Looking Ahead at the Effects of Automation in an Economy with Matching Frictions *

Luís Guimarães[†] & Pedro Mazeda Gil[‡]

December 10, 2020

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. Therefore, they may reduce long-run employment 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, our model suggests that an automation tax trumps a robot tax.

JEL classification: E24; J64; L11; O33.
Keywords: Automation; Employment; Labor-Market Frictions; Technology Choice.

*We thank the helpful comments by Marta Alvaro-Taus, Murat Çelik, Alper Çenesiz, Davide Debortoli, Sarah Guillou, Miguel León-Ledesma, Paul Levine, 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 Portuguese public funds through FCT - Fundação para a Ciência e a Tecnologia, I.P., in the framework of the project UIDB/04105/2020 (cef.up – Center for Economics and Finance at University of Porto).

[†] Queen's University Belfast and cef.up. E-mail address: l.guimaraes@qub.ac.uk. Address: Queen's Management School, Queen's University Belfast, 185 Stranmillis Rd, Belfast BT9 5EE, United Kingdom.

[‡] Faculdade de Economia, Universidade do Porto and cef.up. E-mail address: pgil@fep.up.pt.

1 Introduction

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

We use a general-equilibrium model to answer this question. In the model, we conceptualize advances in AI and Robotics as an aggregate increase in the comparative advantage of machines relative to labor, which motivates firms to automate production. Our model borrows several features from previous models in the literature but presents a unique flexibility. In the literature, among the models that address automation, none (best of our knowledge) offers a qualitatively flexible response of employment to an automation-related shock. In these models, either employment always falls (e.g., Zeira, 1998; Acemoglu and Restrepo, 2018; Prettner and Strulik, 2019) or employment always increases (Guimarães and Gil, 2019). Our model, on the contrary, offers a framework in which automation has an ambiguous effect on employment, which is, in our view, an advantage as there is no consensus about how automation will affect aggregate employment. That is, our model allows us to anticipate different scenarios for the future of employment and the conditions in which economic policies (robot and automation tax) agree and disagree with their intended goals.

The narrative and assumptions of our model broadly agree with those in Acemoglu and Restrepo (2018). In our model, 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

[1] A growing empirical literature suggests that recent technologies with the purpose of automating production have increased employment by creating more jobs than destroying (see, e.g., Bessen, 2016, Autor and Salomons, 2018, and Gregory, Salomons and Zierahn, 2018; see also Bessen et al., 2020 for a review of this literature). A notable exception is Acemoglu and Restrepo (2019*b*), who find that robot adoption depresses employment and wages at the commuting-zone level. Yet, Acemoglu and Restrepo abstract from the indirect effects of robot adoption in one commuting zone on the other commuting zones that may render a positive effect of robot adoption at the aggregate level. Thus, in a way, Acemoglu and Restrepo abstract from the indirect positive effects of robot adoption on employment found by other studies (e.g., Autor and Salomons and Gregory, Salomons and Zierahn), which more than compensate for its job-displacing effects.

only machines/robots.[2]

Yet, to properly take into account the idiosyncrasies of the labor market, we fundamentally deviate from Acemoglu and Restrepo and build a model with matching frictions based on the Diamond-Mortensen-Pissarides setup. This allows us to realistically model the long-term firm-worker relationship and bring us closer to Hornstein, Krusell and Violante (2007) and to our previous work in Guimarães and Gil (2019). We, however, depart from our previous work by assuming that jobs are endogenously destroyed as firms continuously contrast their value using the manual technology and the option to move to the automated technology. In this sense, our model is closer to Hornstein, Krusell and Violante because they also endogenize job destruction. Yet, our model and focus also differ from theirs in important aspects. Hornstein, Krusell and Violante build a model with vintage capital to study capital-embodied technological change. We, on the other hand, consider the dichotomy of manual and automated technologies to study the consequences of AI and Robotics.

Our assumptions imply that a rise in the comparative advantage of machines affects employment by changing both job creation and job destruction. This is an important deviation from the literature that assumes flexible labor markets, which cannot offer insights regarding how the flows in the labor market react to shocks and determine employment fluctuations. And it is precisely this deviation from the literature that lends our model its flexibility regarding the response of employment to an *automation-augmenting shock* – a shock that increases the productivity of all machines/robots, and hence comparative advantage, relative to that of labor.

To simulate the effects of improvements in AI and Robotics on employment, we calibrate the automation-augmenting shock relying on the predictions in the literature about the future share of automated jobs. In all our simulations, job creation and job destruction increase.[3,4] Job destruction increases because it is more profitable to invest

[2]By allowing firms to choose whether to invest in the manual or in the automated technology, our model relates to a long literature of technology choice that we review more extensively in Guimarães and Gil (2019). In our model and in several contributions within this literature, the technology choice depends explicitly on a firm-specific (or task-specific) exogenous feature (e.g., Zeira, 1998, 2010; Acemoglu and Zilibotti, 2001; Acemoglu, 2003; Acemoglu and Restrepo, 2018; Alesina, Battisti and Zeira, 2018; and Guimarães and Gil, 2019). This feature then determines, *ceteris paribus*, the firm's overall productivity or cost level in using each technology.

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

[4]The increase in the job destruction rate after the automation-augmenting shock together with other features of our model seem to concur with the definition of Machine Learning – AI's state-of-the-art. Machine learning refers to algorithms allowing machines to *learn* from data by identifying data patterns (Naudé, 2020). This form of AI is, thus, not comparable to human intelligence and justifies our assumption that workers retain their comparative advantage in producing new and complex tasks. Furthermore, as data is collected, the comparative advantage of workers in producing current and future tasks should tend to vanish more rapidly than before, permanently increasing the job-destruction rate as predicted by our model. The effect of Machine Learning on employment is, however, ambiguous as new tasks and jobs emerge due to these technologies.

in the automated technology and so more firms destroy jobs and automate production. Job creation increases because of one or a combination of two mechanisms. First, as in Guimarães and Gil (2019), an automation-augmenting shock increases job creation if firms can choose technology at the time of entry after paying a sunk entry cost. In this scenario, a productivity effect emerges in general-equilibrium after the automation-augmenting shock: motivated by the increase in productivity of the automated technology, firm entry surges, which ultimately rises employment. 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 new tasks tend to be produced by workers, 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.[5]

Even though both flows increase if machines become more productive, their absolute and relative magnitudes crucially depend on the calibration of the model. In our benchmark calibration, we follow Acemoglu and Restrepo (2018) and mute general-equilibrium effects (as explained below); in this case, we find that job destruction increases more than job creation, lowering employment. But in other calibrations, the opposite occurs and employment rises. The relative magnitudes of the changes in the flows depend crucially on the parameter controlling entry-related general-equilibrium effects, i.e., the degree of technology constraints at the time of entry. If more firms are not technologically constrained and are able to choose technology at the time of entry, then employment falls less and may even rise as both mechanisms promoting job creation (discussed in the previous paragraph) are active. More surprisingly, the relative magnitudes of the changes in the flows also depend on the future share of automated jobs: an increase in this share does not monotonically dampen employment; indeed, a very pessimistic view of the future share of automated jobs coincides with more employment. The reason is that, in the model, a higher calibrated share of automated jobs requires a greater usage of machines/robots and, thus, a larger automation-augmenting shock. The latter, in turn, has a scale effect by increasing the productivity of all machines/robots. Put simply, a larger share of automated jobs amplifies the scale effect, benefiting job creation. The response of job creation is, then, convex in the future share of automated jobs, leading to an increase in employment when a sufficiently large share of future jobs are automated.

Our result that technology affects both job creation and job destruction flows echoes the analysis by Mortensen and Pissarides (1998), who study the relation between the rate of technological progress and employment in a model with capital-embodied technological change and matching frictions.[6] In Mortensen and Pissarides, a higher growth rate increases job destruction as wages grow faster due to rapidly-improving outside

[5]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.

[6]Their paper is a precursor of Hornstein, Krusell and Violante (2007).

options for workers; the effects of the growth rate on job creation are ambiguous, depending on the size of renovation costs (a cost that if paid allows firms to update their capital stock without laying-off the worker). They conclude that there is a threshold for the renovation cost above which employment falls with technological progress. Differently, we study the effects of advances in AI and Robotics, i.e., the increase in productivity of a technology that substitutes labor instead of complementing it as in Mortensen and Pissarides. But we also find that the calibration of the model affects the qualitative response of employment because of its effect on labor market flows.

Akin to Mortensen and Pissarides (1998), we find that the increase in wages after the automation-augmenting shock plays a very important role in shaping the response of employment.[7] In tighter labor markets (as observed in our model after the shock), workers demand higher wages for two reasons. One is that the outside option of manual firms of looking for an alternative worker is more costly and another is that the automation-augmenting shock increases jobs' continuation value. When we counterfactually assume that wages are orthogonal to labor market tightness (and to the productivity of the automated technology), 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 make it harder for a firm to find a worker suitable for the task, 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 relevance of human interactions. A simple case is the one of sellers and vending machines. Both broadly sell (they perform the same task) but consumers do not necessarily find the same task performed by one or the other perfect substitutes. In the scenario in which they are imperfect substitutes, a widespread use of machines increases the price of the tasks produced by workers relative to the price of the tasks produced by machines, 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.

In face of the automation threat, there is a growing literature assessing whether economic policy improves welfare and employment. For example, Jaimovich et al. (2020) assess retraining programs and redistributive transfers (e.g., increase in unemployment insurance and universal basic income), while Gasteiger and Prettner (2020) and Guerreiro, Rebelo and Teles (2020) assess the implications of a robot tax – a tax that is paid for buying a robot to produce a task. Furthermore, Prettner and Strulik (2019) study

[7]Importantly, the empirical literature also finds that an increase in robots leads to higher average wages; see, e.g., Autor and Salomons (2018) and Graetz and Michaels (2018). Yet, despite higher average wages after the automation-augmenting shock, workers performing tasks with less comparative advantage continue to earn relatively lower wages in our model. This agrees with the findings in Arnoud (2018) that workers in occupations with higher higher probability of automation earn lower wages.

the implications of redistribution policies financed by labor income taxes and robot taxes. We add to this literature by assessing how a robot tax and an automation tax – a tax that is paid for buying a robot to replace a worker in producing a task – affect employment and wages in our model. A robot and an automation tax are similar; the only difference is that all firms that buy a robot must pay a robot tax but only the firms that replace a worker with a robot must pay an automation tax. We find that both unambiguously lower wages but have an ambiguous effect on employment. The automation tax is, however, less pernicious on employment and wages as it does not directly affect the entry-related general equilibrium effects discussed above.

Our paper relates to Prettner and Strulik (2019), Basso and Jimeno (2018), Berg, Buffie and Zanna (2018), and Caselli and Manning (2019) (and again with Acemoglu and Restrepo, 2018) in that these papers also assess how automation-related shocks may affect either wages or employment in the future. Prettner and Strulik build a life-cycle model in which machines complement high-skill labor but substitute low-skill labor. They conclude that innovation asymptotically increases automation and inequality. And in an extension, they show that innovation always reduces low-skill employment due to greater automation and the high costs of acquiring skills for some workers. Basso and Jimeno assess the effect of demographical changes in a life-cycle model in which R&D investment may be directed to innovation (new tasks) or automation (of current tasks). They conclude that the demographic transition in the United States and Europe promoted higher wages in the beginning of 2000's but lower wages afterwards. Berg, Buffie and Zanna build a model with a nested CES (constant-elasticity of substitution) production function in which standard capital complements a composite of labor and robots; this composite assumes that labor and robots are substitutes. They conclude that robot-augmenting shocks can only benefit labor in the very long run. Caselli and Manning study how innovation affects real wages in economies with constant returns to scale, constant real interest rate, and multiple types of labor. They conclude that average wages increase as long as the price of capital falls more than that of consumption goods. Under this condition, they also conclude that all wages increase if the supply of labor types is perfectly elastic. But their model, as well as the models in Basso and Jimeno and Berg, Buffie and Zanna, abstracts from the impacts of shocks on employment as labor supply is assumed inelastic. More generally, our model differs from all these models because they assume perfectly competitive labor markets.

Our paper also relates to our previous paper, Guimarães and Gil (2019), and to Cords and Prettner (2019), Leduc and Liu (2019), and Jaimovich et al. (2020) as the three papers include models with matching frictions and automation. But there are important differences regarding the objects of study and models used. In Guimarães and Gil, we do not try to understand the future of employment and rather try to understand its past evolution. In particular, we study the evolution of the US economy from 1967 to 2007 and conclude that an acceleration in automation-augmenting shocks was an important driver of the fall in the US labor share after 1987. On the other hand, Cords and Prettner study how an increase in the stock of robots affects low- and high-skill employment; Leduc and Liu study the implications of automation for the business cycle; and

Jaimovich et al. focus on the welfare loss of workers in routine occupations caused by automation. In contrast with the four papers, in this paper, we assess the long run implications of expected advances in AI and Robotics on the employment rate, studying the conditions in which these innovations lead to higher and lower employment. Regarding the modeling strategy, our goal in Guimarães and Gil is to build a very stylized version of a model with matching frictions and automation that agrees with empirical studies suggesting that automation has increased employment in the past. Thus, and as mentioned above, we abstract from endogenous fluctuations in the job destruction rate. Cords and Prettner, Leduc and Liu, and Jaimovich et al. also abstract from endogenous changes in the job destruction rate, allowing only for exogenous fluctuations. This is in stark contrast with our current paper, in which endogenous job destruction is key for the ambiguity of the effects of automation-augmenting shocks on employment and to capture the consequences of advances in AI and robotics.

The remainder of this paper is organized as follows. We start by detailing our model in Section 2. In Section 3, we calibrate our model and study numerically how advances in AI and Robotics are expected to disturb future employment. In Section 4, we assess how the *human touch* affects the disruption in the labor market caused by technological innovations. In Section 5, we discuss how a robot tax and an automation tax affect employment and wages. In Section 6, we conclude.

2 The Model

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 $\frac{\kappa_L}{\mu(\theta)}$ to match with a worker and it must bargain wages with the worker.[8]

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.[9] 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 job-filling probability and the job-finding probability as, respectively, $\mu(\theta) \equiv \chi\theta^{-\eta}$ and $f(\theta) \equiv \chi\theta^{1-\eta}$.

[8] 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, Battisti and Zeira (2018), and Guimarães and Gil (2019).

[9] 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).

Each task has a stochastic idiosyncratic productivity, z, in the interval $[z_{min}, \bar{z}]$ according to a probability distribution function $G(z)$. 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.

Firms' technological choice depends on the task's idiosyncratic productivity, z. In Figure 1, we summarize the timeline of how z affects the distribution of firms between the technologies. In Acemoglu and Restrepo (2018), labor has the highest comparative advantage in producing new tasks because newly created tasks have the highest index. We assume a more general environment. Of the number of new tasks created each period, a proportion $1 - \lambda_e$ has the highest productivity, $\bar{z}$, 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 have their productivity drawn from the distribution $G(z)$ of productivity levels over the interval $[z_{min}, \bar{z}]$ and firms choose technology according to the present-discounted values of the technologies. Thus, λ_e is a measure of technology constraints at the time of entry. Producing tasks with higher z is more profitable if 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 technology and below which firms prefer the automated technology at the time of entry.[10]

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. If it becomes too low, manual firms prefer to destroy the job and automate the production of the task. This line of events further echoes the setting in Acemoglu and Restrepo (2018). In their model, tasks previously performed by labor can be automated as the tasks' (relative) productivity falls due to the expansion of the technological frontier over time and the implied gradual obsolescence of existing manual tasks. We also find a similar mechanism in the model of Hornstein, Krusell and Violante (2007). They build a model in which a unit of vintage capital is matched with a worker. As technology evolves, firms that use the oldest vintage of capital prefer to scrap their capital and, as in our model, destroy the job. Yet, in the models of both Acemoglu and Restrepo and Hornstein, Krusell and Violante, the fall in the task's idiosyncratic productivity (relative to the technology frontier) is de-

[10] If $\lambda_e > 0$, entry in the model is, at least, partially undirected, which is our assumption in Guimarães and Gil (2019). In that paper, we motivated this assumption by reviewing the literature on entrepreneurship and venturing. This literature emphasizes a (costly) learning stage about the market and menu of technologies prior to technology-choice and production; we capture this learning stage in our model as a productivity draw, z, after the payment of the sunk entry cost, Ω.

Figure 1: Timing of technological constrains and technology choice

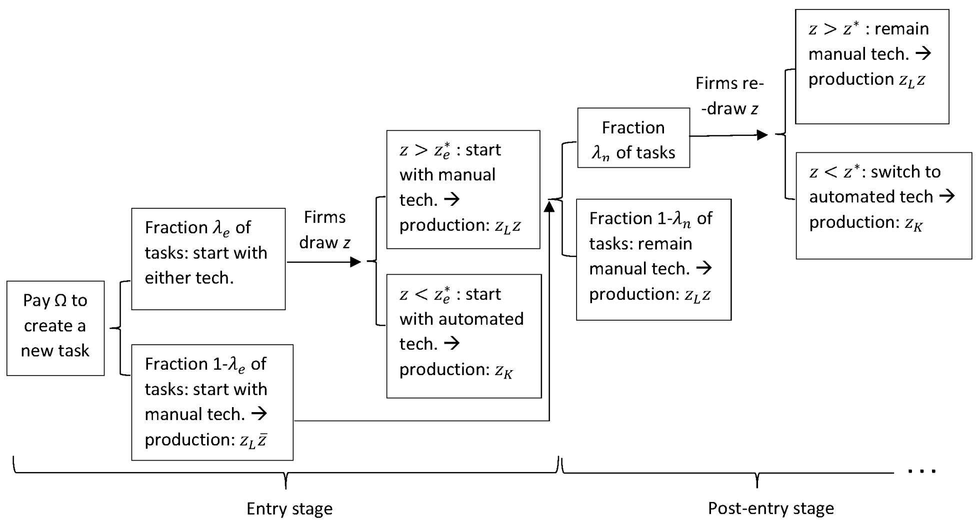

terministic while, in our model, we assume it to be stochastic.[11] To model the evolution of z, we build on Mortensen and Pissarides (1994). After production takes place, a proportion $1 - \lambda_n$ of manual firms sees no change in their tasks' idiosyncratic productivity and, thus, in their position relative to the technology frontier, $\bar{z}$. 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.[12]

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

[11] We assume it to be stochastic for two reasons. One is that it is a convenient assumption that does not demand us to keep track of how far or close a task is from being automated. The other, and more important, is that tasks may differ on the speed at which they are automated; thus, we find it more realistic to assume that the transition from manual to automated is random rather than deterministic.

[12] 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.

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.

2.1 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 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:

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

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^*:

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

2.2 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 wages change, 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

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

2.3 Wage Bargaining

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

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

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 $\beta J_L(z) - \frac{\kappa_L}{\mu(\theta)}$.[13] 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). Thus, we summarize the solution to Nash bargaining as

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

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

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. Eq. (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.

2.4 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.

2.4.1 Employment Rate and Number of Firms

We define employment as the number of workers employed at the time of production. As usual, 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

[13] 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.

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

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

Because every manual firm employs one worker, n also represents the number of manual firms. But the number of firms that use the automated technology is more intricate: some firms immediately choose the automated technology; others start with the manual technology and then move to the automated technology. We start by measuring the former. First, only a proportion λ_e of new firms can choose technologies. Second, if the firms can choose technology, they only choose the automated technology if the idiosyncratic productivity is below the cutoff z_e^*; this occurs with a probability $G(z_e^*)$. Third, the proportion of those that enter and choose the manual technology is $\lambda_e(1-G(z_e^*)) + 1 - \lambda_e$, which corresponds to the number of firms choosing the manual technology: $f(\theta)(1-n)$. Thus, every period, there is $\frac{\lambda_e G(z_e^*)}{\lambda_e(1-G(z_e^*))+1-\lambda_e} f(\theta)(1-n)$ firms that start production immediately using the automated technology.

Now we measure the other source of automated firms: those that start with the manual technology and change technology. To measure this, we must determine the number of firms that endogenously fire their workers every period. Given that there are n manual firms, there is a probability δ_L that the job is exogenously destroyed, there is a probability λ_n that the productivity of the task changes, and there is a probability $G(z^*)$ that a firm that redraws productivity moves to the automated technology, then the number of firms that automate the production of their respective tasks is $(1-\delta_L)\lambda_n G(z^*)n$.

Additionally, denoting n_K as the stock of automated firms, there are $\delta_K n_K$ automated firms destroyed every period. Thus, there are

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

automated firms.

2.4.2 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 $\frac{\lambda_e(1-G(z_e^*))}{\lambda_e(1-G(z_e^*))+1-\lambda_e} f(\theta)(1-n)$. This follows from two factors. First, every period, $f(\theta)(1-n)$ new manual firms are created. Second, these firms split between those that do not draw productivity (in proportion $1-\lambda_e$ of all new firms) and those that draw productivity and prefer the manual technology (in proportion $\lambda_e(1-G(z_e^*))$ of all new firms). Furthermore, the number of manual firms that redraw productivity and obtain z above z_e^* is $(1-\delta_L)\lambda_n(1-G(z_e^*))$ 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 $\delta_L+(1-\delta_L)\lambda_n G(z_e^*)$. Combining the flows into and out of n_e^* implies after a few derivations:

$$n_e^* = \frac{(1-\delta_L)\lambda_n(1-G(z_e^*))n}{\delta_L + (1-\delta_L)\lambda_n} + \frac{\frac{\lambda_e(1-G(z_e^*))f(\theta)(1-n)}{\lambda_e(1-G(z_e^*))+1-\lambda_e}}{\delta_L + (1-\delta_L)\lambda_n}. \tag{13}$$

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

$$n^* = \frac{(1-\delta_L)\lambda_n}{\delta_L + (1-\delta_L)\lambda_n}(G(z_e^*) - G(z^*))n. \tag{14}$$

Having established the number of firms, we quantify output as

$$y = n_K z_K + (n - n^* - n_e^*)z_L\bar{z} + n_e^* \frac{1}{1-G(z_e^*)} \int_{z_e^*}^{\bar{z}} z \mathrm{d}G(z) + n^* \frac{1}{G(z_e^*) - G(z^*)} \int_{z^*}^{z_e^*} z \mathrm{d}G(z), \tag{15}$$

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:

$$LS = \frac{(n - n^* - n_e^*)w(\bar{z}) + n_e^* \frac{1}{1-G(z_e^*)} \int_{z_e^*}^{\bar{z}} w(z)\mathrm{d}G(z) + n^* \frac{1}{G(z_e^*)-G(z^*)} \int_{z^*}^{z_e^*} w(z)\mathrm{d}G(z)}{y}. \tag{16}$$

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 the task's idiosyncratic productivity from a uniform distribution, i.e., $G(z) = \frac{z - z_{min}}{\bar{z} - z_{min}}$, with $z_{min} = 1$ and $\bar{z} = 1.92$; this calibration agrees with the average TFP ratio between an industry's 90th and 10th percentile plants (within four-digit SIC industries) in the US manufacturing sector (Syverson, 2004). To calibrate b, we assume it is 70% of the productivity of the firm that draws $z = z_{min} + \frac{\bar{z} - z_{min}}{2}$. 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 normalize z_K, κ_K, and κ_L to unity and set $\delta_K = 0.01$. The remaining five parameters, Ω, z_L, χ, δ_L, and λ_n, are calibrated such that our steady-state matches five targets. We target the prime-age (aged 25-54) workers' employment rate and the labor share in the US from 1977 until 2018;[14] 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 between 2002 and 2017 in the US. Based on Jaimovich and Siu (2020), we target $SAJ = 0.042$. Finally, we 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]

[14] We target the employment rate of prime-age workers because our model abstracts from demographic changes.

[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 our model abstracts from much of the churn in the labor market and focuses on the motivation to create jobs composed of new tasks (more broadly, occupations) and their automation.

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:	$z_{min} = 1$
Maximum productivity draw:	$\bar{z} = 1.92$
Nonemployment income:	$b = 0.7 z_L \left(z_{min} + \frac{\bar{z} + z_{min}}{2} \right)$
Rate of automated-firm destruction:	$\delta_K = 0.01$
Cost of capital/robot:	$\kappa_K = 1$
Job-filling cost:	$\kappa_L = 1$
Automation productivity (current):	$z_K = 1$
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:	$SAJ = 0.042$
Job destruction rate:	$JD = 0.016$
Future steady-state target	
Share of automated jobs within a decade:	$SAJ_{future} = 0.09$

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; Gregory, Salomons and Zierahn, 2018; and Bessen et al., 2020). 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 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. Therefore, as data is collected, Machine Learning should permanently facilitate machines catching up with labor in producing each task, increasing the pace of automation.[16]

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.[17] 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, Gregory and Zierahn (2017) use the estimates in Frey and Osborne 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, Gregory and Zierahn look at the various tasks composing each job while Frey and Osborne look only at the main occupation; 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 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, Gregory and Zierahn. 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 4.2% to 9% between our two steady-states.

3.3 The Future of Employment

Table 2 summarizes our main findings: the calibrated increase in z_K reduces employment in our benchmark but increases it in other calibrations, increases both job-finding and job-destruction probabilities, increases the wage in jobs with productivity $\bar{z}$ (all new jobs when $\lambda_e = 0$), and lowers the labor share.

Automation-augmenting shocks in the economy affect employment through changes in both job creation and job destruction. The change in job destruction is imposed by

[16]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. A potential criticism of both our approaches is that future jobs might demand new skills that are not as easily automated, restoring the comparative advantage of workers for a longer period. In our model, this could be captured by a combination of higher total factor productivity and a fall in z_K, implying that the productivity of the automated technology grows less than the aggregate. This would generate the opposite effects that we report in this paper.

[17]Although it would also be interesting to study the transition between steady-states, it is beyond the scope of our paper. As, for example, Jaimovich et al. (2020) and others, we instead focus on the differences between steady-states, which, in our case, are caused by improvements in AI and Robotics.

[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.

Table 2: Main results

	Δn	$\Delta f(\theta)$	ΔJD	ΔLS	$\Delta w(\bar{z})$
Benchmark	-0.36	0.79	2.45	-0.93	0.11
$\lambda_e = 0.2$	0.21	3.42	2.45	-1.61	0.34
Pareto	-0.47	0.28	2.45	-0.93	0.05
$\bar{z} = 2.5z_{min}$	-0.45	0.39	2.45	-0.91	0.06
$\eta = 0.4$	-0.32	0.98	2.45	-0.93	0.11
$\kappa_K = 0.5$	-0.36	0.79	2.45	-0.93	0.11
$\kappa_L = 0.5$	-0.37	0.77	2.45	-0.94	0.08
$\phi = 0.4$	-0.38	0.69	2.45	-0.94	0.08
$SAJ = 0.06$	-0.22	0.54	1.52	-0.58	0.08

Note: This table shows the effects of an automation-augmenting shock calibrated such that $SAJ_{future} = 0.09$. The five columns show, respectively, the percentage change in employment, job-finding probability, job-destruction probability, labor share, and wages (in firms with draw $z = \bar{z}$). 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 3.12.

our SAJ_{future} target, but the change in job creation depends heavily on the calibration of the model. Thus, we must understand the latter to understand the change in employment.[19] 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 the additional surplus with workers. Yet, employment falls in our benchmark because the change in job destruction exceeds that of job creation.

Looking at the various robustness checks included in Table 2, we conclude that the results are essentially robust. The only striking change occurs when we deviate from Acemoglu and Restrepo (2018) and assume that part of the new tasks are not technologically constrained and may start as automated depending on the draw of z ($\lambda_e = 0.2$). In this case, employment rises because of a significant surge in job creation. Furthermore, wages rise substantially more than in the benchmark. The reason for this result

[19]We note that absent the rise in $f(\theta)$, employment would fall 0.54%. This is not as substantial as a doubling of SAJ might suggest because automation is only a part of the reason for jobs to be destroyed. There are other significant exogenous reasons for job separation, including factors like shifts in tastes or in technology unrelated with automation.

is 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. But an automation-augmenting shock also increases the expected value of a firm, which incentivizes firm entry.[20] 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.

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, Gregory and Zierahn (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, Figure 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.

This result is 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. To understand these results, we have to take into account 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^*. Then, 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

[20]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).

Figure 2: Robustness checks to different SAJ_{future}

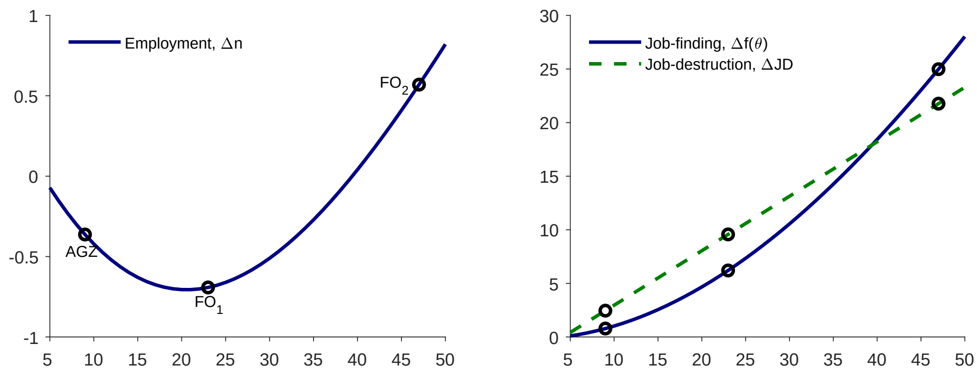

Note: This figure shows the effects of an automation-augmenting shock calibrated such that our model reaches the different targets of SAJ_{future} in the horizontal axis (in percentage terms). The left panel shows the percentage change in employment. The right panel shows the percentage change in the 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, Gregory and Zierahn (2017). The other two (marked with FO) correspond to the lower and upper bounds of the estimates in Frey and Osborne (2017), respectively. See the discussion in Section 3.2 for more details.

z^*), amplifying its effect on job creation and giving rise to a convex change in the job-finding probability. Thus, the (negative) elasticity of n with respect to z_K gets smaller (in modulus) with the rise in z_K, changing sign for a sufficiently large z_K.

3.5 Dissecting the Mechanism: Ad hoc Function for Wages

Table 2 shows that 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. 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: $w(z) = (1 - \phi_{nb})b + \phi_{nb} z_L z$ $(0 < \phi_{nb} < 1)$. 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.

Table 3 contrasts the effect of the rise in the share of automated jobs in the model with Nash bargaining and in the model with ad hoc wages under two calibrations of η.[21] 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 open more vacancies, creating more jobs. This, in turn, leads to a positive change in employment in almost all our calibrations. Put differently, if wages are orthogonal to z_K and θ, firms have a much greater incentive to hire a worker as their operational profits remain unchanged.

Table 3: Nash bargaining vs ad hoc wage

	A: Baseline (Nash)			B: Ad hoc			C: Ad hoc ($\eta = 0.3$)		
	Δn	$\Delta f(\theta)$	ΔJD	Δn	$\Delta f(\theta)$	ΔJD	Δn	$\Delta f(\theta)$	ΔJD
Benchmark	-0.36	0.79	2.45	0.35	4.09	2.45	1.49	9.79	2.45
$\lambda_e = 0.2$	0.21	3.42	2.45	3.03	18.28	2.45	7.26	47.95	2.45
Pareto	-0.47	0.28	2.45	-0.21	1.47	2.45	0.22	3.46	2.45
$\bar{z} = 2.5z_{min}$	-0.45	0.39	2.45	-0.09	2.05	2.45	0.50	4.85	2.45
$\kappa_K = 0.5$	-0.36	0.79	2.45	0.35	4.09	2.45	1.50	9.81	2.45
$\kappa_L = 0.5$	-0.37	0.77	2.45	0.54	5.00	2.45	1.92	12.06	2.45
$\phi = 0.4$	-0.38	0.69	2.45	0.09	2.88	2.45	0.91	6.86	2.45
$SAJ = 0.06$	-0.22	0.54	1.52	0.27	2.79	1.52	1.07	6.64	1.52

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 $\eta = 0.3$. See the note to Table 2 for more details.

Yet, employment still falls under some calibrations of the model. We conjecture that matching frictions are behind this result. As labor market tightness, θ, increases, the costs of a firm to match with a worker also increase, reducing incentives for job creation. We can test this conjecture by checking how our results change with different calibrations of the 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 ($\frac{\kappa_L}{\mu(\theta)} = \frac{\kappa_L \theta^\eta}{\chi}$). Thus, matching frictions are less relevant for job creation, which should magnify the response of the latter to disturbances. Panel C in Table 3 suggests that our conjecture that matching frictions prevent a greater increase in employment is reasonable. This panel shows that job creation increases much more if $\eta = 0.3$ than if $\eta = 0.5$ (when wages are ad hoc), implying a positive change in employment after the increase in z_K

[21] 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 z^* is the same in the models with and without Nash bargaining. Ω, z_L, χ, δ_L, and λ_n, are calibrated such that our steady-state matches the five current steady-state targets as reported in Table 1.

necessary to reach our target for SAJ_{future}. Therefore, in our model, the reduction in employment in many of our calibrations is explained by a combination of higher wages and 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 it. The worker (seller) can offer a more personal (*human*) touch to the task whereas the machine (vending machine) offers an impersonal service.[22] 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 automated and manual technologies, where y is an index of final consumption (i.e., a bundle of goods and services demanded by consumers). In this setup, the elasticity of substitution is ϵ, and this model nests our baseline model if $\epsilon = \infty$. In particular, the CES takes the following form:

$$y = \left[y_K^{\frac{\epsilon-1}{\epsilon}} + y_L^{\frac{\epsilon-1}{\epsilon}} \right]^{\frac{\epsilon}{\epsilon-1}}, \tag{17}$$

[22]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).

where y_K and y_L are the sum of the outputs produced using each type of technology:

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

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

$$p_K = y_K^{-\frac{1}{\epsilon}} y^{\frac{1}{\epsilon}}, \tag{18}$$
$$p_L = y_L^{-\frac{1}{\epsilon}} y^{\frac{1}{\epsilon}}. \tag{19}$$

Thus, a rise in z_K leads to an increase in y_K, which reduces the price of the tasks produced using the automated technology. Furthermore, it also leads to a rise in y, which converts into a higher price of the tasks produced using the manual technology. These two effects clearly affect the motivation to create jobs as well as to fire workers and automate the production of tasks (destroy jobs).

Table 4 contrasts the effect of the rise in the share of automated jobs in the baseline model with that in the model with the CES assuming $\epsilon = 4$. We find that assuming that the outputs of the two technologies are imperfect substitutes qualitatively changes our results as, in the benchmark and in almost all our robustness checks, employment increases. As hinted above, endogenous price movements explain this result. A rise in z_K increases the production of the automated technology both directly and indirectly by steering resources away from the manual technology towards the automated technology (both cutoffs, z^* and z_e^*, rise). This, in turn, leads to increased prevalence of the automated technology, which reduces the price of the tasks produced by it, p_K, and increases that of the tasks produced using the manual technology, p_L. Therefore, workers effectively become more productive, which reinforces the incentives to create jobs.

Imperfect substitutability per se is not enough to qualitatively change our results; ϵ, controlling the degree of substitutatability between the two technologies, is critical both qualitatively and quantitatively as shown in Figure 3. In the limit, if $\epsilon = \infty$, the model with the CES converges to the baseline model and employment falls in our benchmark. Yet, 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

Table 4: The effect of the *Human Touch* (CES)

	A: Baseline			B: CES ($\epsilon = 4$)		
	Δn	$\Delta f(\theta)$	ΔJD	Δn	$\Delta f(\theta)$	ΔJD
Benchmark	-0.36	0.79	2.45	0.12	3.02	2.45
$\lambda_e = 0.2$	0.21	3.42	2.45	0.76	6.09	2.45
Pareto	-0.47	0.28	2.45	-0.07	2.12	2.45
$\bar{z} = 2.5z_{min}$	-0.45	0.39	2.45	0.01	2.52	2.45
$\eta = 0.4$	-0.32	0.98	2.45	0.25	3.63	2.45
$\kappa_K = 0.5$	-0.36	0.79	2.45	0.12	3.02	2.45
$\kappa_L = 0.5$	-0.37	0.77	2.45	0.44	4.56	2.45
$\phi = 0.4$	-0.38	0.69	2.45	0.29	3.84	2.45
$SAJ = 0.06$	-0.22	0.54	1.52	0.05	1.75	1.52

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 4. Panel B, thus, corresponds to the case with *Human Touch*. See the note to Table 2 for more details.

Figure 3: The effect of the *Human Touch* (CES) under different values of ϵ

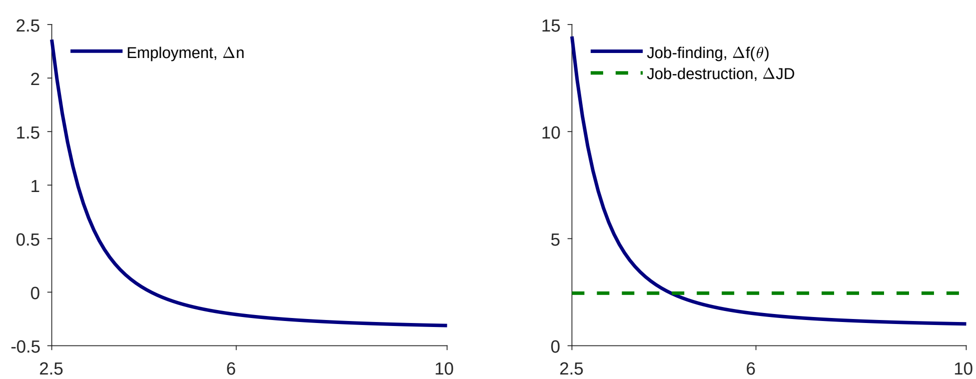

Note: This figure shows the effects of an automation-augmenting shock calibrated such that our model reaches our target of $SAJ_{future} = 0.09$ under different values of ϵ in the horizontal axis. The left panel shows the percentage change in employment. The right panel shows the percentage change in the job-finding probability (solid line) and the job-destruction probability (dashed line).

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 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, 2019; Gasteiger and Prettner, 2020; and Guerreiro, Rebelo and Teles, 2020). In this section, we contribute to the discussion by studying how a robot tax affects employment and wages in our model. We start by recalibrating z_K to reach 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., $\tau = 0.1J_K$, to be paid upfront in addition to the cost of capital/robot, κ_K.

The results are depicted in Figure 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 Vs Robot Tax

An alternative tax that might be implemented by policymakers is what we call *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. Hence, 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 the robot at the time of entry without the need to invest in the manual technology before; these firms do not pay the automation tax but must

Figure 4: Robot Tax for different values of SAJ_{future}

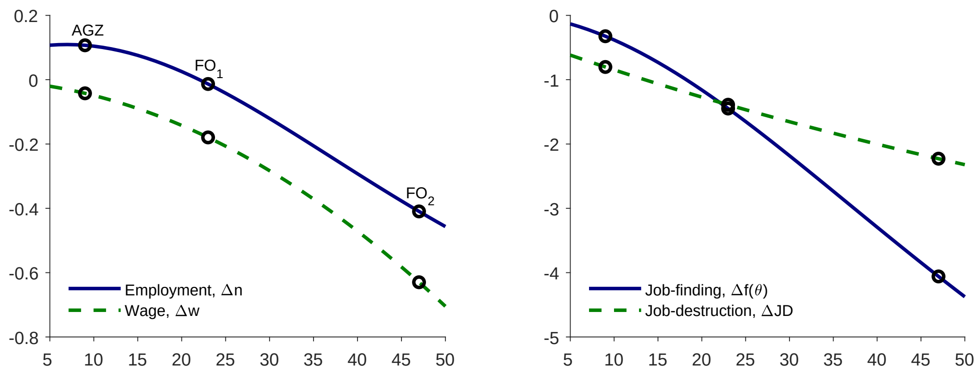

Note: This figure shows the effects of a robot tax. The horizontal axis corresponds to the SAJ_{future} target (in percentage terms). The left panel shows the percentage change in employment (solid line) and in the wage in firms with $z = \bar{z}$ (dashed line). The right panel shows the percentage change 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. See the note to Figure 2 for more details.

pay the robot tax.

To understand the distinction between an automation tax and a robot tax, Figure 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, $\tau = 0.1J_K$, 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 Figure 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 tax only lowers employment if $\lambda_e \geq 0.25$. 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. In sum, our results suggest that an automation tax should be preferred to a robot tax.

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; Gregory, Salomons and Zierahn, 2018; and Bessen et al., 2020). Yet, AI and Robotics are expected to disrupt the labor

Figure 5: Automation Tax Vs Robot Tax

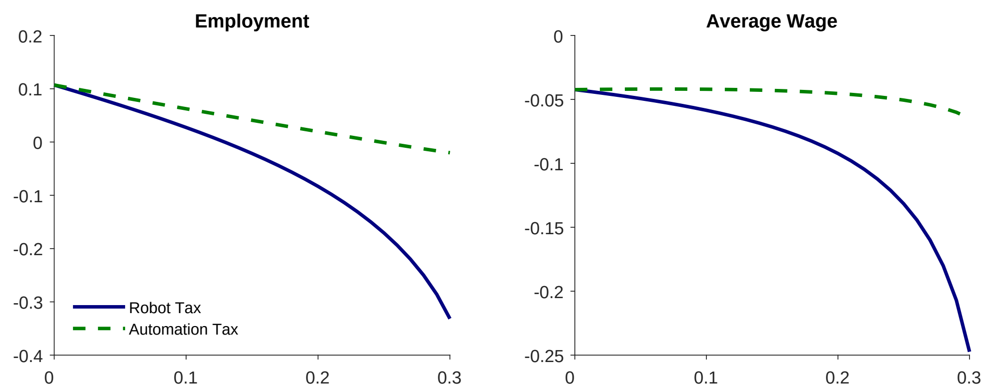

Note: The left (right) panel shows the percentage change in employment (wage when $z = \bar{z}$) when a robot tax is implemented (solid line) and when an automation tax is implemented (dashed line). The horizontal axis corresponds to λ_e.

market by much more than previous automating technologies (e.g., Brynjolfsson and McAfee, 2011; Brynjolfsson, Rock and Syverson, 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, 2019) 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.

Yet, 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). This documented trend is even more relevant given that the fall in labor market flows occurred in a period of increased automation and investment in robots (Prettner and Strulik, 2019; Acemoglu and Restrepo, 2019*a*, Guimarães and Gil, 2019). The downward trend in labor market fluidity, however, seems mostly driven by composition effects that our model abstracts from. Hyatt and Spletzer (2017) document that about half of the decline in hires and separations is accounted for by a significant drop of the prevalence of jobs that start and end in the same quarter; 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.[23] Furthermore, even though aggregate labor market flows fell in all sectors, they fell unevenly across them. Particu-

[23]Looking at their Figure 5, we can see that job finding rates dropped in the final years in their data (2008-14) but that drop was most likely caused by the Great Recession.

larly, Decker et al. (2014) document that labor market flows fell much more in retail and services sectors than in finance and manufacturing sectors – the sectors that arguably were more susceptible to automation. Finally, other evidence points in the direction of increasing flows, concurring with our model. Kambourov and Manovskii (2008) document a positive trend in occupational mobility, which might be related with automation forcing workers to adapt to new tasks and occupations; and Bloom, McKenna and Prettner (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.

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.[24] 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, Rock and Syverson, 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

[24]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. (2020) on the consequences of the disappearance of routine occupations in the last decades in the US.

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, Margo and Rhode, 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 losses in employment. We will consider the consequences of this point in a future paper.

References

Acemoglu, Daron. 2003. “Labor- and Capital-Augmenting Technical Change.” *Journal of European Economic Association*, 1: 1–37.

Acemoglu, Daron, and Fabrizio Zilibotti. 2001. “Productivity Differences.” *Quarterly Journal of Economics*, 116 (2): 563–606.

Acemoglu, Daron, and Pascual Restrepo. 2018. “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment.” *American Economic Review*, 108(6): 1488–1542.

Acemoglu, Daron, and Pascual Restrepo. 2019*a*. “Automation and New Tasks: How Technology Displaces and Reinstates Labor.” *Journal of Economic Perspectives*, 33(2): 3–30.

Acemoglu, Daron, and Pascual Restrepo. 2019*b*. “Robots and Jobs: Evidence from US Labor Markets.” *Journal of Political Economy*. Forthcoming.

Aghion, Philippe, Antonin Bergeaud, Richard W Blundell, and Rachel Griffith. 2019. “The Innovation Premium to Soft Skills in Low-Skilled Occupations.” *Available at SSRN 3489777*.

Alesina, A., M. Battisti, and J. Zeira. 2018. “Technology and Labor Regulations: Theory and Evidence.” *Journal of Economic Growth*, 23(1): 41–78.

Arnoud, Antoine. 2018. “Automation Threat and Wage Bargaining.” *Unpublished Manuscript, Yale University*.

Arntz, Melanie, Terry Gregory, and Ulrich Zierahn. 2017. “Revisiting the Risk of Automation.” *Economics Letters*, 159: 157 – 160.

Atack, Jeremy, Robert A. Margo, and Paul W. Rhode. 2019. "”Automation” of Manufacturing in the Late Nineteenth Century: The Hand and Machine Labor Study." *Journal of Economic Perspectives*, 33(2): 51–70.

Autor, David, and Anna Salomons. 2018. "Is Automation Labor Share–Displacing? Productivity Growth, Employment, and the Labor Share." *Brookings Papers on Economic Activity*, 1–63.

Autor, David H. 2015. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives*, 29(3): 3–30.

Basso, H, and Juan F Jimeno. 2018. "From Secular Stagnation to Robocalypse? Implications of Demographic and Technological Changes." mimeo.

Berg, Andrew, Edward F. Buffie, and Luis-Felipe Zanna. 2018. "Should we Fear the Robot Revolution? (The correct answer is yes)." *Journal of Monetary Economics*, 97: 117 – 148.

Bessen, James E. 2016. "How Computer Automation Affects Occupations: Technology, Jobs, and Skills." *Boston Univ. school of law, law and economics research paper*, , (15-49).

Bessen, James, Maarten Goos, Anna Salomons, and Wiljan van den Berge. 2020. "Automation: A Guide for Policymakers."

Bloom, David E, Mathew McKenna, and Klaus Prettner. 2018. "Demography, Unemployment, Automation, and Digitalization: Implications for the Creation of (Decent) Jobs, 2010–2030." National Bureau of Economic Research Working Paper 24835.

Brynjolfsson, Erik, and Andrew McAfee. 2011. *Race Against the Machine*. Digital Frontier, Lexington, MA.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2017. "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics." National Bureau of Economic Research Working Paper 24001.

Caselli, Francesco, and Alan Manning. 2019. "Robot Arithmetic: New Technology and Wages." *American Economic Review: Insights*, 1(1): 1–12.

Coles, Melvyn G., and Ali Moghaddasi Kelishomi. 2018. "Do Job Destruction Shocks Matter in the Theory of Unemployment?" *American Economic Journal: Macroeconomics*, 10(3): 118–36.

Cords, Dario, and Klaus Prettner. 2019. "Technological Unemployment Revisited: Automation in a Search and Matching Framework."

Davis, Steven J, and John Haltiwanger. 2014. "Labor Market Fluidity and Economic Performance." National Bureau of Economic Research Working Paper 20479.

Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda. 2014. “The Secular Decline in Business Dynamism in the US.” *Unpublished draft, University of Maryland.*

Feigenbaum, James, and Daniel P Gross. 2020. “Automation and the Fate of Young Workers: Evidence from Telephone Operation in the Early 20th Century.” National Bureau of Economic Research Working Paper 28061.

Frey, Carl Benedikt, and Michael A. Osborne. 2017. “The Future of Employment: How Susceptible are Jobs to Computerisation?” *Technological Forecasting and Social Change*, 114: 254 – 280.

Gasteiger, Emanuel, and Klaus Prettner. 2020. “Automation, Stagnation, and the Implications of a Robot Tax.” ECON WPS 02/2020, Vienna:Vienna University of Technology, Institute of Mathematical Methods in Economics.

Graetz, Georg, and Guy Michaels. 2018. “Robots at Work.” *The Review of Economics and Statistics*, 100(5): 753–768.

Gregory, Terry, Anna Salomons, and Ulrich Zierahn. 2018. “Racing With or Against the Machine? Evidence from Europe.”

Guerreiro, Joao, Sergio Rebelo, and Pedro Teles. 2020. “Should Robots be Taxed?” National Bureau of Economic Research Working Paper 23806.

Guimarães, Luís, and Pedro Gil. 2019. “Explaining the Labor Share: Automation vs Labor Market Institutions.” University Library of Munich, Germany MPRA Paper 94236.

Hall, Robert E, and Marianna Kudlyak. 2019. “Job-Finding and Job-Losing: A Comprehensive Model of Heterogeneous Individual Labor-Market Dynamics.” National Bureau of Economic Research Working Paper 25625.

Hall, Robert E., and Paul R. Milgrom. 2008. “The Limited Influence of Unemployment on the Wage Bargain.” *American Economic Review*, 98(4): 1653–1674.

Hornstein, Andreas, Per Krusell, and Giovanni L. Violante. 2007. “Technology-Policy Interaction in Frictional Labour-Markets.” *The Review of Economic Studies*, 74(4): 1089–1124.

Hyatt, Henry R., and James R. Spletzer. 2017. “The Recent Decline of Single Quarter Jobs.” *Labour Economics*, 46: 166 – 176.

Jaimovich, Nir, and Henry E. Siu. 2020. “Job Polarization and Jobless Recoveries.” *The Review of Economics and Statistics*, 102(1): 129–147.

Jaimovich, Nir, Itay Saporta-Eksten, Henry E Siu, and Yaniv Yedid-Levi. 2020. “The Macroeconomics of Automation: Data, Theory, and Policy Analysis.” National Bureau of Economic Research Working Paper 27122.

Kambourov, Gueorgui, and Iourii Manovskii. 2008. “Rising Occupational and Industry Mobility in the United States: 1968-97.” *International Economic Review*, 49(1): 41–79.

Leduc, Sylvain, and Zheng Liu. 2019. “Robots or Workers? A Macro Analysis of Automation and Labor Markets.”

Martellini, Paolo, and Guido Menzio. 2020. “Declining Search Frictions, Unemployment, and Growth.” *Journal of Political Economy*, 128(12). Forthcoming.

Molloy, Raven, Riccardo Trezzi, Christopher L Smith, and Abigail Wozniak. 2016. “Understanding Declining Fluidity in the US Labor Market.” *Brookings Papers on Economic Activity*, 2016(1): 183–259.

Mortensen, Dale T., and Christopher A. Pissarides. 1998. “Technological Progress, Job Creation, and Job Destruction.” *Review of Economic Dynamics*, 1(4): 733 – 753.

Mortensen, Dale T., and Christopher Pissarides. 1994. “Job Creation and Job Destruction in the Theory of Unemployment.” *Review of Economic Studies*, 61(0): 397–415.

Naudé, Wim. 2020. “Artificial Intelligence: Neither Utopian Nor Apocalyptic Impacts Soon.” *Economics of Innovation and New Technology*, 0(0): 1–23.

Petrongolo, Barbara, and Christopher A. Pissarides. 2001. “Looking into the Black Box: A Survey of the Matching Function.” *Journal of Economic Literature*, 39(2): 390–431.

Pissarides, Christopher A. 2009. “The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?” *Econometrica*.

Prettner, Klaus, and Holger Strulik. 2019. “Innovation, Automation, and Inequality: Policy Challenges in the Race Against the Machine.” *Journal of Monetary Economics*.

Shimer, Robert. 2012. “Reassessing the Ins and Outs of Unemployment.” *Review of Economic Dynamics*, 15(2): 127–148.

Syverson, Chad. 2004. “Market Structure and Productivity: A Concrete Example.” *Journal of Political Economy*, 112(6): 1181–1222.

Zeira, Joseph. 1998. “Workers, Machines and Economic Growth.” *Quarterly Journal of Economics*, 113: 1091–1113.

Zeira, Joseph. 2010. “Machines as Engines of Growth.” mimeo.