Looking ahead at the effects of automation in an economy with matching frictions

Guimarães, L., & Mazeda Gil, P. (2022). Looking ahead at the effects of automation in an economy with matching frictions. *Journal of Economic Dynamics and Control*, *144*, Article 104538. https://doi.org/10.1016/j.jedc.2022.104538

Published in:
Journal of Economic Dynamics and Control

Document Version:
Publisher's PDF, also known as Version of record

Queen's University Belfast - Research Portal:
Link to publication record in Queen's University Belfast Research Portal

Publisher rights
Copyright 2022 the authors.
This is an open access article published under a Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution and reproduction in any medium, provided the author and source are cited.

General rights
Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy
The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

Open Access
This research has been made openly available by Queen's academics and its Open Research team. We would love to hear how access to this research benefits you. – Share your feedback with us: http://go.qub.ac.uk/oa-feedback

Download date:18. May. 2026

Journal of Economic Dynamics & Control 144 (2022) 104538

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

Looking ahead at the effects of automation in an economy with matching frictions[☆]

Luís Guimarães[a,*], Pedro Mazeda Gil[b]

[a] Queen's Management School, Queen's University Belfast, UK and cef.up, Portugal
[b] PlanAPP, Faculdade de Economia da Universidade do Porto, and cef.up, Portugal

ARTICLE INFO

Article history:
Received 26 April 2022
Revised 18 September 2022
Accepted 21 September 2022
Available online 26 September 2022

JEL classification:
E24
J64
L11
O33

Keywords:
Automation
Employment
Labor-market frictions
Technology choice

ABSTRACT

We look at how advances in AI and Robotics will affect employment in an economy with matching frictions and endogenous job destruction. In the model, tasks can be produced by workers or by machines. Workers have a comparative advantage in producing advanced tasks but machines tend to catch up with labor, leading to automation. To calibrate the model, we rely on predictions in the literature about the expected share of automated jobs due to AI and Robotics. Our model suggests that these technological innovations will raise job destruction but also job creation because the prospect of automating jobs increases the value of hiring workers. Therefore, long-run employment might fall but not massively. Furthermore, employment will likely rise if consumers value human interactions (*human touch*) as the relative price of labor tasks increases with widespread usage of machines. Regarding policy, we compare the outcomes of a robot tax with alternative policies.

1. Introduction

The future of work is highly uncertain. On the one hand, it may look grim as advances in Artificial Intelligence (AI) and Robotics are expected to disrupt the labor market. This pessimism is entrenched in the *end-of-work* argument in Brynjolfsson and McAfee (2011), who describe a race between workers and machines that ultimately increases unemployment due to automation. Moreover, the predictions in Frey and Osborne (2017) appear to support the pessimism: they anticipate that advances in AI and Robotics (particularly Machine Learning and Mobile Robotics) will destroy almost half of US jobs within a decade or two. On the other hand, one may be positive on the future of work because AI and Robotics

---

[☆] We are deeply grateful for the comments and suggestions of the Editor, B. Ravikumar, the Associate Editor, and the two referees. We also thank the helpful comments by Marta Alvaro-Taus, Murat Çelik, Alper Çenesiz, Davide Debortoli, Sarah Guillou, Miguel León-Ledesma, Paul Levine, Diogo Lourenço, Ricardo Nunes, Ariell Reshef, and Holger Strulik, as well as participants at the Workshop on Economic Growth, Innovation, and Finance (SKEMA Business School, Nice), Workshop on Technological Change and the Future of Jobs and the Labor Share of Income (University of Porto), and conference participants at the University of Surrey. This research has been financed by the European Regional Development Fund through COMPETE 2020 Programa Operacional Competitividade e Internacionalizao (POCI) and by Portuguese public funds through FCT (Fundao para a Cncia e a Tecnologia) in the framework of the project POCI-01-0145-FEDER-006890 and (cef.up Center for Economics and Finance at University of Porto). Luís Guimarães also thanks FCT's financial support (2020.00714.CEECIND). The views expressed in the paper are those of the authors and do not necessarily represent those of PlanAPP.

[*] Corresponding author.
E-mail addresses: l.guimaraes@qub.ac.uk (L. Guimarães), pgil@fep.up.pt (P. Mazeda Gil).

https://doi.org/10.1016/j.jedc.2022.104538

should give rise to new jobs (Autor, 2015) and past automating technologies have presumably created more jobs than they destroyed, increasing employment.[1] These opposing views prompt our question: will advances in AI and Robotics massively reduce long-run employment?

To answer this question, we use a general-equilibrium model that borrows several features from models in the literature. The narrative and assumptions of our model broadly agree with those in Acemoglu and Restrepo (2018). Labor has a comparative advantage in producing new and complex tasks and, thus, new firms tend to invest in, what we call, the *manual* technology and produce using only labor. Machines, however, tend to catch up with labor in producing tasks. Every period, some workers lose their comparative advantage, motivating their employers to fire them and automate the production of the tasks. In this case, firms move to, what we call, the *automated* technology and produce using only machines/robots.[2] Yet, we fundamentally deviate from Acemoglu and Restrepo and build a model with matching frictions based on the Diamond-Mortensen-Pissarides setup. Thus, our setup allows us to realistically model the long-term firm-worker relationship and consider the role of automation in endogenizing job destruction.

We conceptualize advances in AI and Robotics as an *automation-augmenting* shock – i.e., an aggregate increase in the productivity of all machines/robots, which lowers labor comparative advantage and, thus, motivates firms to automate production. To calibrate this shock, we rely on the literature's predictions of the future share of automated jobs (Arntz et al., 2017; Frey and Osborne, 2017), and find that employment will likely fall but not massively.

In all our simulations, job creation and job destruction increase.[3]

Job destruction increases because the automation-augmenting shock makes it more profitable to replace workers with machines, while job creation increases because of two mechanisms. First, an automation-augmenting shock increases job creation if firms can choose technology at the time of entry. In this scenario, motivated by the increase in productivity of the automated technology, firm entry surges in general-equilibrium, which ultimately rises job creation. Second, we present a mechanism (to our knowledge) new to the literature through which automation-augmenting shocks promote job creation. Because firms are forward-looking and workers are usually needed to produce new and complex tasks, firms have a higher incentive to hire a worker upon entry in anticipation of the greater profits when they automate production post-entry. A real-world example confirming the existence of this mechanism is UBER.[4]

Even though both flows increase if machines become more productive, their magnitudes crucially depend on the calibration of the model, rendering a small but ambiguous employment response. One important factor is the degree of technology constraints when firms enter. If more firms are technologically unconstrained, then employment falls less and may even rise as both mechanisms promoting job creation (described above) are active. More surprisingly, the future share of automated jobs does not monotonically dampen employment; indeed, there is a U-relationship as the apparently most pessimistic scenarios coincide with more employment. This occurs because a higher calibrated share of automated jobs requires a large automation-augmenting shock, generating a large scale effect. The increased likelihood of automating a job together with the increased productivity of the automated technology reinforce each other to raise the expected value of hiring a worker; this amplifies job creation, implying more employment when a sufficiently large share of future jobs are automated.

In our model, workers performing tasks with less comparative advantage earn lower wages (agreeing with, e.g., Arnoud, 2018), but, overall, wages increase together with the rise in the share of automated jobs (concurring with, e.g., Autor and Salomons, 2018 and Graetz and Michaels, 2018). In light of this and to better understand the employment response, we consider a new set of experiments in which we counterfactually assume that wages are orthogonal to labor market tightness (and to the productivity of the automated technology). In this case, job creation is seriously magnified to the point that employment increases for a much wider range of calibrations. Employment does, however, still fall in some calibrations because matching frictions imply more congestion among firms (and, thus, higher hiring costs) when job-finding rates rise, discouraging further job creation.

In another variant of the model, we consider the implications of, what we call, *human touch*. Even though both workers and machines can execute the same task, consumers may deem tasks executed by humans and by machines differently due to the distinctive contribution of human interactions. In this scenario, a widespread use of machines increases the relative price of manual tasks, which largely increases job creation for a given magnitude of job destruction. Thus, if many of the tasks produced in the economy are directed to consumers and they find the differentiated *human touch* relevant, then an automation-augmenting shock is likely to increase employment.

Because of the automation threat, there is a growing literature assessing whether economic policy improves welfare and employment (e.g., Prettner and Strulik, 2020; Guerreiro et al., 2021; Jaimovich et al., 2021; Gasteiger and Prettner, 2022). We

---

[1] A growing empirical literature suggests that recent technologies with the purpose of automating production have increased employment by creating more jobs than they destroy (see, e.g., Bessen, 2016, Autor and Salomons, 2018, and Gregory et al., 2021; see also Bessen et al., 2020 for a review of this literature). A notable exception is Acemoglu and Restrepo (2020), who find that robot adoption depresses employment and wages at the commuting-zone level.

[2] Our setup thus assumes the extreme case of a technology that only uses labor and a technology that only uses capital/robots. We share this convenient assumption with, e.g., Zeira (1998, Sec. 7; 2010), Acemoglu and Restrepo (2018), Alesina et al. (2018), and Guimarães and Gil (2022).

[3] This result concurs with the emergence of the so-called *gig economy* (less permanent contracts and more hiring on demand), which Bloom et al. (2018) relate with automation.

[4] UBER's Initial Public Offering prospectus offers a good example of this channel. The prospectus assumes that developing autonomous vehicles importantly contributes to the current valuation of the firm by potentially allowing it to reduce their labor demand in the future. Thus, the possibility of automating tasks in the future contributes to UBER's investment and recruitment in the present.

add to this literature by studying an automation tax – a tax that is paid for buying a robot to replace a worker in producing a task. Our model suggests that it is less pernicious on employment and wages than a (standard) robot tax because it inhibits a larger fall in entry. Yet, the automation tax is likely difficult to implement in practice and, thus, we compare these taxes with a more standard firing tax and hiring subsidy.

Contrasting our paper with the literature, we find that (best of our knowledge) our model is the only framework in which automation has an ambiguous long-run effect on employment. In the literature, either employment is assumed inelastic (e.g., Acemoglu and Restrepo, 2018), or employment always falls (e.g., Zeira, 1998; Prettner and Strulik, 2020) or it always increases (Guimarães and Gil, 2022). Furthermore, our paper relates to Berg et al. (2018), Nakamura and Zeira (2018), Caselli and Manning (2019), and Basso and Jimeno (2021) (and again with Acemoglu and Restrepo, 2018 and Prettner and Strulik, 2020) in that these papers also assess how automation-related shocks may affect either wages or employment in the future. But a key difference relative to our paper is that these papers assume perfectly competitive labor markets. In a different vein, our paper relates to Leduc and Liu (2019), Cords and Prettner (2021), and (again) Guimarães and Gil, 2022 as they include models with matching frictions and automation. Yet, there are important differences regarding the objects of study and models used. First, these papers do not aim to assess the long run implications of expected advances in AI and Robotics on the employment rate.[5] Second, all these papers assume constant job-destruction rates; in contrast, allowing for endogenous changes in the job destruction rate in our model is key for the ambiguity of the effects of AI and Robotics on employment. In this regard, our paper is closer to Mortensen and Pissarides (1998) and Hornstein et al. (2007) who endogenize the effects of technology – but not automation – on job destruction.

The remainder of this paper is organized as follows. We start by detailing our model in Section 2. Then, Section 3 describes our calibration and main results. Section 4 considers the role of the *human touch*, and Section 5 discusses the role of policy. Section 6 concludes our paper.

2. The model

2.1. Entry and environment description

In the model, the aggregate output is the sum of the production of a number of tasks, which can be produced by one of two technologies: an automated technology and a manual technology. At the time of entry, a firm must first create a task, which amounts to an entry cost denoted by Ω. If the firm produces the task using the automated technology, it must pay an additional κ_K, which can be interpreted as a robot investment. If the firm produces the task using the manual technology, it must pay an additional $\kappa_L(\theta)$ to match with a worker and it must bargain wages with the worker.

Entering firms that choose the manual technology must search for workers in the labor market. A Cobb-Douglas matching function determines the number of matches between these firms and the workers that were nonemployed at the beginning of the period.[6] This matching function has constant returns to scale, has as argument labor market tightness, θ, is scaled by matching efficiency, $\chi > 0$, and has an elasticity with respect to nonemployed workers of $0 < \eta < 1$. Thus, we write the hiring cost and the job-finding probability as, respectively, $\kappa_L(\theta) \equiv \frac{\bar{\kappa}_L}{\chi\theta^{-\eta}}$ and $f(\theta) \equiv \chi\theta^{1-\eta}$, where $\bar{\kappa}_L$ is a scalar.

Acemoglu and Restrepo (2018) assume that workers have a comparative advantage in producing more productive (higher-indexed) tasks. We borrow this assumption and assume that the manual technology produces $z_L z$ units of the task, while (as a normalization) the automated technology produces z_K units of the task. Thus, z represents the comparative advantage of workers in producing the respective task, so that highly-productive tasks (high z) tend to be produced by the manual technology and less-productive tasks with the automated technology. The relative productivity of the technologies is, instead, mostly influenced by z_L and z_K.

Figure 1 summarizes the timeline of how z shapes the distribution of firms between the technologies. In Acemoglu and Restrepo (2018), labor always has the highest comparative advantage in producing new tasks. We, however, assume a more general environment. Of the number of new tasks created each period, a proportion $1-\lambda_e$ start with the maximum z, $\bar{z}$, and, thus, workers have the maximum comparative advantage. In this case and in equilibrium, firms choose the manual technology and produce $z_L\bar{z}$ units of the task. Conversely, a proportion λ_e of new tasks draw z at the time of entry from a probability distribution function $G(z)$ in the interval $[z_{\min}, \bar{z}]$, and firms choose technology according to the respective present-discounted values. Thus, λ_e is a measure of technology constraints at the time of entry. Producing tasks with higher z is more profitable when the firm uses the manual technology to take advantage of the higher workers' comparative advantage. As a result, there is an idiosyncratic productivity cutoff, denoted by z_e^*, above which firms prefer the manual rather than the automated technology.[7]

Firms that start production using the manual technology can move to the automated technology in later periods. Their technological choice depends on how the task's idiosyncratic productivity, z, evolves over time, and if it becomes too low,

---

[5] For example, in our previous work in Guimarães and Gil (2022), we do not try to understand the future of employment and instead focus on the past evolution of the US labor share.

[6] The workers that lose their jobs (either exogenously or endogenously) do not produce for at least a period. This agrees with the evidence in Hall and Kudlyak (2019).

[7] If $\lambda_e > 0$, entry in the model is, at least, partially undirected, which is our assumption in Guimarães and Gil (2022).

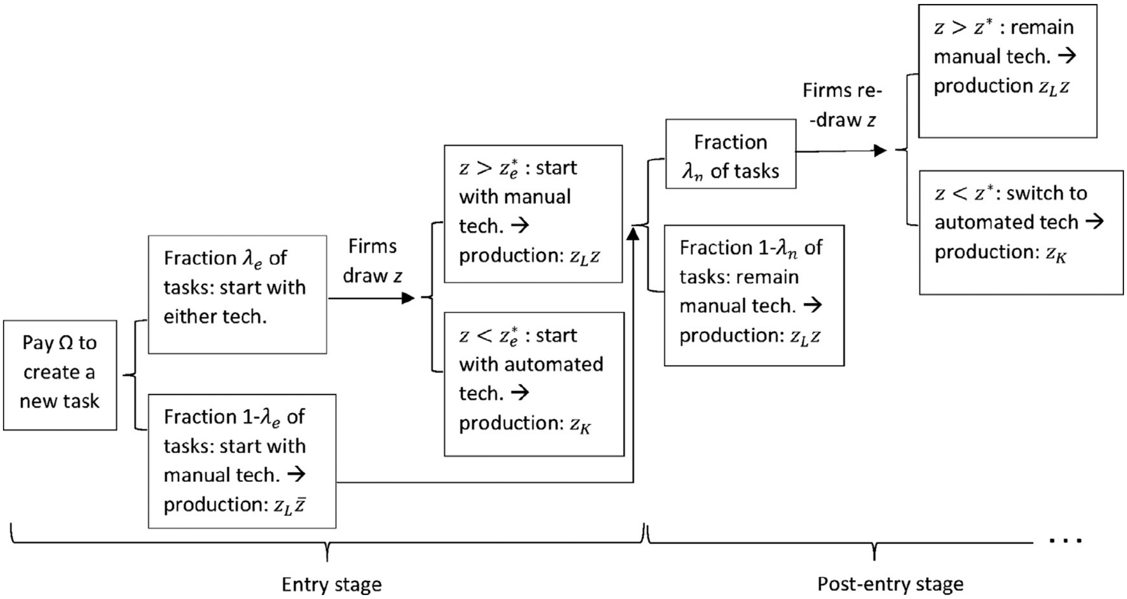

Fig. 1. Timing of technological constrains and technology choice.

manual firms destroy the job and automate production.[8] To model the evolution of z, we build on Mortensen and Pissarides (1994). After production takes place, a proportion $1-\lambda_n$ of manual firms sees no change in their tasks' idiosyncratic productivity. But a proportion λ_n of manual firms redraws the task's idiosyncratic productivity from the same distribution $G(z)$ of productivity levels. If the new idiosyncratic productivity, z, is too low – below the cutoff, which we denote by z^* – the manual firm fires the worker and shifts from the manual to the automated technology.[9]

These assumptions imply that shocks to the economy can change the employment rate by affecting both job creation and job destruction. Thus, this setting allows for a rich environment to study how automation-augmenting (rise in z_K) shocks affect the employment rate.

In writing the equations below, we omit the time subscripts as we are only interested in steady-states. Yet, within a period, there is an order of events that we must further clarify before laying out the equations. 1) New firms pay Ω to create a task and enter the market until a free-entry condition is satisfied. 2) A proportion λ_e of new firms and a proportion λ_n of manual firms (re)draw the task's idiosyncratic productivity, z. 3) Depending on the productivity draw, z, and anticipating wage bargaining, firms decide which technology to use in the following period. If an incumbent manual firm decides to automate the production of the task, it must fire the worker, pay κ_K, and wait a period to resume production. 4) Matching between new manual firms and workers occurs. 5) Production takes place and manual firms bargain wages with their workers. 6) A proportion δ_L of the tasks produced by active (producing within the period) manual firms and a proportion δ_K of the tasks produced by active automated firms are exogenously destroyed.

2.2. Firms

An active firm using the manual technology to produce a task with idiosyncratic productivity z has the following present-discounted value $J_L(z)$:

$$J_L(z) = z_L z - w(z) + \beta(1-\delta_L)\left\{(1-\lambda_n)J_L(z) + \lambda_n\left[G(z^*)(\beta J_K - \kappa_K) + \int_{z^*}^{\bar{z}} J_L(z)\mathrm{d}G(z)\right]\right\}, \tag{1}$$

where we assume a discount factor of β. This firm produces $z_L z$ units of the task (and, thus, of the output) and pays the wage $w(z)$ to its worker. There is a probability $1-\delta_L$ that it will keep producing in the following period. And if it does produce, its value remains unchanged with a probability $1-\lambda_n$ and changes due to the redraw of the idiosyncratic productivity, z, with a probability λ_n. Those that draw a productivity below z^* prefer to fire the worker and change to the

---

[8] This echoes the settings in Hornstein et al. (2007) and Acemoglu and Restrepo (2018) as, in their models, the tasks (relative) productivity gradually falls due to the expansion of the technological frontier over time and the implied progressive obsolescence of existing tasks/firms. Slightly different from them, we assume that the progressive obsolescence of manual tasks is stochastic, in part, for convenience (we do not need to keep track of z for each task) but, more importantly, tasks may (in reality) differ on the speed at which they are automated, resembling a random process.

[9] Naturally, some firms also draw a higher z. We can interpret this as a form of technological catching up of the task. In any case, the most relevant aspect for the mechanism of the model is that these firms remain manual.

automated technology; in this case, because they already paid Ω and it takes one period to shift technologies, their value equals the discounted value of the automated technology, βJ_K, reduced of the technology-specific cost κ_K. If they draw a productivity above z^*, they choose to maintain the manual technology; in this case, their value equals the unconditional expected value of the manual technology between z^* and $\bar{z}$. This intuitively implies that z^* is determined by the following indifference condition:

$$J_L(z^*) = \beta J_K - \kappa_K. \tag{2}$$

The present-discounted value of the automated technology, J_K, is much simpler as its productivity is constant:

$$J_K = z_K + \beta(1-\delta_K)J_K. \tag{3}$$

At the time of entry, all firms pay Ω to create a new task. A proportion λ_e of the new firms draws the task's idiosyncratic productivity; the other firms start with the manual technology with idiosyncratic productivity $\bar{z}$. Among the firms that draw idiosyncratic productivity, a proportion $G(z_e^*)$ chooses the automated technology and the remaining firms choose the manual technology. These assumptions allow us to write the free-entry condition in our model:

$$\lambda_e\left[G(z_e^*)(\beta J_K - \kappa_K) + \int_{z_e^*}^{\bar{z}} (\beta J_L(z) - \kappa_L(\theta))\mathrm{d}G(z)\right] + (1-\lambda_e)(\beta J_L(\bar{z}) - \kappa_L(\theta)) = \Omega, \tag{4}$$

where the present-discounted values, J_K and $J_L(z)$, are discounted by β because it takes one period for firms to start production. New firms that draw productivity are only indifferent between either technology if their values net of the technology-specific entry cost are equal. This occurs when the task's idiosyncratic productivity equals z_e^*:

$$\beta J_L(z_e^*) - \kappa_L(\theta) = \beta J_K - \kappa_K. \tag{5}$$

2.3. Workers

In our model, there is a unit measure of risk-neutral workers who are either employed or nonemployed. The lifetime income of an employed worker is given by $E(z)$:

$$E(z) = w(z) + \beta\left\{(1-\delta_L)\left[(1-\lambda_n)E(z) + \lambda_n\left(G(z^*)U + \int_{z^*}^{\bar{z}} E(z)\mathrm{d}G(z)\right)\right] + \delta_L U\right\}. \tag{6}$$

$E(z)$ increases with the wage $w(z)$, which varies with the idiosyncratic productivity of the task the worker is producing at the firm. $E(z)$ falls with the probability that the job is exogenously destroyed and the worker is back to nonemployment. In this case, the lifetime income is given by U. $E(z)$ also changes with the future productivity draw of the firm: if the new productivity draw is low – below z^* –, the firm fires the worker and the lifetime income returns to U; if the new productivity draw exceeds z^*, then the wage changes, shifting the lifetime income of employment.

If nonemployed, a worker enjoys income $b \geq 0$ and finds a job with a probability $f(\theta)$. In equilibrium, nonemployed workers only match with new firms to produce new tasks. But new tasks vary in idiosyncratic productivity. A proportion $1-\lambda_e$ of new tasks start with idiosyncratic productivity $\bar{z}$ and, thus, are produced by labor. On the other hand, a proportion λ_e of new tasks have their idiosyncratic productivity drawn from $G(z)$ and the firms producing the tasks only hire a worker if the draw exceeds z_e^*. As a result, we write the lifetime income of a nonemployed worker as

$$U = b + \beta\left\{f(\theta)\left[(1-\lambda_e)E(\bar{z}) + \frac{\lambda_e}{1-G(z_e^*)}\int_{z_e^*}^{\bar{z}} E(z)\mathrm{d}G(z)\right] + (1-f(\theta))U\right\}. \tag{7}$$

2.4. Wage bargaining

Workers and firms bargain over wages such that the bargained wage maximizes the Nash product:

$$w(z) = \arg\max\,[E(z)-U]^{\phi}[J_L(z) - \max\,(\beta J_L(z) - \kappa_L(\theta), \beta J_K - \kappa_K)]^{1-\phi}, \tag{8}$$

where the parameter $0 < \phi < 1$ measures the worker's bargaining power. A firm that employs a worker has two outside options. It may fire the worker and look for a new one, which generates a value of $\beta J_L(z) - \kappa_L(\theta)$.[10] Alternatively, it may fire the worker and adopt the automated technology, which generates a value of $\beta J_K - \kappa_K$. We infer that there is an idiosyncratic productivity cutoff that makes the manual firm indifferent between the two outside options, which turns out to be the same as the entry cutoff, z_e^*, in Eq. (5).[11] Thus, we summarize the solution to Nash bargaining as

$$E(z) - U = \frac{\phi}{1-\phi}[J_L(z) - (\beta J_L(z) - \kappa_L(\theta))] \text{ if } \bar{z} > z \geq z_e^*; \tag{9}$$

---

[10] Since the productivity z is idiosyncratic, it implies that if firms decide to look for another worker, they do not have to redraw productivity. This prevents workers from capturing a large share of the surplus generated by greater productivity.

[11] If firms draw $z \in [z^*, z_e^*]$ post-entry, they do not break an existing match because the hiring cost is sunk and, thus, the value of the match is positive. Therefore, in this range, the outside option in wage bargaining shifts from replacing the worker with another worker to replacing the worker with a machine. If firms draw $z \in [z_{\min}, z^*]$, the value of the match is negative and firms automate production; yet, the same logic applies as firms automate production knowing that the wage would be set according to the outside option of automating production.

$$E(z) - U = \frac{\phi}{1-\phi}[J_L(z) - (\beta J_K - \kappa_K)] \text{ if } z_{min} < z < z_e^*. \tag{10}$$

In both cases, workers retain a proportion ϕ of the surplus, which is an increasing function of the idiosyncratic productivity, z, only due to $J_L(z)$. As a result, wages increase with z but less than proportionately. Equation (9), for example, implies that wages increase in proportion $\frac{\phi(1-\beta)}{\phi(1-\beta)+1-\phi} < 1$ of $z_L z$. This confirms our anticipation that greater idiosyncratic productivity implies greater profits, guaranteeing that only the least productive firms in using the manual technology prefer to use the automated technology.

2.5. Equilibrium

The equilibrium of the model is defined at the aggregate level of the economy and is characterized by the vector $(\theta, z^*, z_e^*, w(z))$, which satisfies the free-entry condition, Eq. (4), and the two indifference conditions, Eqs. (2) and (5), and solves Nash bargaining.

2.5.1. Employment rate

We define employment as the number of workers employed at the time of production. In equilibrium, employment is determined by the balance between the flows from employment to nonemployment and the flows from nonemployment to employment. Using n to denote the employment rate, the flows from nonemployment to employment sum up to $f(\theta)(1-n)$: a proportion $f(\theta)$ of the nonemployed workers, $(1-n)$, find jobs every period. The flows from employment to nonemployment take two forms because workers may lose their jobs exogenously and endogenously. There is a probability δ_L that employed workers lose their jobs for exogenous reasons. From those that do not lose their jobs for exogenous reasons, there is a probability λ_n that the productivity of the task changes. And there is a probability $G(z^*)$ that the new productivity is below the cutoff z^*, leading the firm to move to the automated technology and fire the worker. Thus, after some algebra, we get an equilibrium employment rate of

$$n = \frac{f(\theta)}{f(\theta) + \delta_L + (1-\delta_L)\lambda_n G(z^*)}. \tag{11}$$

2.5.2. Number of firms

The number of manual firms is n because every manual firm employs one worker. The number of automated firms, n_K, consists of (i) the number of existing automated firms that are not exogenously destroyed, $(1-\delta_K)n_K$; (ii) the number of existing manual firms that survive and choose to automate, $n(1-\delta_L)\lambda_n G(z^*)$; and (iii) new firms that choose to automate upon entry. To express the latter, note that only new firms hire workers; hence, $f(\theta)(1-n)$ is the number of new manual firms. We also know that new manual firms are a fraction $1-\lambda_e G(z^*)$ of the total number of new firms. Therefore, we can conclude that the total number of new firms is $\frac{f(\theta)(1-n)}{1-\lambda_e G(z^*)}$, of which a fraction $\lambda_e G(z^*)$ choose to automate upon entry. In equilibrium, the number of automated firms is

$$n_K = \frac{(1-\delta_L)\lambda_n G(z^*)}{\delta_K} n + \frac{\frac{\lambda_e G(z_e^*)}{1-\lambda_e G(z_e^*)}}{\delta_K} f(\theta)(1-n). \tag{12}$$

3. Results

3.1. Calibration

We calibrate the model to monthly US data and summarize our benchmark calibration in Table 1. We set $\beta = 0.996$, which implies an annual discount rate of 4.91%. We follow Petrongolo and Pissarides (2001) and set $\eta = 0.5$. We also set $\phi = 0.5$. In our model, firms draw z – which captures the comparative advantage of workers – from a uniform distribution, i.e., $G(z) = \frac{z - z_{\text{min}}}{\bar{z} - z_{\text{min}}}$, with the parameters calibrated internally as explained below.[12] To calibrate b, we assume it is 70% of the productivity of the firm that draws $z = z_{\text{min}} + \frac{\bar{z} - z_{\text{min}}}{2}$. This is similar to what we find in many studies in the literature (including Hall and Milgrom, 2008; Pissarides, 2009; and Coles and Kelishomi, 2018) that assume that $b \approx 0.7 z_L$ in models with homogeneous firms.

In choosing our benchmark calibration, we follow Acemoglu and Restrepo (2018) and assume that all new tasks are technologically constrained and, thus, are initially produced using the manual technology, implying that $\lambda_e = 0$. We also choose $\delta_K = 0.0087$, which implies a 10% robot depreciation rate as in Graetz and Michaels (2018). Furthermore, we normalize z_{min} to unity.

---

[12] We also check the robustness of our results by replacing the uniform distribution with a Pareto distribution; in those exercises, we also calibrate the parameters of the Pareto internally (see Appendix B).

Table 1
Benchmark Calibration.

Parameters	
Discount factor:	$\beta = 0.996$
Matching function elasticity:	$\eta = 0.5$
Workers' bargaining power:	$\phi = 0.5$
Minimum productivity draw:	$z_{\min} = 1$
Nonemployment income:	$b = 0.7 z_L (z_{\min} + (\bar{z} + z_{\min})/2)$
Rate of automated-firm destruction:	$\delta_K = 0.0087$
Share of firms technologically unconstrained:	$\lambda_e = 0$
Current steady-state targets	
Employment rate:	$n = 0.78$
Labor share:	$LS = 0.61$
Labor market tightness:	$\theta = 0.72$
Share of automated jobs within a decade:	$SAJ = 0.045$
Share of nonroutine jobs:	$(n - n_e^* - n^*)/n = 0.58$
Job destruction rate:	$JD = 0.016$
Output per worker (numeraire):	$y/n = 1$
Cost of capital/robot:	$\kappa_K = 12w(z^*)$
Hiring costs:	$\kappa_L(\theta) = 20/160 z_L \bar{z}$
Future steady-state target	
Share of automated jobs within a decade:	$SAJ_{future} = 0.09$

The remaining nine parameters, Ω, z_K, z_L, $\bar{z}$, κ_K, $\bar{\kappa}_L$, χ, δ_L, and λ_n, are internally calibrated such that our steady-state matches nine targets. We target the prime-age (aged 25–54) workers' employment rate and the labor share in the US from 1977 until 2018;[13] this implies that $n = 0.78$ and $LS = 0.61$. Following Pissarides (2009), we target labor market tightness in the US so that $\theta = 0.72$. We also target the share of jobs that are automated. In particular, we target the rate at which jobs are automated within a decade ("share of automated jobs"), $SAJ \equiv 10 \times 12(1 - \delta_L)\lambda_n G(z^*)$, such that it equals the proportion of routine jobs that disappeared (a proxy for automated jobs) between 2002 and 2017 in the US. Based on Jaimovich and Siu (2020), we target $SAJ = 0.045$. We further rely on Jaimovich and Siu, 2020 to set a target for $\frac{n - n_e^* - n^*}{n}$, i.e., the share of jobs in which workers have the maximum comparative advantage;[14] we proxy this share using the share of nonroutine jobs in the US economy (58%). We also impose that the steady-state probability that a firm-worker match breaks equals the average occupational mobility in the US. For this, we rely on Kambourov and Manovskii (2008), who document that it is about 18% per year at the three-digit level in PSID data; thus, we set the job destruction rate in the model as $JD \equiv \delta_L + SAJ/120 = 0.016$.[15]

To calibrate the technology-specific entry costs, κ_K and $\bar{\kappa}_L$, we turn to Barron and Bishop (1985) and Fornino and Manera (2022). The survey evidence in Barron and Bishop, 1985 suggests that it takes approximately 20 recruiting hours to hire one worker; assuming that the recruiters' productivity is $z_L \bar{z}$ and that an employed worker works an average of 160 hours per month, we impose $\kappa_L(\theta) = 20/160 z_L \bar{z}$.[16] Fornino and Manera (2022) estimate that the price of a robot is close to the annual salary of production-line workers in manufacturing; hence, we target $\kappa_K = 12w(z^*)$. Finally, we use output per worker as the numeraire, $y/n = 1$.

3.2. Our approach

Almost all of the recent empirical studies on the effects of automating technologies point to a net increase in employment in the last four decades (see, e.g., Bessen, 2016; Autor and Salomons, 2018; Bessen et al., 2020; and Gregory et al., 2021). These studies suggest that the direct labor-displacing (job destruction) effect has been outweighed by indirect effects that ultimately lead to job creation. But do these results hold under all circumstances? In other words, can we expect a different future?

In answering this question, we use our model to assess the future implications of automating technologies like Artificial Intelligence (AI) and Robotics. We conceptualize advances in these technologies as an automation-augmenting shock, i.e., a rise in z_K. This shock increases the productivity of the automated technology, which lowers the comparative advantage of

---

[13] We target the employment rate of prime-age workers because our model abstracts from demographic changes. For the derivation of the labor share in our model, see Appendix A.

[14] n_e^* denotes the number of manual firms that draw $z \in [z_e^*, \bar{z}]$, while n^* denotes those that draw $z \in [z^*, z_e^*)$. See Appendix A for derivations.

[15] In business cycle applications of matching models, the average job destruction rate is typically 3.6% based on Shimer (2012). We, however, prefer to target average occupational mobility because we focus on the effect of automation on long run employment and, thus, our model abstracts from much of the churn in the labor market, concentrating instead on the motivation to create jobs composed of new tasks (more broadly, occupations) and their automation. In a similar vein, the literature addressing the effects of labor market turbulence on the European unemployment rate (e.g., Ljungqvist and Sargent, 2004) also calibrate their matching models with low churn as their goal is to study long-run unemployment. Despite this argument, our robustness checks includes the case of $JD = 0.03$ (the proportion of prime-age workers that move out of employment in CPS data from 1978 to 2012), and we find that our message does not change (see Table B1 in Appendix B).

[16] Barron and Bishop's (1985) evidence is commonly used in the literature (e.g., Fujita and Ramey, 2012).

Table 2
Main results.

	$\%\Delta n$	$\%\Delta f(\theta)$	$\%\Delta JD$	$\%\Delta LS$	$\%\Delta w(\bar{z})$	$\%\Delta z_K$
Benchmark	-0.28	1.01	2.29	-1.27	0.05	3.87
$\lambda_e = 0.2$	0.25	3.46	2.29	-1.90	0.14	8.05
Ad hoc wage ($\eta = 0.5$)	0.69	5.55	2.29	-1.24	0.00	1.56
Ad hoc wage ($\eta = 0.3$)	2.21	13.44	2.29	-1.24	0.00	1.56

*Note:*This table shows the effects of an automation-augmenting shock calibrated such that $SAJ_{future} = 0.09$. The six columns show, respectively, the percentage change in employment, job-finding probability, job-destruction probability, labor share, wages (in firms with draw $z = \bar{z}$), and the productivity of the automated technology, z_K. The first line uses the benchmark calibration in Table 1. The second line changes the benchmark calibration by assuming that 20% of the firms are technologically unconstrained when they enter. The third and fourth lines show the results of the variant of our model with an ad hoc wage instead of Nash bargaining described in Section 3.5.

labor and, *ceteris paribus*, increases the cutoff z^* and the share of automated jobs (see Eq. (2)). By raising z^*, advances in AI and Robotics effectively increase the job destruction rate, which, in our view, is a key feature of these new automating technologies. For example, in reality, Machine Learning – AI's state-of-the-art – allows machines to adapt to new tasks as long as enough data is available to identify patterns in it (Naudé, 2020). Therefore, a rise in z_K (an improvement in Machine Learning) should permanently facilitate machines catching up with labor in producing each task as data is collected, increasing the pace of automation.[17]

Our approach is simple: first, we recalibrate z_K such that our model targets the expected share of automated jobs in the future, $SAJ_{future} > SAJ$; then, we look at the consequences of the implied rise in z_K on the steady-state of employment and other variables of interest. To set SAJ_{future}, we rely on recent studies predicting the share of jobs that will be automated in the next decades due to rapid improvements in AI and Robotics. The most prominent example is Frey and Osborne (2017). After estimating occupation-specific probabilities of automation, they conclude that about 47% of US jobs have a probability of automation in excess of 70% within a decade or two,[18] putting them at a high risk of automation. On the other hand, Arntz et al. (2017) use the estimates in Frey and Osborne (2017) but conclude that only 9% of US jobs are at a high risk of automation also within a decade or two. The main distinction between the two papers lies on the approach: Arntz et al. (2017) look at the various tasks composing each job while Frey and Osborne (2017) look only at the main task; as most jobs include tasks that are hard to automate (*e.g.*, face-to-face interactions with customers), the empirical results in Frey and Osborne (2017) tend to overestimate the actual risk of automation. Therefore, in setting a benchmark target of SAJ_{future}, and despite our simplification in the model that each job corresponds to only one task, we rely more heavily on the estimates in Arntz et al. (2017). In particular, we set $SAJ_{future} = 0.09$ as our benchmark (their upper bound), implying that the percentage of jobs automated every decade increases from 4.5% to 9% between our two steady-states.

3.3. *The future of employment*

The first line in Table 2 summarizes our main findings: the calibrated increase in z_K slightly reduces employment, increases both job-finding and job-destruction probabilities, slightly increases the wage in jobs with maximum comparative advantage ($z = \bar{z}$) (all new jobs when $\lambda_e = 0$), and lowers the labor share.[19]

Automation-augmenting shocks affect employment through changes in both job creation and job destruction. The change in job destruction is imposed by our SAJ_{future} target, but the change in job creation depends on the calibration of the model. Thus, we focus on job creation to understand the change in employment. Because $\lambda_e = 0$ in our benchmark calibration, all firms with new tasks must invest in the manual technology and can only take advantage of the increased productivity of the automated technology inasmuch as they later automate production. Therefore, hiring a worker is a mandatory first step, which promotes job creation when z_K rises. Formally, this effect can be seen looking at Eqs. (1)–(4). For given wages, the rise in z_K increases the continuation value of manual firms in Eq. (1), which is reinforced by the rise in the cutoff z^* as the least productive jobs are endogenously replaced by the automated technology (see Eq. (2)). This increased value of manual firms then implies, by the free-entry condition, Eq. (4), that the labor market tightens even though wages rise as firms share

---

[17] In the model of Acemoglu and Restrepo (2018), a rise in automation directly implies that, absent other shocks, workers permanently perform the same tasks for a shorter period. In this regard, the main difference between our model and theirs is that they assume a competitive labor market and, thus, the fall in the average duration of jobs is not relevant per se.

[18] The methodology used in Frey and Osborne (2017) does not allow them to specify a time-interval for jobs to actually be automated, but the authors judge it to be within a decade or two.

[19] We also find that the share of workers with maximum comparative advantage increases, which given our interpretation of these jobs as nonroutine, suggests that the documented increase in the share of nonroutine jobs in recent decades should persist. Furthermore, despite the increase in higher average wages, workers performing tasks with less comparative advantage continue to earn relatively lower wages in our model, which agrees with the findings in Arnoud (2018).

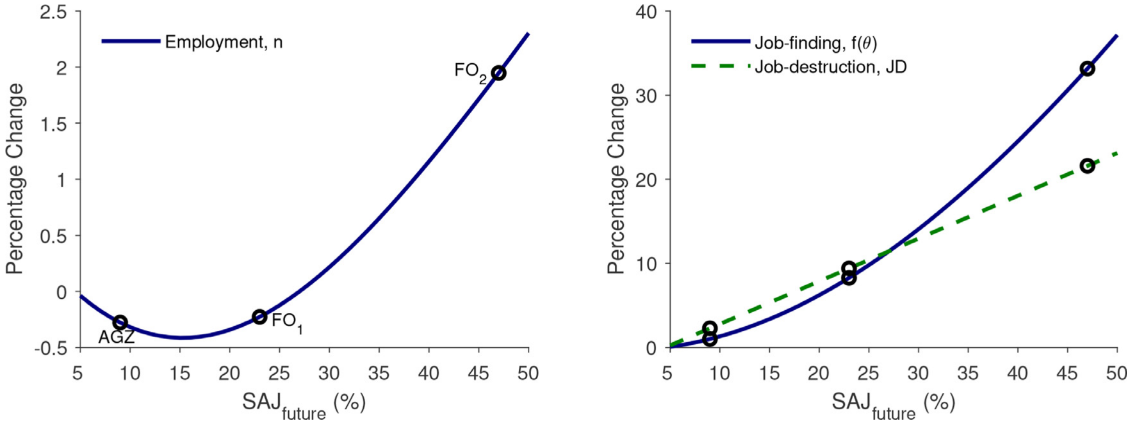

Fig. 2. Robustness checks to different SAJ_{future} *Note:* This figure shows the effects of an automation-augmenting shock calibrated such that our model reaches the different targets of SAJ_{future} in the horizontal axis (in percentage terms). The left panel shows the percentage change in employment. The right panel shows the percentage change in the job-finding probability (solid line) and the job-destruction probability (dashed line). Each circle identifies different estimates of SAJ_{future} found in the literature. The circles on the left (marked with *AGZ*) correspond to our benchmark, $SAJ_{future} = 0.09$, based on Arntz et al. (2017). The other two (marked with FO_1 and FO_2) correspond to the lower and upper bounds of the estimates in Frey and Osborne (2017), respectively.

the additional surplus with workers. Yet, employment falls in our benchmark because the change in job destruction exceeds that of job creation.

Among the robustness checks that we performed, the only striking change occurs when we deviate from Acemoglu and Restrepo (2018) and assume that some new tasks are not technologically constrained and may start as automated ($\lambda_e = 0.2$; second line in Table 2).[20] In this case, employment rises because of a significant surge in job creation. Furthermore, wages rise substantially more than in the benchmark. These results are partially explained by the larger necessary increase in z_K to reach our target for SAJ_{future}; yet, even if z_K was to increase 3.87% as in the benchmark, employment would still rise 0.06%. This is due to a general-equilibrium effect. As expected, a rise in z_K increases the value of the automated technology, leading to a reallocation effect: some entering firms steer away from the manual technology and invest instead in the automated technology (z_e^* increases); for a given number of entering firms, job creation shrinks. Yet, firm entry does not remain fixed. An automation-augmenting shock also increases the expected value of a firm, which incentivizes firm entry.[21] The free-entry condition, Eq. (4), is only satisfied if the value of the manual technology drops, which occurs in our model through higher wages and, most importantly, greater labor market tightness and job creation. This general-equilibrium effect promoting job creation adds to the incentives already identified above. Thus, $f(\theta)$ rises more, and employment increases.

Irrespective of the calibration, the message from our model is that even if the share of automated jobs doubles in the future, the change in employment is, at most, mildly negative. In part this happens because automation is only a fraction of job destruction; indeed absent the rise in $f(\theta)$, employment would fall 0.5% in our benchmark.[22] But an important part is the contribution of job creation that attenuates and can even revert the response of employment.

3.4. An apparently pessimistic future

In Section 3.2, we have discussed our choice of SAJ_{future}, the share of automated jobs in the future. We have chosen the most pessimistic estimate in Arntz et al. (2017) but we have excluded the even more pessimistic estimates in Frey and Osborne (2017). Therefore, in this section, we assess how our results change with SAJ_{future}.

Under our benchmark calibration, Fig. 2 reports how different values of SAJ_{future} affect employment on the left panel and job-finding and job-destruction probabilities on the right panel. This figure shows that employment falls for a large range of values of SAJ_{future} because the change in the job-destruction probability exceeds that of the job-finding probability. Yet, it also shows that the increase in the share of automated jobs might increase employment, showing a change of the sign of the employment elasticity with respect to the productivity of the automated technology, z_K.

---

[20] See Table B1 in Appendix B for several robustness checks. This includes a different matching elasticity, a different bargaining power, and replacing the uniform distribution with a Pareto distribution. It also includes different targets for the share of non-routine jobs, hiring costs, robot installation costs, job-destruction rates, and the initial share of automated jobs.

[21] The expected value of a firm (prior to entry) surges because a higher z_K directly increases the expected value of automated firms and, *ceteris paribus*, indirectly increases the expected value of manual firms. Furthermore, the expected value of a firm increases even further because the productivity of the tasks produced with manual technology is heterogeneous and the firms drawing the least productive of these tasks prefer the automated technology when z_K increases (z_e^* increases).

[22] This is not as substantial as a doubling of *SAJ* in the future might suggest because there are other (exogenous) reasons for job separation, including factors like shifts in tastes or in technology unrelated with automation.

This result may be surprising because the most pessimistic views about the automating abilities of AI and Robotics actually coincide with the most optimistic views of our model about the effects of these technologies on employment. Indeed, the most pessimistic scenario that we consider ($SAJ_{future} = 0.5$) would imply a staggering 5% fall in employment absent the change in job creation (contrasting with 0.5% in our benchmark with $SAJ_{future} = 0.09$). To understand our results, we have to consider that an increase in z_K has mainly two effects on job-creating decisions (excluding its effect on wages) when $\lambda_e = 0$. First, a rise in z_K motivates firms to replace workers with more profitable machines (z^* rises), which increases the continuation value of manual firms and, thus, motivates firm entry and job creation. Second, an increase in z_K allows firms to reap an even higher return from the automated technology in the circumstances (draws of z) in which they automate, giving rise to a scale effect. This scale effect also increases the continuation value of manual firms (promoting job creation) and is key to our results. A marginal increase in SAJ_{future} requires marginally higher z_K and z^*, which reinforce the effect of each other in job creation. As SAJ_{future} rises, the scale effect of a marginally higher z_K applies to a larger range of draws of z (because of higher z^*), amplifying its effect on job creation and giving rise to a convex change in the job-finding probability. Eventually, the required increase in z_K to reach a high SAJ_{future} is so large that job creation increases more than job destruction.

3.5. *Dissecting the mechanism: Ad hoc function for wages*

In our baseline model with the benchmark calibration (the first line in Table 2), the increase in the share of automated jobs leads simultaneously to lower employment and higher wages, which is uncommon in the literature and would not occur under a competitive labor market with risk-neutral agents. In our model, wages increase due to Nash bargaining: workers capture a share of the surplus, which increases due to more congestion in the labor market and the larger value of manual firms (namely because of a better outside option to move to the automated technology). Yet, the worker's productivity remains unchanged, implying that the rise in z_K squeezes the operational profits in the manual technology. So we ask: if wages were only a function of the task's productivity, how would the job-creation margin react to an increase in z_K? In other words, if wages would not increase with the rise in z_K, what would happen to employment?

To answer this question, we build a new version of the model in which we replace Nash bargaining with an ad hoc functional form for wages: $w(z) = (1 - \phi_{nb})b + \phi_{nb} z_L z$ $(0 < \phi_{nb} < 1)$.[23] Wages are the weighted sum of a constant term and the tasks' productivity. In this case, the improvement in the worker's and firm's outside option have no effect on the wage. Importantly, a rise in z_K has no effect on wages.

The third and fourth lines in Table 2 show the results in the model with ad hoc wages under two calibrations of η.[24] As expected, $f(\theta)$ rises substantially more in the model with ad hoc wages than in the model with Nash bargaining. As firms are not obliged to share the additional surplus generated by higher z_K, they are willing to create more jobs. This, in turn, leads to a positive change in employment in almost all our calibrations.

Yet, in our robustness checks, we find that employment still falls under some calibrations of the model (see Panel B in Table B2 in Appendix B). Matching frictions seem to be the cause as a calibration with lower hiring costs implies a substantially small reduction or a larger increase in employment (see Table B2). Furthermore, we reran our simulations with a lower matching function elasticity, η. If η is low, then the costs of a firm to match with a worker are less sensitive to labor market tightness ($\kappa_L(\theta) = \frac{\bar{\kappa}_L \theta^\eta}{\chi}$); in other words, for a given number of new manual firms, a low η reduces congestion, which is a symptom of matching frictions. The fourth line in Table 2 (and Panel C in Table B2) confirm that reducing the elasticity of $\kappa_L(\theta)$ improves employment outcomes, further suggesting that matching frictions prevent larger increases in employment. Therefore, in our model, the reduction in employment in many of our calibrations seems to be explained by a combination of higher wages and by the congestion caused by matching frictions.

Our experiments with the model assuming the ad hoc wage equation work as counterfactuals to understand the dynamics in our baseline model. But these experiments do not seem to be a good account of how AI and Robotics might affect employment in the future. Unless the historical positive relationship between labor market tightness and wage increments definitely breaks in the future, these innovations will increase wages, which may promote the negative employment effects that we obtain using our baseline model.

4. CES Aggregator: *Human touch*

In our baseline model, we assume that the tasks produced by workers and by machines are perfect substitutes. In this section, we instead build a model assuming that – from the perspective of consumers – they are imperfect substitutes. Our motivation for this setup is to take into account that consumers may deem differently a task produced by a machine or by a worker, a factor that we call *human touch*. For example, both a vending machine and a seller sell goods and, thus, they broadly perform the same task. Nonetheless, consumers may value the task differently on the basis of who is performing

---

[23] In this variant of the model, we abstract from the workers' side of the labor market and assume that they never endogenously break the match. We note, however, that if we considered the workers' side of the labor market, those with lowest wages would find it optimal to quit their jobs.

[24] In all cases, we continue to target SAJ_{future} to recalibrate z_K as explained in Section 3.2. To calibrate the model with the ad hoc wage, we fix ϕ_{nb} such that n^* (share of jobs with $z \in [z^*, z^*_e]$) is the same in the models with and without Nash bargaining. The other parameters are set internally to reach the nine current steady-state targets reported in Table 1.

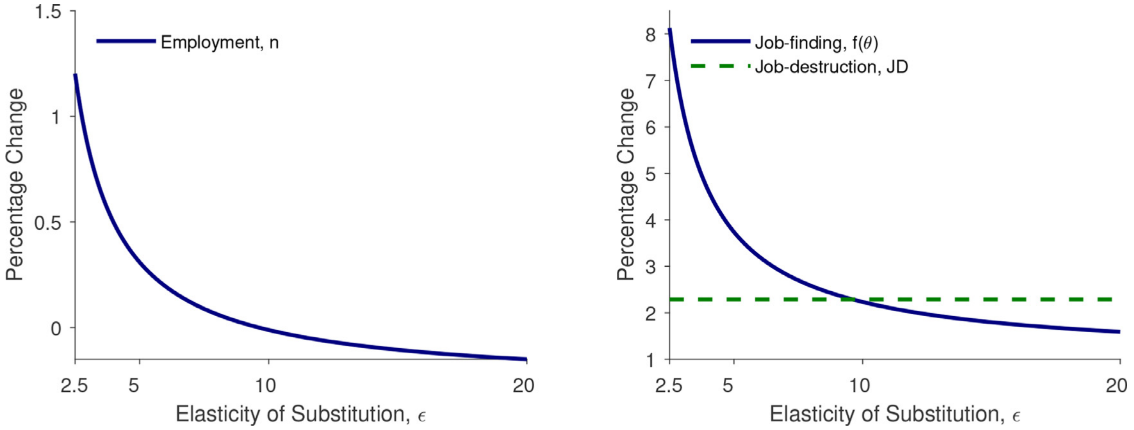

Fig. 3. The effect of the *Human Touch* (CES) under different values of ϵ *Note:* This figure shows the effects of an automation-augmenting shock in the variant with a CES aggregate production function, which captures the role of the *Human Touch*. The shock is calibrated such that this variant also reaches our target of $SAJ_{future} = 0.09$ under different values of the elasticity of substitution between the outputs of the automated and manual technologies, ϵ. The left panel shows the percentage change in employment. The right panel shows the percentage change in the job-finding probability (solid line) and the job-destruction probability (dashed line). A large ϵ implies the same results as in the baseline model.

it. The worker (seller) can offer a more personal (*human*) touch to the task whereas the machine (vending machine) offers an impersonal service.[25] This naturally renders machine and worker imperfect substitutes, from the perspective of the consumer. An ubiquitous use of the automated technology may, then, change the relative price of the tasks produced by machines and workers as consumers look for the differentiated offer of the manual technology. Our goal is to assess how the presence of the *human touch* (imperfect substitutability) affects the wrestle between the job-finding and job-destruction margins in determining how AI and Robotics affect employment. In particular, can the *human touch* reverse our prediction of lower employment in the benchmark and most robustness checks?

We implement this model by assuming a CES aggregator of the outputs of the tasks produced by the automated and manual technologies, y_K and y_L. In particular, the CES takes the following form:

$$y = A\left[\alpha_K y_K^{\frac{\epsilon-1}{\epsilon}} + \alpha_L y_L^{\frac{\epsilon-1}{\epsilon}}\right]^{\frac{\epsilon}{\epsilon-1}}, \tag{13}$$

where y is an index of final consumption (i.e., a bundle of goods and services demanded by consumers), ϵ is the elasticity of substitution, A is a scale parameter, and α_K and α_L are parameters affecting factor shares.[26] In this model, ϵ controls for the relevance of the *human touch* and is key for our results as shown in Figure 3. In the limit, when the outputs of the two technologies are perfect substitutes ($\epsilon = \infty$) and the *human touch* is nonexistent, the model with the CES reduces to our baseline model and employment falls. But, for roughly $\epsilon < 10$, there is a sufficiently important role for the *human touch* such that employment rises.

These results are explained by endogenous price movements. In the model with the CES and finite ϵ, a rise in z_K increases the relative price of the tasks produced by the manual technology relative to those produced by the automated technology. This is as if workers become more productive, which reinforces the incentives to create jobs for given job destruction. Thus, these experiments with the CES aggregator show that consumers play an important role in determining the effects of automation-augmenting shocks. If a large proportion of the tasks are directed to consumers, their preference for the *human touch* may severely reduce – and even reverse – the negative effects of automation on employment.

5. Policy

5.1. Robot tax

The expected surge in automation and, thus, job destruction has led many prominent figures to suggest a robot tax, i.e., that the usage/purchase of robots or automating technologies are taxed to prevent a surge in job losses. This has, in turn, prompted research towards understanding the consequences of the robot tax and how to optimize them (e.g., Prettner and Strulik, 2020; Guerreiro et al., 2021; and Gasteiger and Prettner, 2022). In this section, we contribute to the discussion by studying how a robot tax affects employment and wages in our baseline model. We start by recalibrating z_K to reach

---

[25] A different but related perspective is the distinction between hard and soft skills. Arguably, machines substitute more easily the hard than the soft skills. Furthermore, soft skills are highly valuable as suggested by the evidence in Aghion et al. (2019).

[26] See Appendix C for more details about this version of the model.

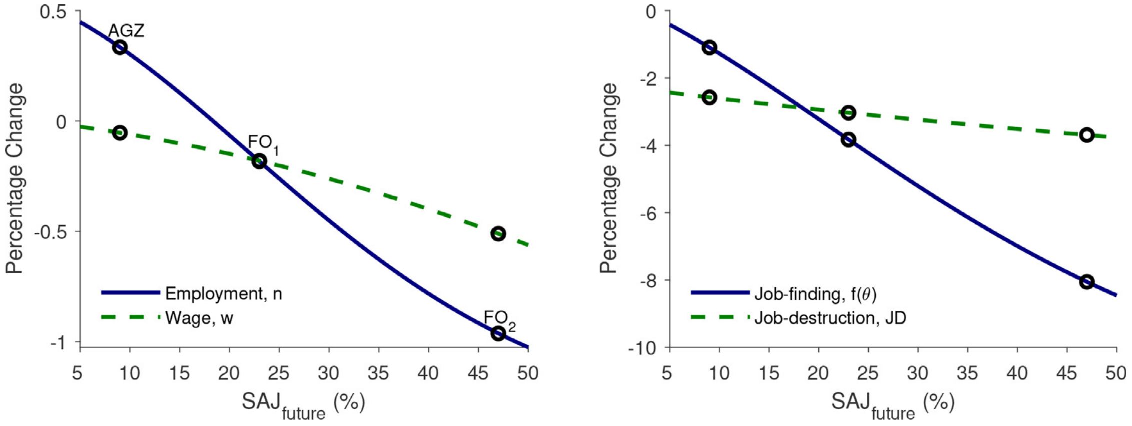

Fig. 4. Robot Tax for different values of SAJ_{future} *Note:* This figure shows the effects of a robot tax. The left panel shows the percentage change in employment (solid line) and in the wage in firms with $z = \bar{z}$ (dashed line). The right panel shows the percentage change in the job-finding probability (solid line) and the job-destruction probability (dashed line). Each circle identifies different estimates of SAJ_{future} found in the literature. The circles on the left (marked with *AGZ*) correspond to our benchmark, $SAJ_{future} = 0.09$, based on Arntz et al. (2017). The other two (marked with FO_1 and FO_2) correspond to the lower and upper bounds of the estimates in Frey and Osborne (2017), respectively.

the target for SAJ_{future}. Then, we conjecture that the government introduces a robot tax equal to 10% of the value of the automated technology, i.e., $\tau = 0.1 J_K$, to be paid upfront in addition to the cost of capital/robot, κ_K.

The results are depicted in Fig. 4 for different values of SAJ_{future}. The left panel shows the effects of the robot tax on employment and wages, while the right panel shows its effects on the job-finding and job-destruction probabilities. As expected, the robot tax dampens job destruction: given that it is more costly to automate a task, firms automate and fire less. Yet, the robot tax also lowers job creation due to the fall in the continuation value of jobs. To make matters worse, the robot tax might inadvertently lower employment in the scenarios in which it seems more important; i.e., when there is more automation (SAJ_{future} is higher). This result echoes our findings in Section 3.4. A higher SAJ_{future} is only achievable if z_K rises substantially and workers lose much of their comparative advantage. When this occurs, the elasticity of employment with respect to z_K becomes positive because the scale effect of higher $\beta J_K - \kappa_K$ overcomes its substitution effect. Therefore, a robot tax (which effectively lowers the value of investing in the automated technology) when z_K is large dampens employment.

Figure 4 also shows that wages (evaluated at $\bar{z}$) fall irrespective of SAJ_{future} due to the robot tax. Put simply, policymakers face a trade-off: using a robot tax, they can only raise employment at the expense of lower wages. Given that the robot tax lowers $\beta J_K - \kappa_K$ in a similar way as a lower z_K would, the robot tax has the mirror implications of the automation-augmenting shock studied so far, which might raise or reduce employment but always rises wages.

5.2. Automation tax & other policies

Policymakers might, alternatively, consider two policies that are common in their toolbox: hiring subsidies and firing taxes. Figure D2 in Appendix D shows that a hiring subsidy that lowers $\bar{\kappa}_L$ increases employment and wages irrespective of λ_e and SAJ_{future}, which may make it a better option than the robot tax, albeit at higher (direct) budgetary costs. Firing taxes, on the other hand, have a largely detrimental effect on employment in our model. They lower job creation substantially due to the fall in the value of hiring workers, but they barely affect job destruction because most separations are exogenous.

Given the implications of firing taxes, policymakers might instead consider an *automation tax*. An automation tax is a tax paid by a firm if and when it automates the production of the task and replaces a worker by a robot; in other words, the automation tax is a firing tax that only applies to endogenous (automation-related) separations.[27] The automation tax lowers the outside option of the manual firm in Eqs. (8) and (10) by the tax amount, τ. In our benchmark, with $\lambda_e = 0$, an automation tax and a robot tax are the same because whenever a firm invests in a robot it is also automating production and replacing a worker. But the two taxes differ when we let $\lambda_e > 0$. In this case, some firms are not technologically constrained and invest in robots immediately when they enter; these firms do not pay the automation tax but must pay the robot tax.

To understand the distinction between an automation tax and a robot tax, Fig. 5 compares their effect on employment and wages for different levels of λ_e with $SAJ_{future} = 0.09$. The two taxes are of the same amount, $\tau = 0.1 J_K$, and their consequences are studied in the future steady-state (i.e., after we recalibrate z_K to reach the target for SAJ_{future}). When $\lambda_e = 0$, the results are the same as the ones reported in Fig. 4. But when $\lambda_e > 0$, the robot tax may lead to much less employment and wages than the automation tax. In fact, if $SAJ_{future} = 0.09$, a robot tax lowers employment if $\lambda_e \geq 0.13$, while an automation

---

[27] The automation tax is in the same spirit as the Trade Adjustment Assistance Program in the US, which aids workers whose labor market outcomes are negatively affected due to increased imports.

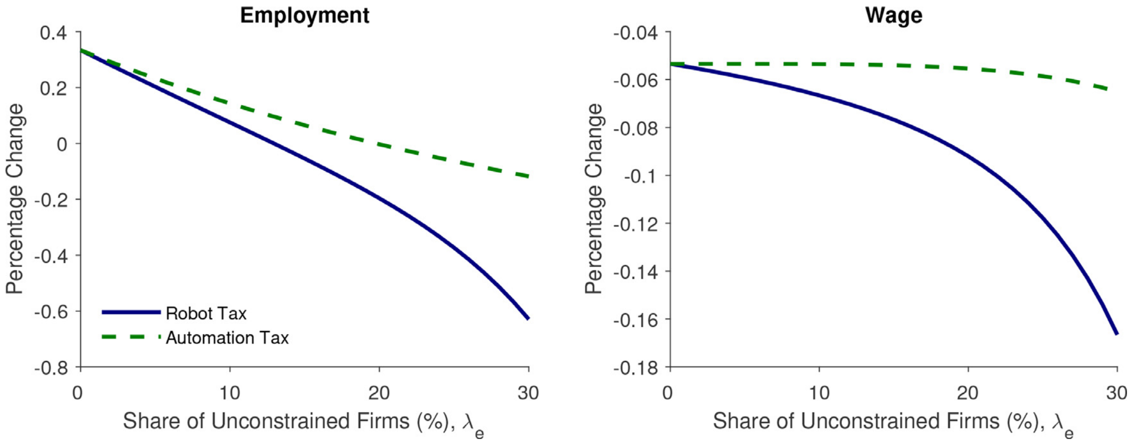

Fig. 5. Automation Tax Vs Robot Tax *Note:* The left (right) panel shows the percentage change in employment (wage when $z = \bar{z}$) when a robot tax is implemented (solid line) and when an automation tax is implemented (dashed line).

tax only lowers employment if $\lambda_e \geq 0.2$. The reason is that a robot tax also lowers the general-equilibrium effects of the rise in z_K, which, as shown in Section 3.3, greatly increases employment and wages relative to when $\lambda_e = 0$. On the contrary, an automation tax is directed towards preventing the automation of a task performed by a worker; thus, it does not directly reduce the entry-related general-equilibrium effects.

Our results suggest that an automation tax is preferred to both a firing tax and a robot tax. Moreover, given that it increases government revenues, the automation tax might also be a better option than the hiring subsidy. Yet, a problem with the automation tax is its practicality. Machines are usually not able to replace workers in all tasks, making the one-to-one relation between machines and workers blurry in practice. Furthermore, it is difficult (if not impossible) to distinguish an automation-related separation from exogenous separations. For this reason, we consider the possibility that some exogenous separations are erroneously considered automation-related in Fig. D1, in Appendix D. This figure shows that if $\lambda_e = 0.3$ and 1% or more of exogenous separations are considered automation-related, then the robot tax has less of a negative effect on employment than the automation tax. To conclude, taking all these results together, the best option might be a policy mix. For example, policymakers might use robot (or automation) taxes to finance hiring subsidies as this mix could increase employment even if the taxes directly lower it.[28]

6. Concluding remarks

In face of the expected progress in AI and Robotics, we build a general-equilibrium model to take on the task of predicting the future of employment. The empirical literature growingly suggests that past automating technologies have favored employment growth (e.g., Bessen, 2016; Autor and Salomons, 2018; Bessen et al., 2020; and Gregory et al., 2021). Yet, AI and Robotics are expected to disrupt the labor market by much more than previous automating technologies (e.g., Brynjolfsson and McAfee, 2011; Brynjolfsson et al., 2017; Frey and Osborne, 2017). Our model suggests that employment will likely fall but not massively.

In our model, advances in AI and Robotics increase the dynamism in the economy, which enlarges labor market flows. On the one hand, the shock increases job destruction because of the higher probability of automating production. On the other hand, due to either a sort of complementarity at the time of entry (as in Guimarães and Gil, 2022) or because hiring a worker is a crucial first step in starting the production of a task (as in Acemoglu and Restrepo, 2018), firm entry and job creation also increase. In turn, this increase in job creation prevents massive unemployment and may even imply more employment in the future.

The increase in labor market flows predicted by our model contrasts with US data showing a downward trend in flows for the last decades (Davis and Haltiwanger, 2014). Yet, this trend seems mostly driven by composition effects that our model abstracts from. For example, Hyatt and Spletzer (2017) and Pries and Rogerson (2022) document the fall in labor market fluidity reflects a substantial reduction of short-lived employment spells, and Molloy et al. (2016) document that after controlling for demographics and education, there is no apparent downward trend in job separation and job finding rates. Furthermore, other evidence points in the direction of increasing flows, concurring with our model. Fujita (2018) documents a positive trend in occupational mobility from 1976 to 2016, which might be related with automation forcing workers to adapt to new tasks and occupations; and Bloom et al. (2018) relate the growing *gig economy* (in which permanent jobs are less common and workers are hired on demand) with the widespread use of automating technologies.

---

[28] Another option with likely similar consequences would be to mix entry subsidies (that reduces Ω) with robot taxes.

In our baseline model, the increase in job destruction is usually larger than the increase in job creation, reducing employment. But, a key factor that significantly affects the predictions of our model is the relevance and prevalence of what we call *human touch. Human touch* refers to a consumers' preference for diversity in the producer/provider of the task itself: in a world with widespread usage of machines to offer multiple services to consumers, they may value the differentiated service of a human (i.e., human interactions). If that is the case, the dissemination of machines/robots increases the relative price of the tasks produced by workers. Hence, for a given job destruction, *human touch* promotes job creation, reversing the prediction of less employment for a large range of parameters.

There are multiple elements that might affect the future of employment but are out of scope of our paper. One is that we abstract from the consequences of improvements in AI and Robotics in the transition between steady-states. As we only study the long-run effects of automating technologies, our analysis might miss massive unemployment in the transition because those that lose jobs are not necessarily those who find new ones.[29] On the other hand, massive unemployment in the transition seems less likely as large implementation lags in AI and Robotics are apparently slowing the transition relative to expectations (Brynjolfsson et al., 2017; Naudé, 2020). Indeed, the record employment rates in the US and other advanced economies prior to the Covid-19 pandemic are a testament of the unexpected behavior in the labor market. Yet, the Covid-19 pandemic might boost automation and unemployment. Jaimovich and Siu (2020) document that almost all of the jobs lost in routine occupations occurred in downturns. The mounting news of firms investing in automating technologies raise concerns that the pattern will repeat once again as implementation lags of AI and Robotics seem to have been shortened during the pandemic.

Looking at the long-run effects of progress in AI and Robotics – the focus of our paper – there are a few more reasons for optimism and pessimism not considered in our analysis. One reason for optimism is that advances in AI might increase matching efficiency (e.g., by reducing the costs of screening job applicants), which raises the quality of firm-worker matches and/or employment (Martellini and Menzio, 2020). Another reason for optimism is that AI automates production but also directly increases the value of final goods and services. For example, durable goods like cameras integrate AI to improve color accuracy at different times of the day and resolution at different distances, and washing machines integrate AI to adjust speed, water consumption, and other factors to the type of clothes being washed. This increased quality and value of final goods raises the value of all tasks: inasmuch as final goods and services rely on multiple tasks, some produced by labor and others by machines/robots, the value of labor tasks goes up, which might increase employment and wages beyond the predictions of our model. A reason for pessimism, however, is that automation might lead to accelerated deskilling (see the discussion in Atack et al., 2019). In an economy in which workers are often forced to adapt to new tasks and occupations (due to more automation and larger flows), the value of the human capital owned by workers might depreciate more rapidly, fostering employment losses. Related to this, there is also the concern of heterogeneous effects of automation, where the workers that lose their jobs because of automation become long-term unemployed, while others benefit from higher labor demand. We will consider the consequences of these two points in a future paper.

Appendix A. Output and the labor share

To quantify output, we only need to sum the output produced by manual and automated firms because we assume that tasks are perfect substitutes. The output of automated firms is $z_K n_K$ as all these firms produce z_K. But it is not as simple to determine the output of manual firms because they are not distributed according to $G(z)$ from z^* to $\bar{z}$. To measure output, we need to distinguish between three groups of manual firms: we need to calculate how many manual firms produce tasks with productivity (i) $\bar{z}$ from the moment they were created and have not redrawn productivity afterwards, (ii) above z_e^* (by means of draws or redraws of z), and (iii) between z^* and z_e^* (by means of redraws of z). We denote the latter two as n_e^* and n^*, respectively. And we obtain the number of firms producing tasks with productivity $\bar{z}$ from inception as the residual: $n - n_e^* - n^*$.

There are two ways in which a manual firm may produce a task with idiosyncratic productivity above z_e^* and belong to n_e^*: either the productivity of the task was drawn at the time of entry or it was later redrawn in the interval $[z_e^*, \bar{z}]$. The number of manual firms that draw productivity at the time of entry is $\frac{\lambda_e(1-G(z_e^*))}{1-\lambda_e G(z_e^*)} f(\theta)(1-n)$. This follows from two factors. First, every period, $f(\theta)(1-n)$ new manual firms are created, which correspond to a fraction $1-\lambda_e G(z^*)$ of all new firms (the remaining fraction chooses automation at entry). Second, only a proportion $\lambda_e(1-G(z_e^*))$ of all new firms draw z and choose the manual technology. Furthermore, the number of manual firms that redraw productivity and obtain z above z_e^* is $(1-\delta_L)\lambda_n(1-G(z_e^*))n$ given that a proportion $1-\delta_L$ of manual firms survive exogenous shocks and a proportion λ_n redraw productivity. But some of these firms were already included in n_e^*; thus, the net inflow of firms by redrawing productivity into n_e^* is only $(1-\delta_L)\lambda_n(1-G(z_e^*))(n-n_e^*)$.

There are also two ways in which a manual firm leaves n_e^*: either the firm ends exogenously or it draws productivity below z_e^*. These exit flows sum to $(\delta_L + (1-\delta_L)\lambda_n G(z_e^*))n_e^*$. Combining the flows into and out of n_e^* implies after a few

---

[29] A few examples are the cases of artisans in the first industrial revolution and US telephone operators in the early 20th century (Feigenbaum and Gross, 2020). See also the evidence in Jaimovich et al. (2021) on the consequences of the disappearance of routine occupations in the last decades in the US.

derivations:

$$n_e^* = \frac{(1-\delta_L)\lambda_n(1-G(z_e^*))n}{\delta_L+(1-\delta_L)\lambda_n} + \frac{\frac{\lambda_e(1-G(z_e^*))f(\theta)(1-n)}{1-\lambda_e G(z_e^*)}}{\delta_L+(1-\delta_L)\lambda_n}. \tag{A1}$$

We can apply a similar logic to find the firms that produce tasks with idiosyncratic productivity between z^* and z_e^*. Making the necessary adjustments and taking into account that no firm starts in the manual technology with productivity between z^* and z_e^*, we obtain

$$n^* = \frac{(1-\delta_L)\lambda_n}{\delta_L+(1-\delta_L)\lambda_n}(G(z_e^*)-G(z^*))n. \tag{A2}$$

Having established the number of firms, we quantify output as

$$y = n_K z_K + (n-n^*-n_e^*)z_L\bar{z} + n_e^*\frac{1}{1-G(z_e^*)}\int_{z_e^*}^{\bar{z}} z\mathrm{d}G(z) + n^*\frac{1}{G(z_e^*)-G(z^*)}\int_{z^*}^{z_e^*} z\mathrm{d}G(z), \tag{A3}$$

in which we multiply the number of firms in each group by its respective average output. The labor share then is ratio of the number of workers in each group of manual firms (recall that every manual firm employs one worker) multiplied by its respective average wage relative to output:

$$LS = \frac{(n-n^*-n_e^*)w(\bar{z}) + n_e^*\frac{1}{1-G(z_e^*)}\int_{z_e^*}^{\bar{z}} w(z)\mathrm{d}G(z) + n^*\frac{1}{G(z_e^*)-G(z^*)}\int_{z^*}^{z_e^*} w(z)\mathrm{d}G(z)}{y}. \tag{A4}$$

Appendix B. Other robustness checks

We consider several robustness checks of the calibration of the baseline model in Table B1. First, we report the results with the benchmark calibration and with $\lambda_e = 0.2$ as in the first two lines in Table 2. Then, we show that the remaining

Table B1
Robustness checks of the baseline model.

	$\%\Delta n$	$\%\Delta f(\theta)$	$\%\Delta JD$	$\%\Delta LS$	$\%\Delta w(\bar{z})$	$\%\Delta z_K$
Benchmark	-0.28	1.01	2.29	-1.27	0.05	3.87
$\lambda_e = 0.2$	0.25	3.46	2.29	-1.90	0.14	8.05
$\eta = 0.4$	-0.23	1.22	2.29	-1.27	0.05	3.86
$\phi = 0.4$	-0.20	1.36	2.29	-1.28	0.04	3.43
Pareto ($\xi = 4$)	-0.22	1.26	2.29	-1.51	0.04	0.68
Non-routine= 65%	-0.20	1.35	2.29	-1.28	0.06	4.93
Half Hiring Costs	-0.14	1.63	2.29	-1.28	0.04	3.44
Half Robot Inst. Cost	-0.27	1.02	2.29	-1.08	0.05	5.66
$JD = 0.03$	-0.23	0.21	1.25	-0.88	0.02	1.79
$SAJ = 0.06$	-0.17	0.73	1.52	-0.85	0.03	2.41

*Note:*This table shows the effects of an automation-augmenting shock calibrated such that $SAJ_{future} = 0.09$. The six columns show, respectively, the percentage change in employment, job-finding probability, job-destruction probability, labor share, wages (in firms with draw $z = \bar{z}$), and the productivity of the automated technology. The first line uses the benchmark calibration in Table 1. Each other line changes one parameter or steady-state target. The exception is the Pareto line, in which we replace the uniform distribution with a Pareto distribution with shape parameter equal to $\xi = 4$.

Table B2
Robustness checks of the model with ad hoc wage*Note:* This table shows the effects of an automation-augmenting shock calibrated such that $SAJ_{future} = 0.09$ in two models, one with Nash Bargaining (Panel A) and another with an ad hoc wage (Panels B and C). Panels B and C only differ in terms of the value of η; Panel B assumes $\eta = 0.5$, whereas Panel C assumes $\eta = 0.3$. The change in job destruction (unreported) is the same in all panels. See Table B1 and the note to the table for more details.

	A: Baseline (Nash)		B: Ad hoc ($\eta = 0.5$)		C: Ad hoc ($\eta = 0.3$)	
	$\%\Delta n$	$\%\Delta f(\theta)$	$\%\Delta n$	$\%\Delta f(\theta)$	$\%\Delta n$	$\%\Delta f(\theta)$
Benchmark	-0.28	1.01	0.69	5.55	2.21	13.44
$\lambda_e = 0.2$	0.25	3.46	3.34	19.92	7.84	52.78
$\phi = 0.4$	-0.20	1.36	0.79	6.05	2.44	14.70
Pareto ($\xi = 4$)	-0.22	1.26	0.07	2.62	0.82	6.21
Non-routine= 65%	-0.20	1.35	1.06	7.42	3.05	18.17
Half Hiring Costs	-0.14	1.63	1.80	11.24	4.66	28.21
Half Robot Inst. Cost	-0.27	1.02	0.70	5.62	2.24	13.62
$JD = 0.03$	-0.23	0.21	-0.03	1.12	0.30	2.64
$SAJ = 0.06$	-0.17	0.73	0.53	4.00	1.64	9.58

parameters have a slight quantitative effect but not a qualitative effect. The same is true when we replace the uniform distribution with the Pareto distribution.

Our robustness checks show that if workers have a lower bargaining power ($\phi = 0.4$), then employment falls less as firms share less of the gains in their value due to the rise in z_K with workers. This agrees with the results in Section 3.5. The results with $\eta = 0.4$ also concur with our findings in Section 3.5; yet, the effect of η is severely mitigated because of Nash bargaining. We also consider a Pareto distribution with a shape parameter of $\xi = 4$ (because we only obtained convergence in our model if $\xi > 3.4$); the results are not much sensitive to the change in distribution. The same is true when we increase the share of non-routine jobs and, especially, when we consider half the robot installation costs as our target ($\kappa_K = 6w(z^*)$).

We do find a bigger change when our target for hiring costs is half that in the benchmark ($\kappa_L(\theta) = \frac{10}{160} z_L \bar{z}$) as hiring costs become relatively less important in hiring decisions. Moreover, a higher value of *SAJ* in the current steady-state lowers the change in employment because the required shock to reach $SAJ_{future} = 0.09$ is smaller. Finally, a higher job destruction target ($JD = 0.03$) leads to less employment response as automation becomes less relevant in the overall flows in the labor market. We consider $JD = 0.03$ and not other values because, according to our calculations, that is the proportion of prime-age workers that move from employment to nonemployment (both unemployment and out of the labor force) in CPS data from 1978 to 2012.

We conclude this section with several robustness checks of the calibration of the model with the ad hoc wage in Table B2.

Appendix C. Model with *Human Touch*

The model with human touch is similar to the baseline, except that we assume the CES aggregator of the outputs of the tasks produced by automated and manual technologies in Eq. (13). The sum of the outputs produced using each type of technology are:

$$y_K = z_K n_K,$$

$$y_L = z_L \left[(n - n^* - n_e^*)\bar{z} + n_e^* \frac{1}{1 - G(z_e^*)} \int_{z_e^*}^{\bar{z}} z \mathrm{d}G(z) + n^* \frac{1}{G(z_e^*) - G(z^*)} \int_{z^*}^{z_e^*} z \mathrm{d}G(z). \right]$$

Assuming competitive markets in the intermediate goods y_K and y_L and a profit-maximizing final-good producer, we get:

$$p_K = \alpha_K y_K^{-\frac{1}{\epsilon}} y^{\frac{1}{\epsilon}}, \tag{C1}$$

$$p_L = \alpha_L y_L^{-\frac{1}{\epsilon}} y^{\frac{1}{\epsilon}}. \tag{C2}$$

To calibrate this variant of our model, we fix z_K and $\bar{z}$ at values that satisfy our current steady-state targets when $\epsilon = \infty$. For finite values of ϵ, we fix $\alpha_L = 1$ and use A and α_K to replace the roles of z_K and $\bar{z}$ in reaching our current steady-state targets as specified in Table 1. Table C1 complements Fig. 3 by showing the effect of the *human touch* under different calibrations of the model when $\epsilon = 10$.

Table C1
Robustness checks of the model with the *human touch* (CES).

	A: Baseline		B: CES ($\epsilon = 10$)	
	$\%\Delta n$	$\%\Delta f(\theta)$	$\%\Delta n$	$\%\Delta f(\theta)$
Benchmark	-0.28	1.01	-0.01	2.23
$\lambda_e = 0.2$	0.25	3.46	0.64	5.31
$\eta = 0.4$	-0.23	1.22	0.09	2.70
$\phi = 0.4$	-0.20	1.36	0.19	3.17
Pareto ($\xi = 4$)	-0.22	1.26	1.15	7.85
Non-routine= 65%	-0.20	1.35	0.08	2.64
Half Hiring Costs	-0.14	1.63	0.32	3.79
Half Robot Inst. Cost	-0.27	1.02	-0.15	1.57
$JD = 0.03$	-0.23	0.21	-0.11	0.77
$SAJ = 0.06$	-0.17	0.73	0.01	1.57

*Note:*This table shows the effects of an automation-augmenting shock calibrated such that $SAJ_{future} = 0.09$ in two models, our baseline (Panel A) and another with a CES (Panel B) with an elasticity of substitution equal to 10. Panel B, thus, corresponds to the case with *Human Touch*. The change in job destruction (unreported) is the same in both panels. See Table B1 and the note to the table for more details.

Appendix D. Other policies

In this appendix, we extend the automation tax with the possibility that some exogenous separations are erroneously considered automation-related; if this occurs, then the firm must pay the automation tax even if it does not replace the worker with a machine/robot. Figure D1 shows how the consequences of the automation tax change with the proportion of exogenous separations that are deemed automation-related. When this proportion is zero, the reported changes in employment and wages are the same as in Fig. 5 (when $\lambda_e = 0.3$). In the other limit in which all exogenous separations imply a tax (a standard firing tax; not reported), employment and wages fall massively as hinted by the case in which 10% of exogenous separations mistakenly lead to the automation tax. Figure D1 also includes the effects of a robot tax, which does not depend on the proportion of exogenous separations that are deemed automation-related as it is always paid when a robot is purchased.

Finally, Fig. D2 shows how the effects of a hiring subsidy (equal to 10% of $\kappa_L(\theta)$) on employment and wages depend on SAJ_{future} and λ_e.

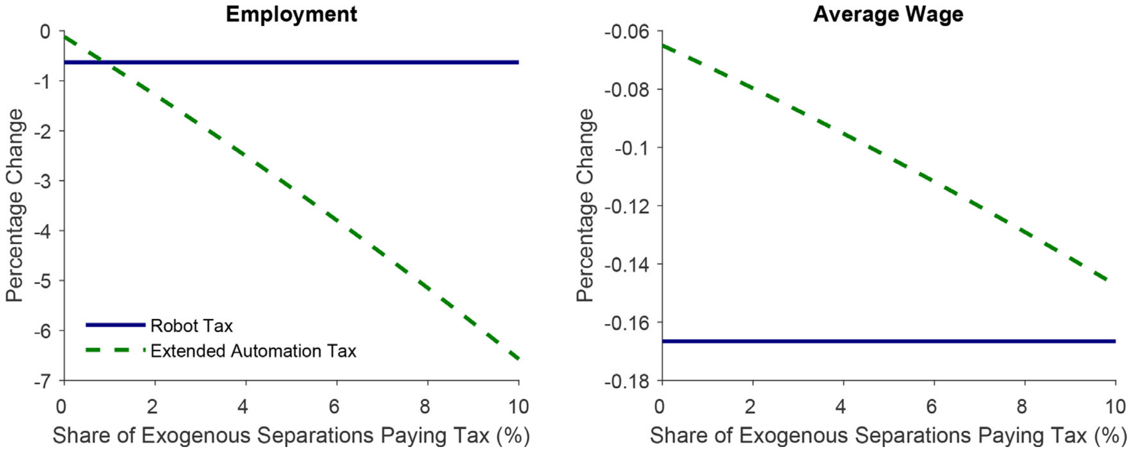

Fig. D1. Extended Automation Tax Vs Robot Tax ($\lambda_e = 0.3$) *Note:* The left (right) panel shows the percentage change in employment (wage when $z = \bar{z}$) when a robot tax is implemented (solid line) and when an extended automation tax is implemented (dashed line). In the case of the extended automation tax, all endogenous separations and a share of the exogenous separations pay taxes. If the share is zero, the tax is our baseline automation tax; if the share is 100%, it is a standard firing tax.

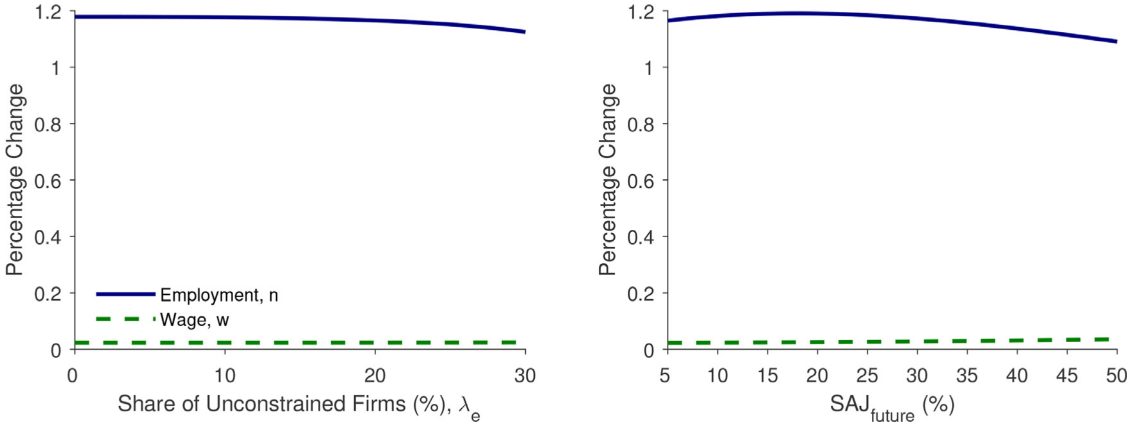

Fig. D2. Hiring subsidy (10% of $\bar{\kappa}_L$) *Note:* This figure shows the percentage change in employment (solid lines) and wage (when $z = \bar{z}$; dashed line) when the hiring subsidy is implemented.

References

Acemoglu, D., Restrepo, P., 2018. The race between man and machine: implications of technology for growth, factor shares, and employment. Am. Econ. Rev. 108 (6), 1488–1542. doi:10.1257/aer.20160696.

Acemoglu, D., Restrepo, P., 2020. Robots and jobs: evidence from US labor markets. J. Polit. Economy 128 (6), 2188–2244. doi:10.1086/705716.

Aghion, P., Bergeaud, A., Blundell, R.W., Griffith, R., 2019. The Innovation Premium to Soft Skills in Low-Skilled Occupations. Available at SSRN 3489777.

Alesina, A., Battisti, M., Zeira, J., 2018. Technology and labor regulations: theory and evidence. J. Econ. Growth 23 (1), 41–78.

Arnoud, A., 2018. Automation Threat and Wage Bargaining. Unpublished Manuscript. Yale University.

Arntz, M., Gregory, T., Zierahn, U., 2017. Revisiting the risk of automation. Econ. Lett. 159, 157–160. doi:10.1016/j.econlet.2017.07.001.

Atack, J., Margo, R.A., Rhode, P.W., 2019. "Automation" of manufacturing in the late nineteenth century: the hand and machine labor study. J. Econ. Perspect. 33 (2), 51–70. doi:10.1257/jep.33.2.51.

Autor, D., Salomons, A., 2018. Is automation labor share–displacing? Productivity growth, employment, and the labor share. Brookings Pap. Econ. Act. 1–63.

Autor, D.H., 2015. Why are there still so many jobs? the history and future of workplace automation. J. Econ. Perspect. 29 (3), 3–30. doi:10.1257/jep.29.3.3.

Barron, J.M., Bishop, J., 1985. Extensive search, intensive search, and hiring costs: new evidence on employer hiring activity. Econ. Inq. 23 (3), 363–382. doi:10.1111/j.1465-7295.1985.tb01773.x.

Basso, H.S., Jimeno, J.F., 2021. From secular stagnation to robocalypse? Implications of demographic and technological changes. J. Monet. Econ. 117, 833–847. doi:10.1016/j.jmoneco.2020.06.004.

Berg, A., Buffie, E.F., Zanna, L.-F., 2018. Should we fear the robot revolution? (The correct answer is yes). J. Monet. Econ. 97, 117–148. doi:10.1016/j.jmoneco.2018.05.014.

Bessen, J., Goos, M., Salomons, A., van den Berge, W., 2020. Automation: A guide for policymakers.

Bessen, J.E., 2016. How Computer Automation Affects Occupations: Technology, Jobs, and Skills. Boston Univ. School of Law, Law and Economics Research Paper 15-49.

Bloom, D.E., McKenna, M., Prettner, K., 2018. Demography, Unemployment, Automation, and Digitalization: Implications for the Creation of (Decent) Jobs, 2010–2030. Working Paper 24835. National Bureau of Economic Research doi:10.3386/w24835. http://www.nber.org/papers/w24835

Brynjolfsson, E., McAfee, A., 2011. Race Against the Machine. Digital Frontier, Lexington, MA.

Brynjolfsson, E., Rock, D., Syverson, C., 2017. Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics. Working Paper 24001. National Bureau of Economic Research doi:10.3386/w24001. http://www.nber.org/papers/w24001

Caselli, F., Manning, A., 2019. Robot arithmetic: new technology and wages. Am. Econ. Rev. 1 (1), 1–12. doi:10.1257/aeri.20170036.

Coles, M.G., Kelishomi, A.M., 2018. Do job destruction shocks matter in the theory of unemployment? Am. Econ. J. 10 (3), 118–136. doi:10.1257/mac.20150040.

Cords, D., Prettner, K., 2021. Technological unemployment revisited: automation in a search and matching framework. Oxf. Econ. Pap. doi:10.1093/oep/gpab022. https://academic.oup.com/oep/advance-article-pdf/doi/10.1093/oep/gpab022/38664661/gpab022.pdf

Davis, S.J., Haltiwanger, J., 2014. Labor Market Fluidity and Economic Performance. Working Paper 20479. National Bureau of Economic Research doi:10.3386/w20479. http://www.nber.org/papers/w20479

Feigenbaum, J., Gross, D.P., 2020. Automation and the Fate of Young Workers: Evidence from Telephone Operation in the Early 20th Century. Working Paper 28061. National Bureau of Economic Research doi:10.3386/w28061. http://www.nber.org/papers/w28061

Fornino, M., Manera, A., 2022. Automation and the future of work: assessing the role of labor flexibility. Rev. Econ. Dyn. 45, 282–321. doi:10.1016/j.red.2021.07.002.

Frey, C.B., Osborne, M.A., 2017. The future of employment: how susceptible are jobs to computerisation? Technol. Forecast. Soc. Change 114, 254–280. doi:10.1016/j.techfore.2016.08.019.

Fujita, S., 2018. Declining labor turnover and turbulence. J. Monet. Econ. 99, 1–19. doi:10.1016/j.jmoneco.2017.12.005.

Fujita, S., Ramey, G., 2012. Exogenous versus endogenous separation. Am. Econ. J. 4 (4), 68–93.

Gasteiger, E., Prettner, K., 2022. Automation, stagnation, and the implications of a robot tax. Macroecon. Dyn. 26 (1), 218–249. doi:10.1017/S1365100520000139.

Graetz, G., Michaels, G., 2018. Robots at work. Rev. Econ. Stat. 100 (5), 753–768. doi:10.1162/rest_a_00754.

Gregory, T., Salomons, A., Zierahn, U., 2021. Racing with or against the machine? Evidence on the role of trade in Europe. J. Eur. Econ. Assoc. doi:10.1093/jeea/jvab040.

Guerreiro, J., Rebelo, S., Teles, P., 2021. Should robots be taxed? Rev. Econ. Stud. 89 (1), 279–311. doi:10.1093/restud/rdab019.

Guimarães, L., Gil, P.M., 2022. Explaining the labor share: automation vs labor market institutions. Labour Econ. 75, 102146. doi:10.1016/j.labeco.2022.102146.

Hall, R.E., Kudlyak, M., 2019. Job-Finding and Job-Losing: A Comprehensive Model of Heterogeneous Individual Labor-Market Dynamics. Working Paper 25625. National Bureau of Economic Research doi:10.3386/w25625.

Hall, R.E., Milgrom, P.R., 2008. The limited influence of unemployment on the wage bargain. Am. Econ. Rev. 98 (4), 1653–1674.

Hornstein, A., Krusell, P., Violante, G.L., 2007. Technology-policy interaction in frictional labour-markets. Rev. Econ. Stud. 74 (4), 1089–1124. http://www.jstor.org/stable/4626174

Hyatt, H.R., Spletzer, J.R., 2017. The recent decline of single quarter jobs. Labour Econ. 46, 166–176. doi:10.1016/j.labeco.2016.10.001.

Jaimovich, N., Saporta-Eksten, I., Siu, H., Yedid-Levi, Y., 2021. The macroeconomics of automation: data, theory, and policy analysis. J. Monet. Econ. 122, 1–16. doi:10.1016/j.jmoneco.2021.06.004.

Jaimovich, N., Siu, H.E., 2020. Job polarization and jobless recoveries. Rev. Econ. Stat. 102 (1), 129–147. doi:10.1162/rest_a_00875.

Kambourov, G., Manovskii, I., 2008. Rising occupational and industry mobility in the United States: 1968–97. Int. Econ. Rev. 49 (1), 41–79. http://www.jstor.org/stable/20486788

Leduc, S., Liu, Z., 2019. Robots or workers? A macro analysis of automation and labor markets.

Ljungqvist, L., Sargent, T.J., 2004. European unemployment and turbulence revisited in a matching model. J. Eur. Econ. Assoc. 2 (2/3), 456–468. http://www.jstor.org/stable/40004919

Martellini, P., Menzio, G., 2020. Declining search frictions, unemployment, and growth. J. Polit. Economy 128 (12), 4387–4437. doi:10.1086/710975.

Molloy, R., Trezzi, R., Smith, C.L., Wozniak, A., 2016. Understanding declining fluidity in the us labor market. Brookings Pap. Econ. Act. 2016 (1), 183–259.

Mortensen, D.T., Pissarides, C., 1994. Job creation and job destruction in the theory of unemployment. Rev. Econ. Stud. 61 (3), 397–415.

Mortensen, D.T., Pissarides, C.A., 1998. Technological progress, job creation, and job destruction. Rev. Econ. Dyn. 1 (4), 733–753. doi:10.1006/redy.1998.0030.

Nakamura, H., Zeira, J., 2018. Automation and unemployment: help is on the way.

Naudé, W., 2020. Artificial intelligence: neither utopian nor apocalyptic impacts soon. Econ. Innov. New Technol. 0 (0), 1–23. doi:10.1080/10438599.2020.1839173.

Petrongolo, B., Pissarides, C.A., 2001. Looking into the black box: a survey of the matching function. J. Econ. Lit. 39 (2), 390–431. http://www.jstor.org/stable/2698244

Pissarides, C.A., 2009. The unemployment volatility puzzle: is wage stickiness the answer? Econometrica.
Prettner, K., Strulik, H., 2020. Innovation, automation, and inequality: policy challenges in the race against the machine. J. Monet. Econ. 116, 249–265. doi:10.1016/j.jmoneco.2019.10.012.
Pries, M.J., Rogerson, R., 2022. Declining worker turnover: the role of short-duration employment spells. Am. Econ. J. 14 (1), 260–300. doi:10.1257/mac.20190230.
Shimer, R., 2012. Reassessing the ins and outs of unemployment. Rev. Econ. Dyn. 15 (2), 127–148. doi:10.1016/j.red.2012.02.001.
Zeira, J., 1998. Workers, machines and economic growth. Q. J. Econ. 113, 1091–1113.
Zeira, J., 2010. Machines as Engines of Growth. Mimeo.