-2019 CPS data, the second set is calibrated to match targets in Table 5, and the last is predetermined.

Table 5: *Targets*

Type	Data	Model
Parameters on labor $\{\eta_l, \eta_m, B_l, B_h\}$:		
Employment shares in hours (low, middle)	$\{.195, .518\}$	$\{.228, .522\}$
Employment shares in workers (low, middle)	$\{.172, .523\}$	$\{.142, .523\}$
Parameters on robots $\{\alpha_l, \alpha_h, \hat{i}_l, \hat{i}_m, \hat{i}_h\}$:		
Wage premium (high to low, high to middle)	$\{2.164, 1.395\}$	$\{2.119, 1.215\}$
Labor shares by occupation (low, middle, high)	$\{.525, .526, .674\}$	$\{.526, .517, .673\}$
Parameter on training $\{\Psi\}$:		
Unemployment rate	.167	.164

Note: Targets are estimated from 2006 CPS data.

5 Results

5.1 Shocks to tasks performable by robots

In this section we present an experiment in which there is a set of unexpected (i.e., MIT) technological shocks, one in each occupation, that increase the set of tasks that robots can perform. For illustrative purposes, we start with the *text shocks*. The size of the shock is different across occupations and it is given by what is described in section 4.2. Thus, this experiment allows us to study the quantitative effects of an innovation that allows to suddenly employ, at zero cost, a technology in which the tasks performed by robots are more than with the previous technology.

Figure 3 reports the result for aggregate variables. The shock to tasks creates a productivity effect that increases consumption, investment, the stock of robots and output along the transition to the new steady state. As robots are more productive, there is a larger demand for them and, along the transition, their rental price increases.

Task-based technological change fosters a rise in steady-state unemployment. This is because the reduction in the fraction of workers in low- and middle-skilled occupations is not fully compensated by the increase in the fraction of workers in high-skilled occupations, as shown in Figure 4. However, transitional unemployment *declines*, not increases, with respect to the initial steady state. Along the transition, workers flow out of low- and middle-skilled occupations and into high-skilled occupations, as training time in high-skilled occupations increases. However, the transition out of the first two is relatively slow. By this, we mean that the household could, in principle, speed up such process, by setting the amount of training in low- and middle-skilled occupations to zero. However, although there is a clear drop in training in these two occupations, as shown by the panels in the second row of Figure 4, this drop is equal or less than 10% of the initial steady state level.

As robots increase in all occupations, total labor declines in low- and middle-skilled roles, where the two inputs are substitutes, but increases in the high-skilled occupation, where they are complements. However, all worker types reduce their labor time and see increases in their real wages, both along the transition and in the new steady state.




Figure 3: Shock to \hat{i}_k in all occupations: aggregate variables.
 Note: Time zero is the initial steady state. All variables are in levels. Price refers to the rental price of robots.




Figure 4: Shock to \hat{i}_k in all occupations.

Note: Row 1: workers, $N_{k,t}$; row 2: time spent in training $l_{k,t}^s$; row 3: time spent in productive work, $l_{k,t} - l_{k,t}^s$; row 4: robots $r_{k,t}$; row 5: total labor $N_{k,t} (l_{k,t} - l_{k,t}^s)$; row 6: real wages $w_{k,t}$. Time zero is the initial steady state. All variables are in levels.

Figures 5 and 6 provide comparisons between the *text shocks* and the *GPT shocks*. With the latter, the new steady state is substantially further away than with the former. However, the shape of the transition is similar for most variables. An important exception is unemployment. Overall, with a larger shock, unemployment increases both along the transition and in steady state, but hours worked fall drastically. One marked difference between the two transitions is that, with a large task shock, time spent working experiences a large fall, even for high-skilled occupations. Nonetheless, this fall is compensated by more training of new high-skilled workers, resulting in a larger size of high-skilled occupations and total hours worked in those occupations (see row 1 and 5 of Figure 6). Thus, our model predicts that, in response to large task shocks, unemployment would increase, but there would be a substantial decline in hours worked in all occupations, with a large increase in leisure time.

Figure 7 reports the welfare of individuals in specific occupations and welfare of the household. For readability, the figure assumes that the shock occurs in period 0. Both the small and the large shocks to \hat{i}_k produce an increase in welfare for all types of workers, on impact, along the transition and in the new steady state.³⁴

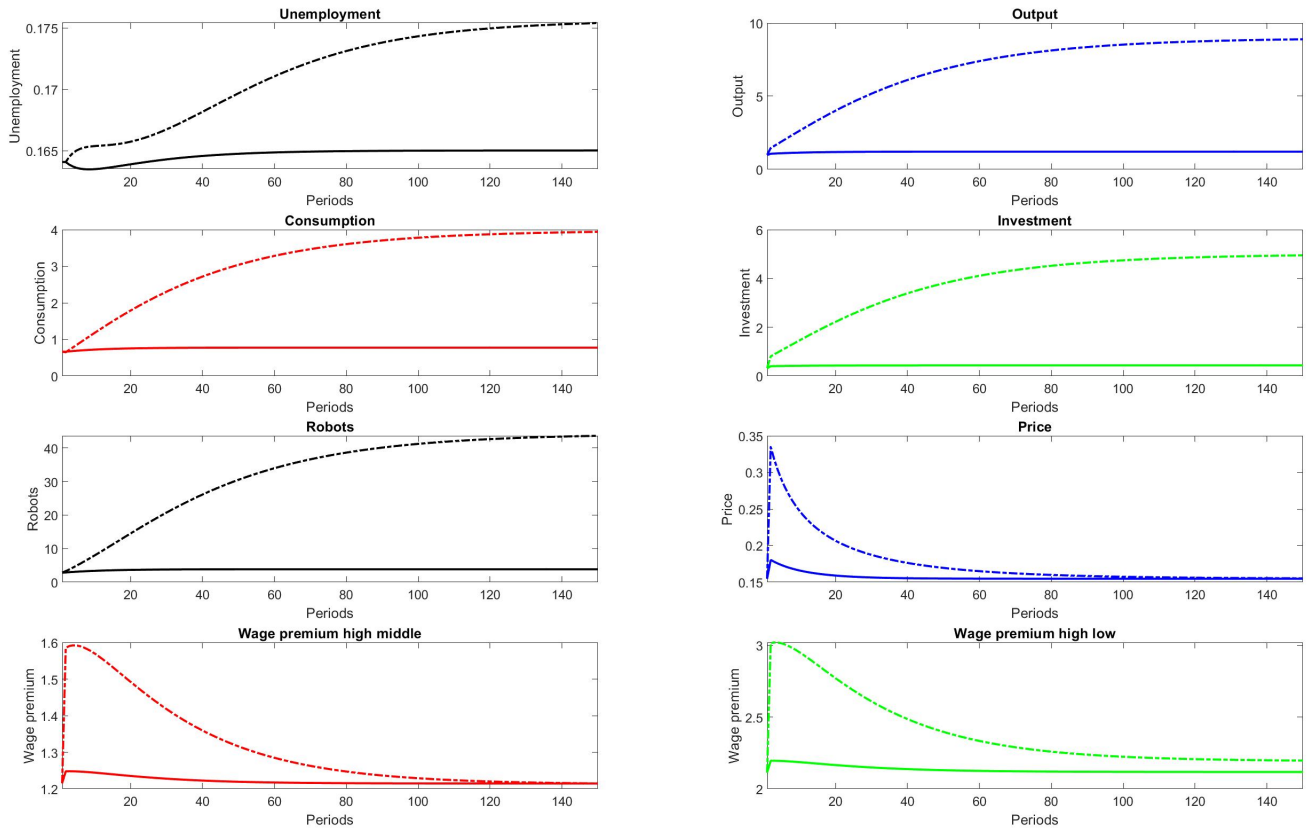


Figure 5: Shock to \hat{i}_k in all occupations: aggregate variables.

Note: The solid lines refer to the *text shock* (lower-bound) and the dashed lines are the *GPT shock* (upper-bound). Time zero is the initial steady state. All variables are in levels. Price refers to the rental price of robots.

³⁴A potential concern here is that these results might be driven by the perfect consumption insurance provided by the household to all types of agents, regardless of the occupation. The result might not hold in a decentralized setting in which agents do not belong to a larger household, and so do not have any insurance device, thus suffering a consumption drop when unemployed. However, the results here suggest that, even in that case, there should exist a transfer scheme that can be implemented such that all agents experience a welfare increase after the shock.

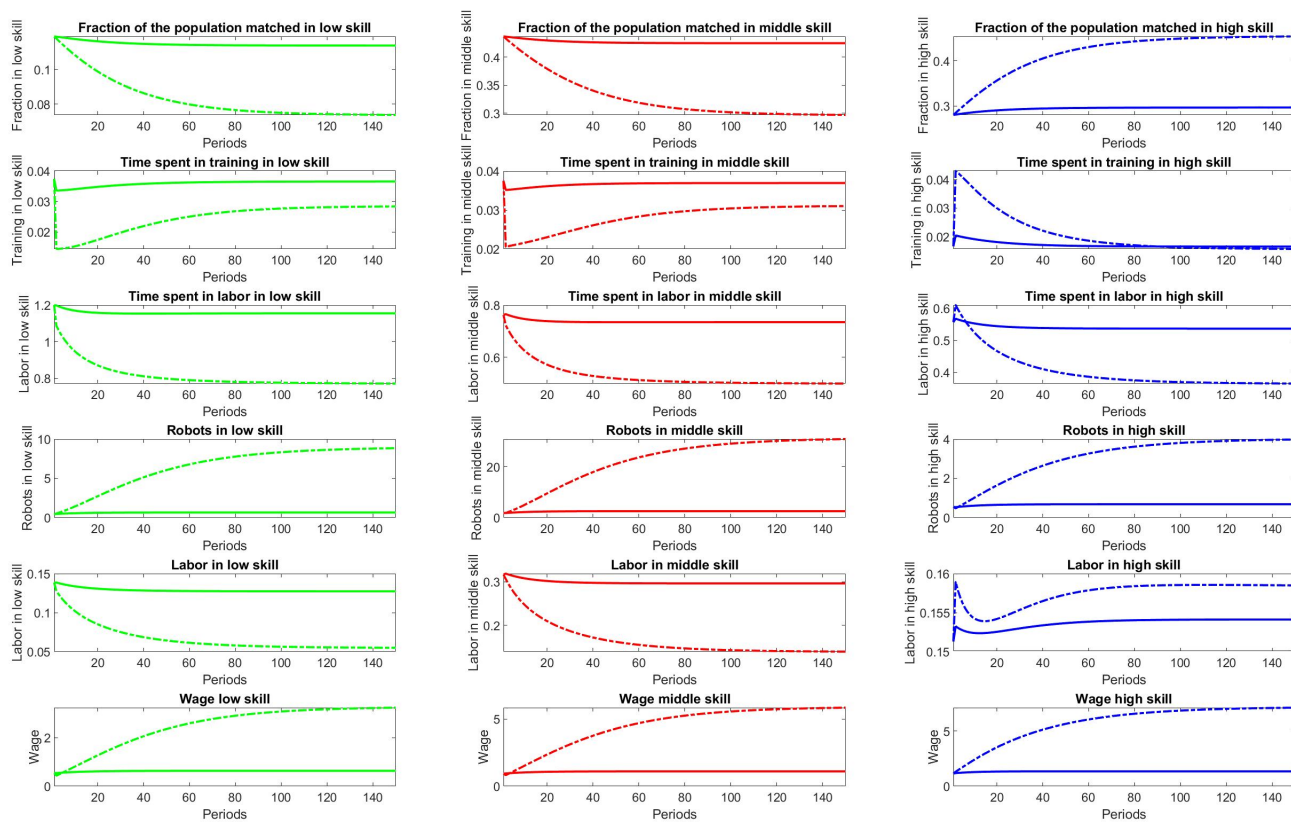


Figure 6: Shock to \hat{i}_k in all occupations.

Note: The solid lines refer to the *text shock* (lower-bound) and the dashed lines are the *GPT shock* (upper-bound). Row 1: workers, $N_{k,t}$; row 2: time spent in training $l_{k,t}^s$; row 3: time spent in productive work, $l_{k,t} - l_{k,t}^s$; row 4: robots $r_{k,t}$; row 5: total labor $N_{k,t} (l_{k,t} - l_{k,t}^s)$; row 6: real wages $w_{k,t}$. Time zero is the initial steady state. All variables are in levels.

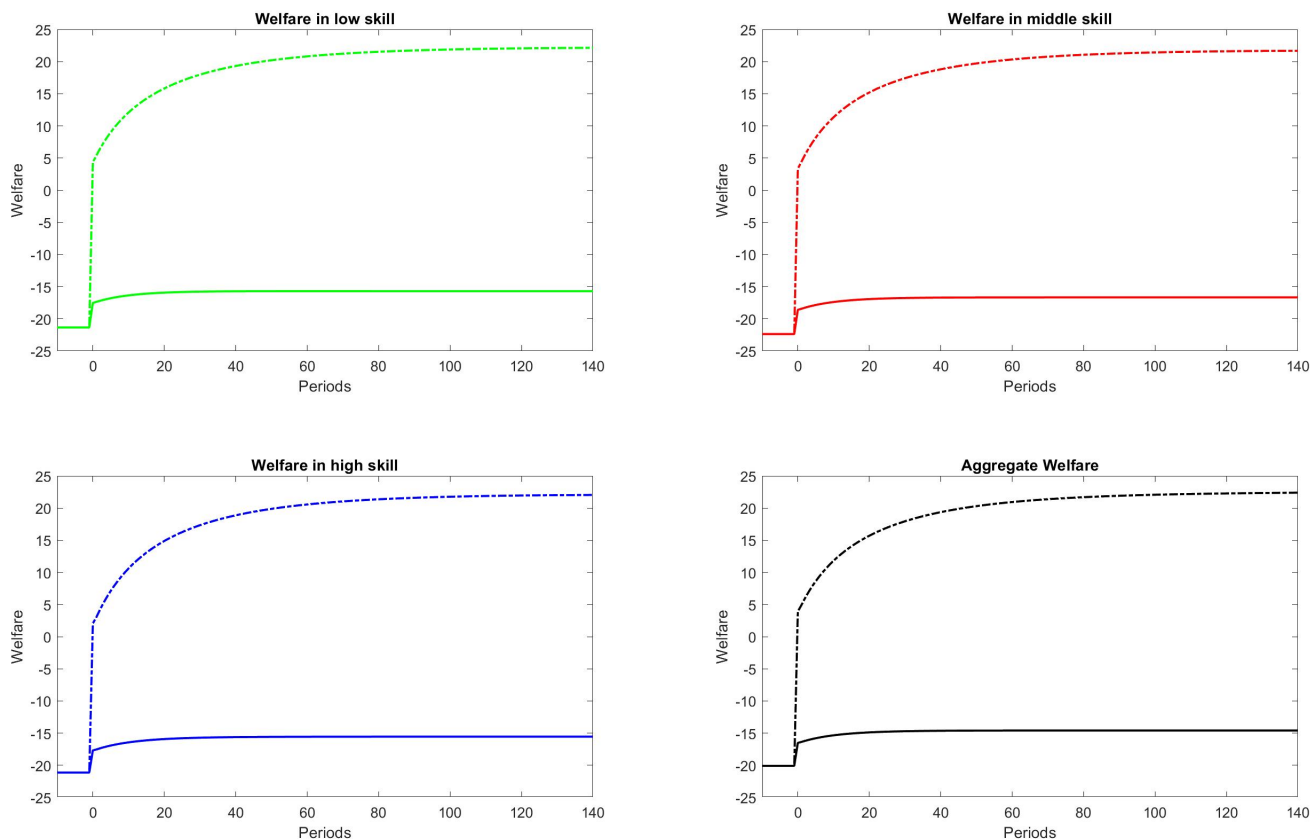


Figure 7: Welfare following a shock to \hat{i}_k in all occupations in period 0.

Note: The solid lines refer to the task shock (lower-bound) and the dashed lines are the GPT shock (upper-bound). Time zero is the initial steady state. All variables are in levels.

5.2 What's different from the past? A comparison with CETC

A concern with a task-based type technological change is that it might disrupt employment in a way that is different from the past because it makes capital/robots perfectly substitutable with specific tasks performed by workers, rather than raising the productivity of capital/robots in a homogeneous (across tasks) way, as capital embodied technological change (CETC) does. In this section we investigate this possibility by comparing the response of the economy following the task-based shocks of Section 5.1 to that following a shock to robots productivity. In the model, this is summarized by the parameter α_k in the occupation production function.

Equation (12) shows that both the fraction of tasks \hat{i}_k and the productivity of robots α_k affect the productivity of robots in occupation k . There are, however, two main differences.

First, changes in \hat{i}_k also affect the productivity of labor. Second, the effect on output of both changes depends on the elasticity of substitution between capital and labor at the occupation levels, but in different ways. More specifically, using equilibrium conditions of the model, we can rewrite total output in equilibrium as:

$$Y^* = A \left[\sum_{k=1}^J \eta_k L_k (\hat{i}_k)^\sigma \right]^{\frac{1}{\sigma}} = A \left[\sum_{k=1}^J \hat{\eta}_k [\hat{\alpha}_k r_k^{\rho_k} + L_k^{\rho_k}]^{\frac{\sigma}{\rho_k}} \right]^{\frac{1}{\sigma}}, \quad (16)$$

where $\hat{\alpha}_k = \alpha_k^{\rho_k} \left(\frac{\hat{i}_k}{(I - \hat{i}_k)} \right)^{(1 - \rho_k)}$ and $\hat{\eta}_k = \eta_k (I - \hat{i}_k)^{\frac{\sigma(1 - \rho_k)}{\rho_k}}$. In Equation (16), $\hat{\alpha}_k$ is the productivity of robots, while $\hat{\eta}_k$ is the productivity of occupation k , and the productivity of labor is normalized to one in the equilibrium output production function. A change in \hat{i}_k affects both the productivity of robots ($\hat{\alpha}_k$) and the productivity of the occupation ($\hat{\eta}_{k,t}$) of equilibrium output Y^* . Instead, a change in α_k affects only the productivity of robots.

To compare the effects of the different types of shocks, we first need to set the size of the two shocks. We proceed as follows. We set the shocks to \hat{i}_k as in section 5.1 (see Table 4). Then, we set the shock to α_k such that the occupational output $L_j(\hat{i}_j)$ in the new steady state is the same as the one generated by the change in \hat{i}_k . In this way, the change in the occupational output in the new equilibrium is equivalent to that induced by the shock on \hat{i}_k .³⁵

Results are reported in Figure (8). While both shock types are set to produce the same output level, consumption and unemployment are higher with the productivity shock. In contrast, investment and the robot stock in both the transition and in the new steady state are lower. The shock to \hat{i}_k first leads to a visible fall in unemployment along the transition, and then an increase in the new steady-state. In contrast, shocks to α_k lead to a very small (close to zero) fall and then a larger *increase* in unemployment. In summary, the model predicts that standard capital embodied technological change produces more unemployment than task-based technological change, not less, both along the transition and in the new steady state.

Figure (9) allows to investigate why this is the case. With the shocks to α_k , real wages increase in the three occupations more than with the shocks to \hat{i}_k . Consequently, time spent in training increases less in the first case for the high-skilled occupation - the occupation receiving unemployed workers along the transition. In the short run, this is offset by a smaller fall in the training-time to the low-skilled occupation. As a consequence, aggregate

³⁵Technically, we have to solve a three-targets/three-parameters calibration problem. Given all the parameters of the model, as calibrated in Section 4, we need to find the values for the three shocks to the α_k 's such that the model replicates the same occupational outputs of the new steady state obtained in the experiment in which we shock the \hat{i}_k .

unemployment does not fall with a shock to α_k , even on impact. But unemployment increases more as aggregate training time falls and total hours worked in the high-skilled occupation remain roughly flat after an initial fall.

Overall, capital embodied technical change leads to a larger welfare increase than task biased technical change (not reported).³⁶ However, the shocks to \hat{i}_k generate larger inequality along the transition, as measured by wages. The bottom row of Figure 8 reports the evolution of the wage premium between high- and middle-skilled and between high- and low-skilled agents. The shocks to \hat{i}_k increase the former ratio by almost 4%, and the second by 2.7%. The shocks to α_k , instead, induces an increase of 2% and 2.2% respectively.³⁷

³⁶The welfare comparison between the \hat{i}_k shock and the α_k shock is similar to the one reported in Figure 7 between the two shocks to \hat{i}_k . Welfare is larger with the α_k shocks both on impact, along the transition and in the steady state.

³⁷Regarding the wage premium between middle- and low-skilled, the shocks to \hat{i}_k induce an increase of 1.2%, while the shocks to α_k do not increase it. This is not reported in Figure 8.

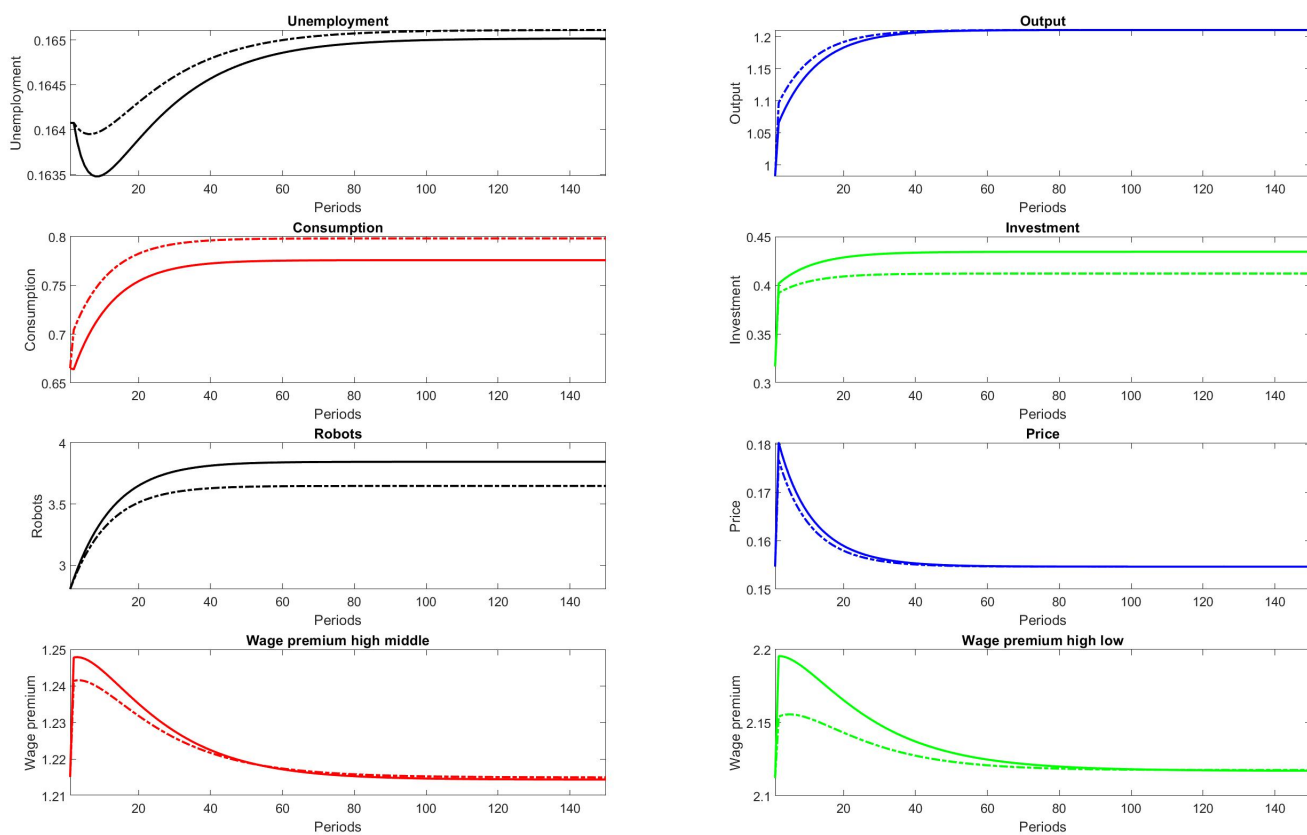


Figure 8: Comparison of a shock to \hat{i}_k in all occupations versus a shock to α_k in all occupations: aggregate variables.

Note: The solid lines refer to the lower-bound shock through text mapping and the dashed lines are the counterfactual CETC shock. Time zero is the initial steady state. All variables are in levels. Price refers to the rental price of robots.




Figure 9: Comparison of a shock to \hat{l}_k in all occupations versus a shock to α_k in all occupations: occupations.

Note: The solid lines refer to the lower-bound shock through text mapping and the dashed lines are the counterfactual CETC shock. Row 1: workers, $N_{k,t}$; row 2: time spent in training $l_{k,t}^s$; row 3: time spent in productive work, $l_{k,t} - l_{k,t}^s$; row 4: robots $r_{k,t}$; row 5: total labor $N_{k,t} (l_{k,t} - l_{k,t}^s)$; row 6: real wages $w_{k,t}$. Time zero is the initial steady state. All variables are in levels.

5.3 What's different from the past? The role of barriers

In this paper, we highlight that training requirements increased over time in the U.S. for high-skilled occupations. This observation represents another difference with respect to the past that has the potential to generate technological unemployment. To investigate this issue, in this section we compare the case described in the previous section to one in which, in addition to the shocks to tasks, there is an unexpected fall in ϕ_h . This parameter governs how difficult it is to train workers to enter the high-skilled occupation. The smaller ϕ_h , the more important the training time of workers already in the occupation, and the less important the current unemployment pool in producing new labor for occupation h . We set

the shock to $\Delta\phi_h = -0.239$ from the estimated change in Equation (15).

Results are reported in Figure 10. The larger training parameter makes unemployment increase, both in the transition and in the steady state. From Figure 11, we observe that the larger training requirement substantially changes the occupational structure in the steady state. The high-skilled occupation is shrinking rather than expanding, and the opposite happens to the middle-skilled one. Recall that the robot shock is largest in the low-skilled occupation. Therefore, in this instance, unemployment is not caused by workers moving from low- and middle-skilled occupations to high-skilled ones. Rather, unemployment is caused by a fall in the training of low- and middle-skilled workers that cannot be offset by the increase in training-time dedicated to high-skilled workers. That is, even though the training of high-skilled workers rises, the entry rate into high-skilled occupations falls in the face of increasing training costs. While the middle-skilled occupation absorbs some of the newly unemployed, this is not enough to prevent a higher unemployment rate. This result highlights the role of training in shaping the occupational structure, both in the steady state and along the transition.




Figure 10: Shock to \hat{i}_k in all occupations and shock to training costs in the high-skilled occupation $\Delta\phi_h = -0.239$: aggregate variables.

Note: The solid lines refer to the *text shock* and the dashed lines are the counterfactual high-training cost transition. Time zero is the initial steady state. All variables are in levels. Price refers to the rental price of robots.

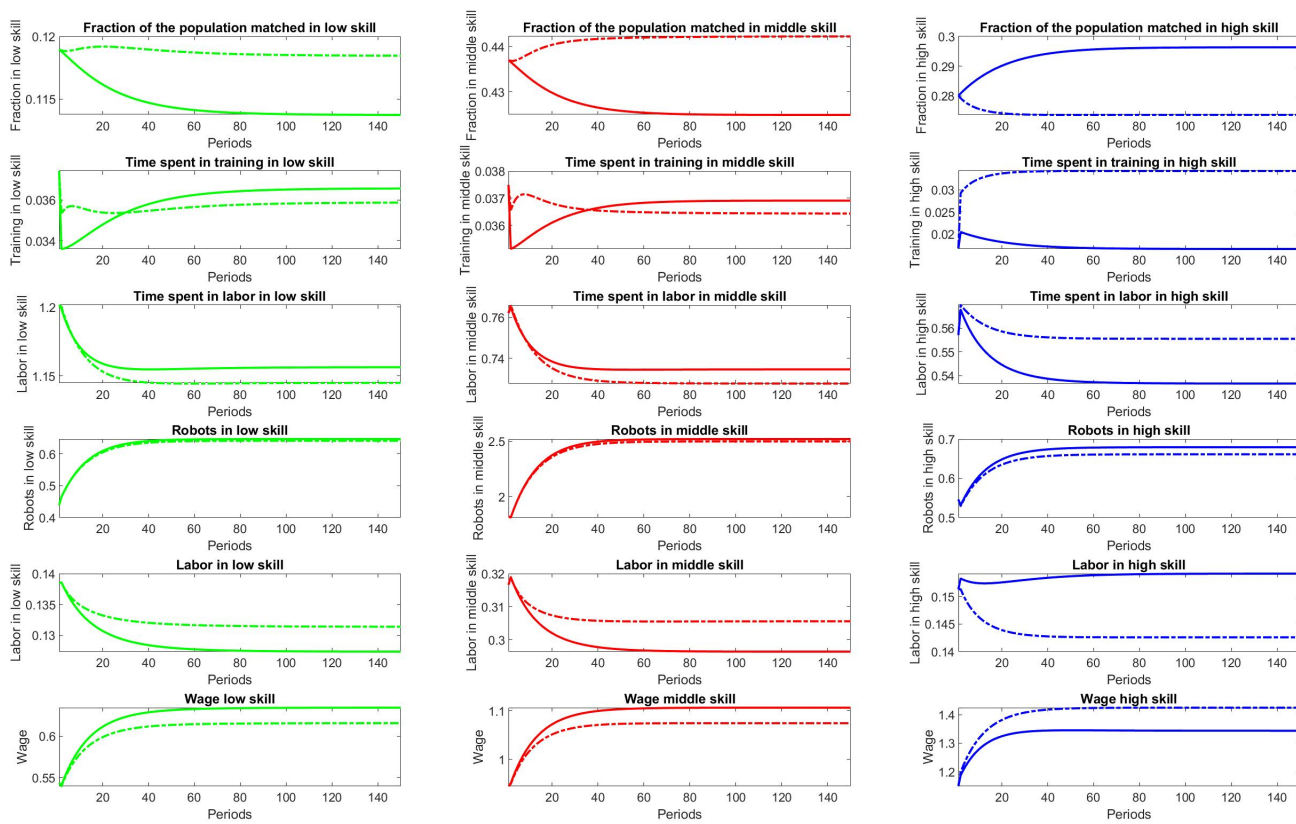


Figure 11: Shock to \hat{i}_k in all occupations and shock to training costs in the high-skilled occupation $\Delta\phi_h = -0.239$.

Note: The solid lines refer to the *text shock* and the dashed lines represent the counterfactual in which the text shock occurs together with the training shock. Row 1: workers, $N_{k,t}$; row 2: time spent in training $l_{k,t}^s$; row 3: time spent in productive work, $l_{k,t} - l_{k,t}^s$; row 4: robots $r_{k,t}$; row 5: total labor $N_{k,t} (l_{k,t} - l_{k,t}^s)$; row 6: real wages $w_{k,t}$. Time zero is the initial steady state. All variables are in levels.

Figures 12 and 13 show the same experiment for the *GPT shock*. In this instance, the two transitions look even closer in terms of output, consumption, investment and the stock of robots. However, the pattern for unemployment is substantially different, with first a drop and then an increase along the transition, before reaching a new higher level steady state. The reason is that, with the *GPT shock*, the increase in time spent in training is so substantial that it dominates the shock to training costs. The large increase in training time in high-skilled occupations is enough to increase the size of this occupation in the periods after the shocks, in a similar vein as in the case without the training shock. On the other hand, the decrease in training time in low- and middle-skilled occupations is less pronounced compared

to the case without the training shock. As a result, in the first part of the transition, the growth in the high-skilled occupation is similar to the case without the training shock, while the decline in the other two occupations is slower, and so unemployment declines. However, in the new steady state, the size of the high-skilled occupation is smaller with the training shock and unemployment is higher. The other marked difference is in the total hours worked of high-skilled workers along the transition (see Figure 13, row 5). A large task shock leads to large robot investment, which offsets the productivity loss of shifting workers' production hours to training hours even in the face of higher training costs.

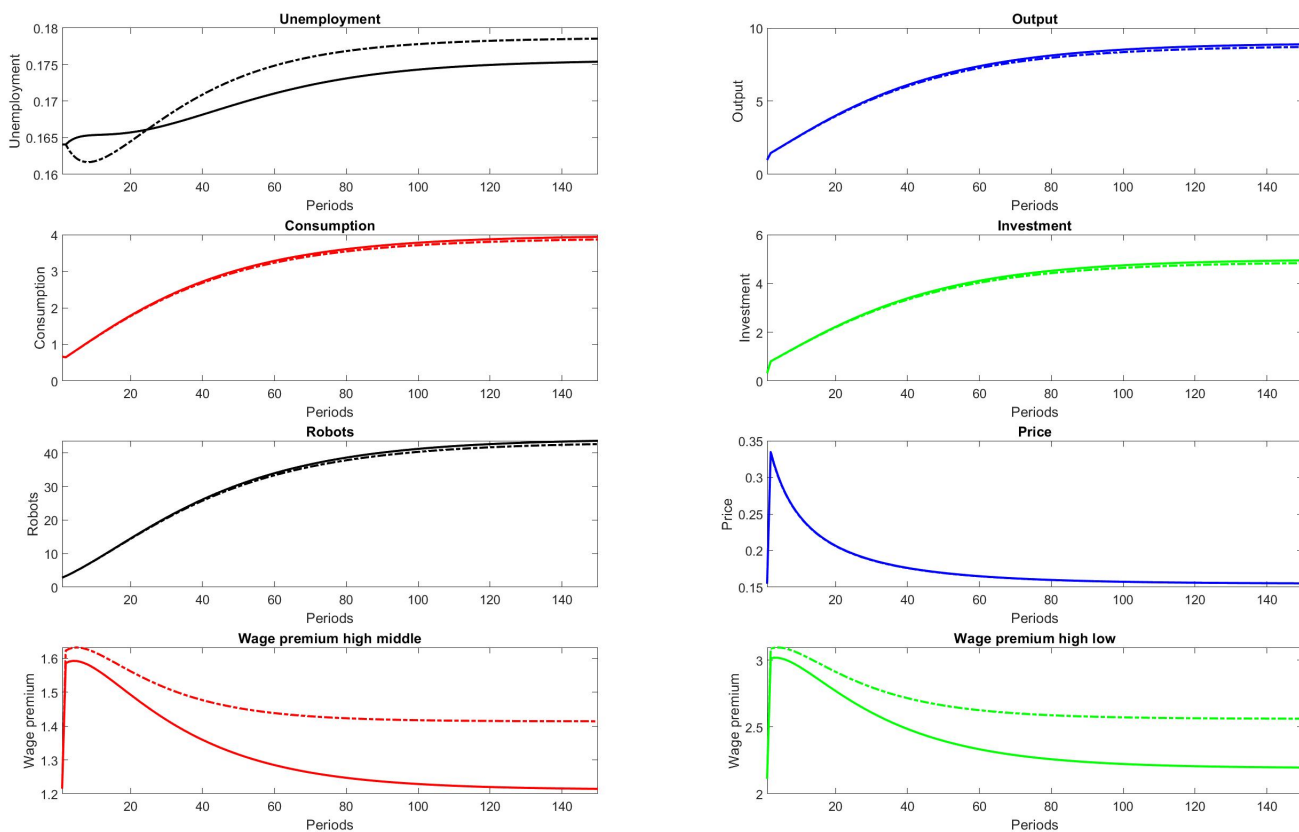


Figure 12: Shock to \hat{i}_k in all occupations and shock to training costs in the high-skilled occupation $\Delta\phi_h = -0.239$: aggregate variables.

Note: The solid lines refer to the *GPT shock* and the dashed lines are the counterfactual high-training cost transition with the same *GPT shock*. Time zero is the initial steady state. All variables are in levels. Price refers to the rental price of robots.

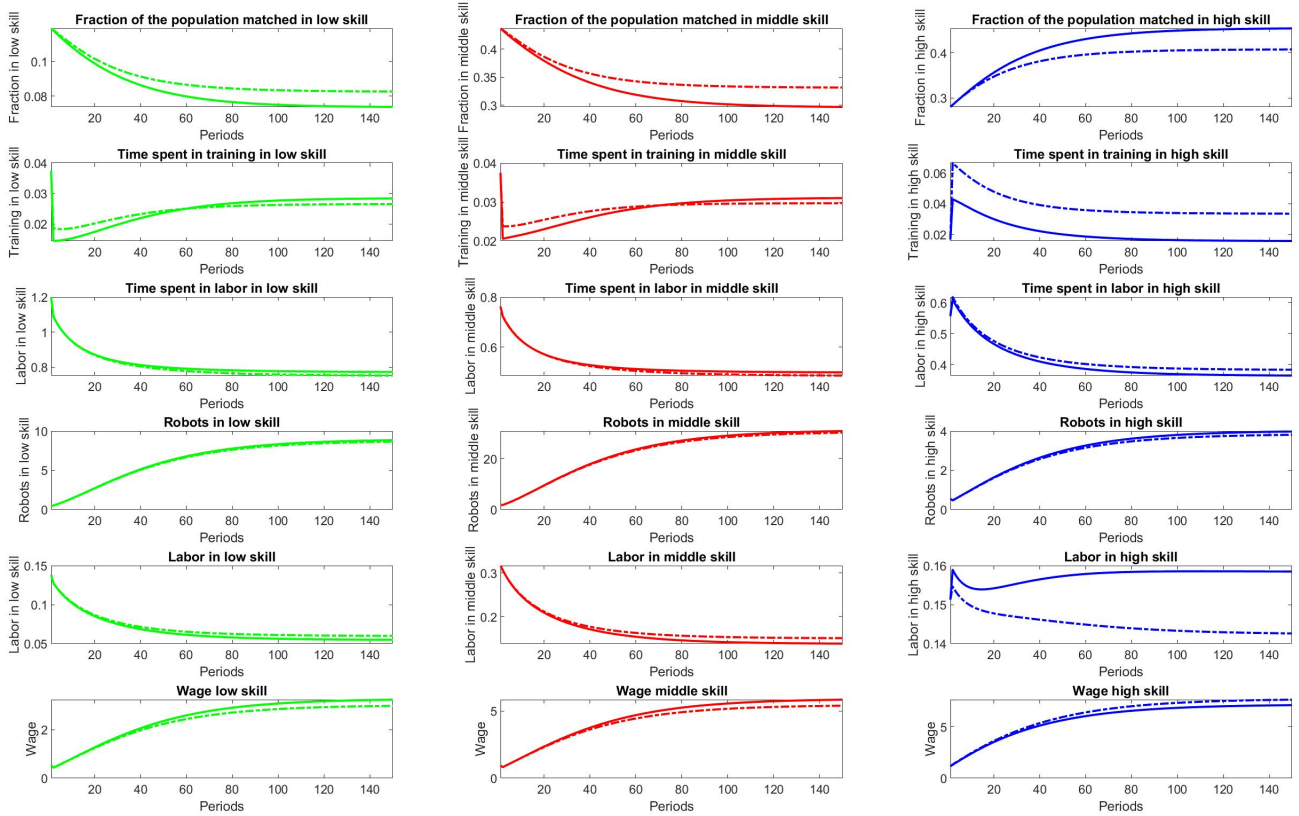


Figure 13: Shock to \hat{i}_k in all occupations with higher training costs $\Delta\phi_h = -0.239$ occupations.

Note: The solid lines refer to the *GPT shock* and the dashed lines represent the counterfactual in which the GPT shock occurs together with the training shock. Row 1: workers, $N_{k,t}$; row 2: time spent in training $l_{k,t}^s$; row 3: time spent in productive work, $l_{k,t} - l_{k,t}^s$; row 4: robots $r_{k,t}$; row 5: total labor, $N_{k,t} (l_{k,t} - l_{k,t}^s)$; row 6: real wages $w_{k,t}$. Time zero is the initial steady state. All variables are in levels.

6 Conclusions

In this paper we studied how the interaction between automation and rising training requirements can shape the emergence of technological unemployment. We documented that the training time required to perform high-skilled occupations in the U.S. has risen substantially between 2006 and 2019, while requirements for middle- and low-skilled occupations have remained broadly flat. To interpret this evidence, we built and calibrated a dynamic multi-occupation growth model in which occupations are sets of tasks, robots can perform a subset of those tasks, and the reallocation of workers across occupations is shaped by an

occupation-specific training technology.

A key prediction of the model focuses on what happens when a task-displacing technological shock coincides with rising training barriers in high-skilled occupations. The unemployment rate rises both along the transition and in the new steady state, and the occupational structure shifts in a way that task shocks alone would not generate. The high-skilled occupation shrinks rather than expanding, because higher training costs dampen entry into that occupation even if training time increases. Unemployment becomes the byproduct of workers flowing out of low- and middle-skilled occupations faster than the high-skilled occupation can absorb them.

Two other counterfactuals clarify why this combination matters and why the current environment might be different relative to earlier waves of technological change. First, comparing a task shock to a more conventional capital-embodied productivity shock calibrated to obtain the same steady-state output in both the initial and final equilibria, we find that the former generates less transitional unemployment, not more: the productivity shock raises real wages by more, which dampens the incentive to train new high-skilled workers and slows reallocation. Second, in the absence of rising training barriers, even a task shock produces a temporary decline in unemployment along the transition, because the household optimally balances the flow out of shrinking occupations against costly training in expanding ones. Thus, in our model it is the combination of task displacement and rising barriers that turns automation into persistent technological unemployment.

These findings have implications that extend beyond labor market aggregates. A well-established literature has documented that job displacement and prolonged unemployment carry severe health consequences, from elevated mortality following displacement (Sullivan and Von Wachter (2009)) to the rise in deaths of despair among less-educated Americans (Case and Deaton, 2015, 2017). Crucially related to our research question, O'Brien, Bair, and Venkataramani (2022) provide causal evidence that robot adoption in U.S. local labor markets has raised working-age mortality, operating through two channels: a direct material effect on displaced workers, and an indirect effect working through the erosion of future economic prospects in affected communities. The mechanics in our model are reminiscent of some of these aspects. The transitional dynamics following an automation shock speak to the material cost of displacement, while the higher steady-state unemployment generated by rising training barriers informs on the structurally diminished prospects of reentry. Training barriers thus emerge as a key friction linking automation to documented health consequences, providing a rationale for why automation shocks can produce the persistent labor market disruptions tied to deteriorating population health.

References

- ACEMOGLU, D., AND D. AUTOR (2011): “Chapter 12 - Skills, Tasks and Technologies: Implications for Employment and Earnings,” vol. 4, Part B of *Handbook of Labor Economics*, pp. 1043 – 1171. Elsevier.
- ACEMOGLU, D., AND P. RESTREPO (2020): “Robots and Jobs: Evidence from US Labor Markets,” *Journal of Political Economy*, 128(6).
- (2022a): “Demographics and automation,” *The Review of Economic Studies*, 89(1), 1–44.
- (2022b): “Tasks, automation, and the rise in US wage inequality,” *Econometrica*, 90(5), 1973–2016.
- AUTOR, D. H., AND D. DORN (2013): “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market,” *American Economic Review*, 103.
- AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): “The skill content of recent technological change: An empirical exploration,” *Quarterly journal of economics*, 118, 4.
- BERAJA, M., AND N. ZORZI (forthcoming): “Inefficient Automation,” *The Review of Economic Studies*.
- CARRILLO-TUDELA, C., AND L. VISSCHERS (2023): “Unemployment and endogenous reallocation over the business cycle,” *Econometrica*, 91(3), 1119–1153.
- CASE, A., AND A. DEATON (2015): “Rising morbidity and mortality in midlife among white non-Hispanic Americans in the 21st century,” *Proceedings of the National Academy of Sciences*, 112(49), 15078–15083.
- (2017): “Mortality and morbidity in the 21st century,” *Brookings papers on economic activity*, 2017, 397.
- CASELLI, F. (1999): “Technological revolutions,” *American economic review*, 89(1), 78–102.
- CAUNEDO, J., D. JAUME, AND E. KELLER (2023): “Occupational exposure to capital-embodied technical change,” *American Economic Review*, 113(6), 1642–1685.
- CERINA, F., E. DIENESCH, A. MORO, AND M. RENDALL (2023): “Spatial polarisation,” *The Economic Journal*, 133(649), 30–69.

- CERINA, F., A. MORO, AND M. RENDALL (2021): “The role of gender in employment polarization,” *International Economic Review*, 62(4), 1655–1691.
- COSSU, F. (2025): “The Skill Inside the Task: How AI and Robotics Reshape the Structure of Work,” Discussion paper.
- CUTLER, D., A. DEATON, AND A. LLERAS-MUNEY (2006): “The determinants of mortality,” *Journal of economic perspectives*, 20(3), 97–120.
- DVORKIN, M. A., AND A. MONGE-NARANJO (2019): “Occupation mobility, human capital and the aggregate consequences of task-biased innovations,” *FRB St. Louis Working Paper*, (2019-13).
- GOLDIN, C. D., AND L. F. KATZ (2008): *The race between education and technology*. harvard university press.
- GUVENEN, F., B. KURUSCU, S. TANAKA, AND D. WICZER (2020): “Multidimensional skill mismatch,” *American Economic Journal: Macroeconomics*, 12(1), 210–244.
- KRUSELL, P., L. E. OHANIAN, J.-V. RÍOS-RULL, AND G. L. VIOLANTE (2000): “Capital-skill complementarity and inequality: A macroeconomic analysis,” *Econometrica*, 68(5), 1029–1053.
- MARIMON, R., AND F. ZILIBOTTI (1999): “Unemployment vs. mismatch of talents: Reconsidering unemployment benefits,” *The Economic Journal*, 109(455), 266–291.
- MODESTINO, A. S., D. SHOAG, AND J. BALLANCE (2020): “Upskilling: Do employers demand greater skill when workers are plentiful?,” *Review of Economics and Statistics*, 102(4), 793–805.
- MOLL, B., L. RACHEL, AND P. RESTREPO (2022): “Uneven growth: automation’s impact on income and wealth inequality,” *Econometrica*, 90(6), 2645–2683.
- MORRIS, B. (2023): “How robots are taking over warehouse work,” *BBC News*, <https://www.bbc.com/news/business-68639533>, Accessed: 2024-06-18.
- O’BRIEN, R., E. F. BAIR, AND A. S. VENKATARAMANI (2022): “Death by robots? Automation and working-age mortality in the United States,” *Demography*, 59(2), 607–628.
- PHELAN, C., AND A. TREJOS (2000): “The aggregate effects of sectoral reallocations,” *Journal of Monetary Economics*, 45(2), 249–268.

- SULLIVAN, D., AND T. VON WACHTER (2009): “Job displacement and mortality: An analysis using administrative data,” *The Quarterly Journal of Economics*, 124(3), 1265–1306.
- TEUBERT, M., M. RENDALL, AND D. DOWE (2024): “Frey and Osborne Revisited: Predicting Task Automation Using LLMs,” Working paper.
- TINBERGEN, J. (1974): “Substitution of graduate by other labour,” *Kyklos: international review for social sciences*.

Appendix

A Training Index

The “Education, Training, and Experience” file from the O*NET database provides frequency data related education, experience, and training levels required for each occupation. Thus, for each element — education, work experience, and training — we have the distribution of levels for specific occupations.

Education: Measures average years of formal education. For example, a Bachelor’s Degree requires 16 years of education, which includes 12 years towards a high-school diploma and 4 years of college. We compute education time based on the years needed to achieve each education level above a high-school diploma. That is, we subtract 12 years of education, the education required to obtain a high school diploma, from each individual’s total years of education to obtain a “required” training time for that individual’s occupation. This allows us to have an education index that refers to the years of education necessary for a specific occupation.

Related Work Experience: Each level specifies an interval of required work experience. For example, for the level specifying *Over 1 year, up to and including 2 years* we take the maximum time of work experience required. Therefore, in this example, two years of work experience is required. The work experience time is estimated accordingly for each category.

On-the-Job Training and On-Site Training: Similarly to work experience, on-the-job and on-site training levels record intervals for the required training time. The maximum time within each interval is taken for estimating the total training time required.

We first compute the training time for each element — education, work experience, and training — and each occupation, using formula (1) in the main text. This provides indices of education, related work experience, on-the-job training, and on-site training for each occupation. Next, summing the four indices at the occupation level, we obtain the total training time required for each occupation. This is what is labeled “training index” in the main text.

The occupational classification that we use is the O*NET-SOC code classification. To construct the training index, as described in Section 2, we use US Census data and compute the employment shares for each occ1990 occupation.

We begin our analysis in 2006, as this is the first year in which the O*NET database provides a sufficient number of occupations to compute training indices for the majority of (at least 300) occ1990 occupations (using for education, training, and experience file).³⁸

³⁸For instance, considering O*NET in 2005 implies that only 220 occ1990 occupations are represented by

Constructing the aggregate training index requires linking occupational information from IPUMS (occ1990 classification) with the O*NET-SOC classification. To do so, we rely on the crosswalks provided by [Guvenen, Kuruscu, Tanaka, and Wiczer \(2020\)](#) which map occ1990 codes into occ2000 and occ2000 into SOC2000. These datasets are crucial, as they enable the crosswalk from occ1990 to occ2000, occ2000 to soc2000, and finally soc2000 to O*NET-SOC, thereby establishing the association between the occ1990 classification and the O*NET-SOC classification. This association represents the starting point that allows us to calculate the aggregate training index. Since O*NET occupations are defined at a more disaggregated level than Census occupations and a single occ1990 occupation may correspond to multiple O*NET-SOC occupations, we compute the training index for each occ1990 occupation as the average of the corresponding O*NET-SOC training measures, resulting in a single index per occ1990 code. In the final step we compute the aggregate training time index, a weighted average where the training time for each occupation is weighted by its employment share. This process is repeated for each year, providing an estimate of the evolution of the aggregate index over time.

Not all occ1990 occupations can be matched to O*NET-SOC codes through this procedure, due to differences in classification systems and the greater level of detail in O*NET. The IPUMS occ1990 classification includes approximately 389 occupations, and we are able to construct the training index for 345 of them. Importantly, our empirical analysis relies on the set of occupations used to compute employment shares over the period 2006–2019. For this set of occupations, we obtain complete coverage: all occ1990 occupations included in the construction of employment shares are successfully matched to O*NET-SOC codes.

As a result, the construction of the training index does not involve any loss of observations within the sample used in the analysis. While the mapping does not cover the full set of possible occ1990 codes, it fully captures the occupational structure relevant for our empirical exercise.

Rationale for using all indices in constructing the training index

To construct the training index as described in Section 2, we use all elements in the Education, Training, and Experience file from O*NET: education, work experience, on-site or in-plant training, and on-the-job training. We use all variables instead of only work experience and training, or only training, as some occupations require a higher level of specific education to be performed, compared to others that may require more practical training or work experience instead of a higher level of education. For example, a qualification of chartered accountancy is acquired through a formal education course, while learning how to

the O*NET classification.

Occupations	Complete Training		Training no educ.		Training no educ. no work experience	
	2006	2019	2006	2019	2006	2019
Managers	10.53	11.74	7.65	8.15	2.95	2.99
Professionals	9.85	10.93	5.28	5.78	2.01	2.42
Technicians	7.93	7.33	5.31	4.57	2.11	1.81
Sales	4.73	4.10	3.69	3.38	1.24	1.08
Admin Services	4.24	4.63	3.29	3.54	1.02	1.20
Low Skill Services	2.71	2.50	2.68	2.43	1.14	1.05
Mechanics	5.21	4.56	5.32	4.52	2.74	2.42
Transport Precision Workers	6.03	6.86	5.56	6.02	2.66	2.28
Machine Operators	2.40	3.54	2.65	3.53	1.16	1.58

Table 6: Training index by nine occupations.

Note: Training constructed for nine broad occupational categories and using various methodologies.

operate a machine is more likely done at a plant. To put it differently, the importance of the four elements is heterogeneous across occupations.

Table 6 highlights this heterogeneity in types of learning by presenting the constructed training time for nine broad occupational categories and using various methodologies. "Complete training" refers to the training constructed as in Section 2, the training constructed by excluding the education element and, finally, the training constructed by excluding both education and work experience (i.e., considering only on-site or in-plant training and on-the-job training). For professional and technicians occupations, the training required when including the education element (9.85 and 7.93 years, respectively in 2006) is substantially higher than that required for precision workers (6.03 years). However, when education is excluded and the training is constructed by considering only work experience and training, or only the training elements, training required is greater for precision workers (5.56 year of training with no education in 2006) compared to professionals and technicians (5.28 and 5.31 years of training with no education in 2006). This suggests that all elements are essential to construct an accurate and comprehensive aggregated training index, as each element significantly contributes to capturing the different training requirements across different occupations.

B Future Tasks

We construct five indicators to assess the potential substitutability of Intermediate Work Activities (IWA) with robots. These indicators differ based on the methodologies used to map tasks to robots’ applications. Our findings indicate that the number of IWA that can be potentially performed by robots depends on the design of these mappings. In section 5 we present results using two of these five measures.

The first measure, ”Map with Condition and Second Search,” is based on the specific descriptions provided by the International Federation of Robotics (IFR). This mapping aims to align the very core of a robot’s function with an occupation’s tasks. However, we note that robot applications often have nuanced specifications. Some robots, for example, are designed to operate on or with specific materials or objects. Thus, in our first mapping, we impose a condition that checks this specific material or object appears both in the robot’s application and in the occupational task. For instance, for the ”welding” domain, we focus on subcategories such as arc welding, spot welding, and others. Our search criteria for this area are “(weld| welding) & (arc|spot|laser| ultrasonic| gas| plasma)”. We classify only tasks containing these terms as those potentially performed by welding robots. However, some tasks also require higher-order skills like “supervising” to be performed. We first assume these tasks are beyond robotic capabilities and refine our results conducting a second search round. Here, we eliminate tasks that contain actions like “supervise; administer; direct; oversee; coordinate; document; investigate; negotiate; employee training” from robots’ capabilities. Results are reported in Table 7 and Figure 14. We consider this mapping as the lower bound of potential substitutability between robots and IWA, and use it as the benchmark in the quantitative part of the paper, labeling it as *Text shock*.

We then assume that robots are already able to perform higher-order skills like supervising, and we remove the second round search from the previous methodology. This gives the “Map with condition - no second search” result in Table 7 and Figure 14. Perhaps not surprisingly, the largest difference compared to the previous measure is observed for high-skilled occupations.

We then construct a measure of the IWA potentially performable by robots by exploiting the mapping between O*NET tasks and IFR applications according to Large Language Models (LLMs) from Teubert, Rendall, and Dowe (2024). In this case, the LLM engine is asked, in several ways, whether a certain task can be performed by a specific type of robot. The authors run whole combinations of tasks and IFR applications with different queries on several LLMs. They then pick the query that gives the most consistent results across all LLMs and the training (hand-labeled) data. This methodology results in a substantially larger

fraction of IWA potentially substitutable by robots than the previous two methodologies, as reported in Table 7 and Figure 14. This is used to construct the *GPT shock* in section 4.2.

To investigate why the results obtained with the LLM are magnitudes larger than with the first two methodologies, we modify our first mapping along two dimensions. First, we map tasks to robot applications without considering specific conditions, but including the second search. This means we separate the action from the object. For example, we consider the word "cleaning" without conditioning it to specific contexts like "windows" or "floors". This approach results in a measure of robot substitutability that is substantially larger than the one obtained when considering specific conditions, and more aligned with the results obtained with the LLMs for low-skilled occupations. However, it is substantially larger for high-skilled occupations.

Second, we neither consider specific conditions nor refine the analysis with a second search. This means we use only generic terms for the mapping (e.g., welding or cleaning) and assume robots can also perform higher-order skills, like supervising and administering, as in the second methodology. This yields the largest share of IWA that can potentially be performed by robots, although, again, the largest effect with respect to the previous methodology (map no condition but with second search) appears to be on high-skilled occupations.

Our analysis reveals that the methodology used to assess robot substitutability greatly influences the outcomes. The design of these mappings plays a crucial role in determining the extent to which robots can already perform specific tasks. We thus conclude this section with an example of the mapping between occupational tasks and robot capabilities, Table 8 reports a set of representative examples linking O*NET occupations, tasks, and robot applications. For each occupation, we present selected tasks that are classified as potentially automatable, along with the corresponding Intermediate Work Activities (IWAs) and the type of robot application that could perform those tasks. The classification of robot applications is based on the International Federation of Robotics (IFR), as described in the World Robotics 2021 – Industrial Robots and Service Robots reports, which provide a detailed taxonomy of robot applications across industrial and service domains. As reported in Table 7, low-skill occupations appear to be relatively more exposed to automation compared to middle- and high-skill occupations.³⁹ This reflects the fact that our mapping includes not only industrial robots but also service robots. Since many low-skill occupations are concentrated in service activities, they often involve tasks such as cleaning, handling objects, or performing routine physical activities, that fall within the scope of service robot applications

³⁹The mapping between O*NET tasks and robot applications is based on the detailed IFR taxonomy of robot capabilities. For example, in the case of medical robotics, the classification includes robots that handle and process samples in medical laboratories. Similarly, in professional cleaning, robots are classified into specific applications such as for example, floor cleaning and window and wall cleaning.

(e.g., professional cleaning). Importantly, our approach does not imply that entire occupations are automated. Rather, the mapping operates at the task level, and only a subset of tasks within each occupation is classified as potentially automatable, while other tasks, especially those requiring interpersonal interaction or complex judgment, remain beyond current robotic capabilities.




Figure 14: Robots measures comparison

Occupation	Map with condition and second search	Map with condition - no second search	Map Chat GPT	Map no condition but with second search	Map no condition no second search
High Skill	0.01	0.03	0.15	0.41	0.54
Middle Skill	0.06	0.06	0.45	0.56	0.59
Low Skill	0.08	0.09	0.54	0.51	0.56

Table 7: Robots measures comparison

O*NET Occupation	Task	IWATitle	robot_application
Waiters and Waitresses	Perform cleaning duties, such as sweeping and mopping floors, vacuuming carpet, tidying up server station, taking out trash, or checking and cleaning bathroom.	Clean tools, equipment, facilities, or work areas.	professional_cleaning
Janitors and Cleaners	Clean windows, glass partitions, or mirrors, using soapy water or other cleaners, sponges, or squeegees.	Clean tools, equipment, facilities, or work areas.	professional_cleaning
Welders, Cutters, and Welder Fitters	Recognize, set up, and operate hand and power tools common to the welding trade, such as shielded metal arc and gas metal arc welding equipment.	Join parts using soldering, welding, or brazing techniques.	welding
Welders, Cutters, and Welder Fitters	Dismantle metal assemblies or cut scrap metal, using thermal-cutting equipment, such as flame-cutting torches or plasma-arc equipment.	Cut materials.	processing
Medical and Clinical Laboratory Technologists	Analyze samples of biological material for chemical content or reaction.	Analyze biological or chemical substances or related data.	medical_robotics
Epidemiologists	Prepare and analyze samples to study effects of drugs, gases, pesticides, or microorganisms on cell structure and tissue.	Analyze biological or chemical substances or related data.	medical_robotics

Table 8: Illustrative Examples of the Mapping Between O*NET Tasks, IWAs, and Robot Applications

C Household Problem

In this appendix, we show that each agent ι has the same amount of consumption at each t , which we can denote by c_t , and so that we can write the problem, as in the main text, considering a unique consumption level for all agents ι .

Recall that the household is made of a unit measure of agents, $\iota \in [0, 1]$, and each agent has instantaneous utility,

$$\hat{u}_{\iota,t} = \frac{c_{\iota,t}^{1-\psi} - 1}{1-\psi} - B_{\iota,k} l_{\iota,t}^\mu, \quad (17)$$

where c_ι is consumption and l_ι is the amount of labor that the agents provide to the household.⁴⁰ Each agent has total lifetime utility given by:

$$u_\iota = \sum_{t=0}^{\infty} \beta^t \hat{u}_{\iota,t}, \quad (18)$$

and utility of the household is:

$$W = \int_0^1 u_\iota d\iota. \quad (19)$$

By substituting (17) and (18) into (19), we can write:

$$W = \int_0^1 \sum_{t=0}^{\infty} \beta^t \left\{ \frac{c_{\iota,t}^{1-\psi} - 1}{1-\psi} - B_{\iota,k} l_{\iota,t}^\mu \right\} d\iota,$$

and by separating the arguments of the integral we obtain:

$$W = \int_0^1 \left\{ \sum_{t=0}^{\infty} \beta^t \frac{c_{\iota,t}^{1-\psi} - 1}{1-\psi} \right\} d\iota - \int_0^1 \left\{ \sum_{t=0}^{\infty} \beta^t B_{\iota,k} l_{\iota,t}^\mu \right\} d\iota.$$

We can now pass the integrals inside the sums to obtain:

$$W = \left\{ \sum_{t=0}^{\infty} \beta^t \int_0^1 \frac{c_{\iota,t}^{1-\psi} - 1}{1-\psi} d\iota \right\} - \left\{ \sum_{t=0}^{\infty} \beta^t \int_0^1 B_{\iota,k} l_{\iota,t}^\mu d\iota \right\}.$$

Considering that at time t : $N_{k,t}$ agents work in occupation k ; $1 - \left[\sum_{k=1}^J N_{k,t} \right]$ agents do not work; and that two agents in the same occupation k work the same amount of time $l_{k,t}$ (if not, the two agents have different disutility of labor, and the household can improve total

⁴⁰Note that, for reasons that become apparent below, we are indexing B_k also with the agent index ι in (17). However, all agents in the same occupation have the same B_k , that is $B_{\iota,k} = B_{\iota',k}$ for agents ι and ι' working in the same occupation.

welfare by moving labor from one agent to the other) we can replace the integral on labor using the fractions of workers in each occupation and the unemployment pool. Total utility of the household becomes:

$$W = \left\{ \sum_{t=0}^{\infty} \beta^t \int_0^1 \frac{c_{\iota,t}^{1-\psi} - 1}{1-\psi} d\iota \right\} - \left\{ \sum_{t=0}^{\infty} \beta^t \left[\sum_{k=1}^J B_k l_{k,t}^{\mu} N_{k,t} + 0 * \left[1 - \left[\sum_{k=1}^J N_{k,t} \right] \right] \right] \right\},$$

subject to the budget constraint at each t ,

$$\int_0^1 c_{\iota,t} d\iota + I_t \leq \sum_{k=1}^J w_{k,t} (l_{k,t} - l_{k,t}^s) N_{k,t} + p_{r,t} r_t,$$

and the laws of motion for occupations size (4).

The Lagrangean for the problem is:

$$\begin{aligned} L = & \sum_{t=0}^{\infty} \beta^t \left\{ \int_0^1 \frac{c_{\iota,t}^{1-\psi} - 1}{1-\psi} d\iota - \left[\sum_{k=1}^J N_k l_k^{\mu} B_k \right] + \right. \\ & - \sum_{k=1}^J \nabla_{k,t} \left[N_{k,t+1} - (1-\kappa) N_k \left[(1-\gamma_k) + \Psi_k \left(1 - \sum_{i=1}^J N_i \right)^{\phi_k} (l_k^s)^{1-\phi_k} \right] \right] + \\ & \left. - \lambda_t \left[\int_0^1 c_{\iota,t} d\iota + r_{t+1} - r_t (1-\delta) - p_{r,t} r_t - \sum_{k=1}^J w_{k,t} N_{k,t} (l_{k,t} - l_{k,t}^s) \right] \right\}. \end{aligned}$$

The FOC with respect to $c_{\iota,t}$ is:

$$\frac{\delta L}{\delta c_{\iota,t}} = 0 \Rightarrow \beta^t c_{\iota,t}^{-\psi} = \lambda_t, \quad (20)$$

which implies that, in equilibrium, all agents consume the same amount at each t , so $c_{\iota,t} = c_{j,t} = c_t$. Thus, in the solution to the problem, it always holds that:

$$\int_0^1 \frac{c_{\iota,t}^{1-\psi} - 1}{1-\psi} d\iota = \frac{c_t^{1-\psi} - 1}{1-\psi}$$

and

$$\int_0^1 c_{\iota,t} d\iota = c_t.$$

It follows that we can write the problem of the household as:

$$\max_{c_t, l_{k,t}, l_{k,t}^s, r_{t+1}, N_{k,t+1}} W = \sum_{t=0}^{\infty} \beta^t \left\{ \frac{c_t^{1-\psi} - 1}{1-\psi} - \left[\sum_{k=1}^J B_k l_{k,t}^{\mu} N_{k,t} \right] \right\},$$

subject to the budget constraint:

$$c_t + r_{t+1} - r_t(1 - \delta) \leq \sum_{k=1}^J w_{k,t} N_{k,t} (l_{k,t} - l_{k,t}^s) + p_{r,t} r_t,$$

and the condition (4) for each k , as in the main text.