LOOKING AHEAD AT THE EFFECTS OF AUTOMATION IN AN ECONOMY WITH MATCHING FRICTIONS *

LUÍS GUIMARÃES† & PEDRO MAZEDA GIL‡

September 13, 2020

Abstract

We study the effects of an automation-augmenting shock in an economy with matching frictions and endogenous job destruction. In the model, tasks can be produced by workers or by machines but workers have a comparative advantage in producing advanced tasks. Firms choose the input at the time of entry. And according to the evolution of the workers' comparative advantage, some firms using labor prefer to fire the worker and automate the task. In our model, an automation-augmenting shock reduces the labor share, increases job creation, and increases job destruction. The effects on employment depend on how rapidly workers may lose their comparative advantage: an automation-augmenting shock increases employment in slow-changing environments but reduces it (possibly catastrophically) in rapid-changing ones.

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 UID/ECO/04105/2019 (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.

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

1 Introduction

In the last five decades, total hours worked and employment rose in developed countries, despite the ubiquitous fall in the labor share. This employment growth looks staggering as it coexisted with the emergence of new technologies that automate production and are supposed to displace labor. But the growing empirical literature suggests that these new technologies – namely, automation – have actually favored employment growth by creating more jobs than they have destroyed.[1] In this paper, we ask: will automation always create more jobs than it destroys or can we expect a different future?

To answer this question, we build a theoretical model that satisfies two criteria. First, in order to be consistent with the past, an automation-augmenting shock – a shock that increases the productivity of all machines/robots – is able to reduce the labor share and simultaneously increase employment. And, second, in order to be insightful about how the future may differ from the past, the model is flexible enough to generate different outcomes from the same sort of shocks. In the literature, among the models that explain the fall in the labor share, none offers a qualitatively flexible response of employment. In these models, either employment always falls (Caballero and Hammour, 1998; Zeira, 1998; Hornstein, Krusell and Violante, 2007; Acemoglu and Restrepo, 2018; Prettner and Strulik, 2019) or employment always increases (Guimarães and Gil, 2019).[2] Our model borrows several features from these models to offer a framework that is consistent with the past and insightful about potential future scenarios. In our model, an automation-augmenting shock reduces the labor share but its effect on employment is ambiguous.

[1]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, 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.

[2]These models do not propose the same mechanism or shock to explain the fall in the labor share. But irrespective of the mechanism, they predict robust directions for employment after the shock that reduces the labor share.

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 only machines/robots.[3]

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).[4] 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.[5] 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 automation-augmenting shocks.

[3]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 using each technology.

[4]In this regard, our paper is also close to Cords and Prettner (2019) as they also build a model with matching frictions but to study how an increase in the stock of robots affects low- and high-skill employment.

[5]To model endogenous job destruction, we particularly rely on Mortensen and Pissarides (1994).

Our assumptions imply that automation-augmenting shocks affect 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 impact of automation-augmenting shocks on employment.

In all our calibrations, job creation and job destruction increase after an automation-augmenting shock. Job destruction increases because the shock makes it more profitable to invest 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, an automation-augmenting shock gives rise to a productivity effect in general-equilibrium: 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.[6]

Even though both flows increase after an automation-augmenting shock, their absolute and relative magnitudes crucially depend on the calibration of the model. In some calibrations, job creation increases more than job destruction, thereby raising employment. In other calibrations, the opposite occurs and employment falls. The

[6]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.

relative magnitudes of the changes in the flows depend crucially on one parameter, which we interpret as a feature intrinsic to each task controlling for how rapidly workers may lose their comparative advantage in producing it. In *slow-changing* environments, in which the comparative advantage of labor in producing each task is relatively stable, job destruction barely shifts after the automation-augmenting shock. In these conditions, job creation increases more than job destruction. Nonetheless, in *rapid-changing* environments, an automation-augmenting shock leads to massive job destruction. This jump in job destruction is not followed by an equal jump in job creation because the increase in labor market tightness makes it more costly to find the right worker and allows workers to demand higher wages. In these scenarios, employment catastrophically drops. These results show how our model can both agree and disagree with the facts documented by the empirical literature on the effects of automation. Thus, our paper conveys an important message: if current and future jobs are made of tasks in which workers rapidly lose their comparative advantage, then automation-augmenting shocks may have dramatically different consequences in the future. This will likely be the case if artificial intelligence techniques allows machines and software to rapidly adjust to new tasks.

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.[7] In Mortensen and Pissarides, a higher growth rate increases job destruction as wages grow faster due to rapidly-improving outside 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. In our case, we study the effects of automation-augmenting shocks, i.e., the increase in productivity refers to a technology that substitutes labor instead of complementing it

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

as in Mortensen and Pissarides. But we also find a threshold (in our case for the pace at which workers lose comparative advantage) above which employment drops with the productivity of the automated technology because of its distinctive effects on the two 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.[8] 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 workers can easily find other jobs. 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 quite rapid-changing environments because matching frictions also play their role. If job creation increases, it becomes harder to find a worker suitable for the job, which increases costs and discourages further job creation. Job destruction, on the other hand, is not much affected by matching frictions and increases significantly in quite rapid-changing environments, leading to the net fall in employment.

We consider two other variants of our model to further dissect the mechanism. In one variant, we deviate from the typical assumption in models with matching frictions that workers must stay nonemployed for at least a period after losing their jobs. This reduces the prevalence of matching frictions and increases the pool of available workers for firms investing in the manual technology. We find that relaxing this assumption does promote greater employment but we also find that it does not have much quantitative impact.

[8] 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.

In another variant, 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 the human touch. 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 reduces job destruction but barely changes job creation. 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 unlikely to catastrophically reduce employment.

Our paper also 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 naturally relates to our previous paper, Guimarães and Gil (2019), and to Leduc and Liu (2019), as both 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 how the future may differ from the past but rather try to understand the past. 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. Leduc and Liu (2019), on the other hand, study the implications of automation for the business cycle and show that accounting for automation is important to match business cycle fluctuations in key labor-market variables. In contrast with the two, in this paper, we study the long run implications of automation-augmenting shocks for the employment rate, studying the conditions in which these shocks 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. Leduc and Liu (2019) 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.

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 the effects of automation-augmenting shocks. In Section 4, we dissect the mechanisms underlying our results, including the role of the *human touch*. In Section 5, 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.[9]

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.[10] 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 $f(\theta) \equiv \chi\theta^{1-\eta}$.

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

[9]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).

[10]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).

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. 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.[11]

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,

[11] 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

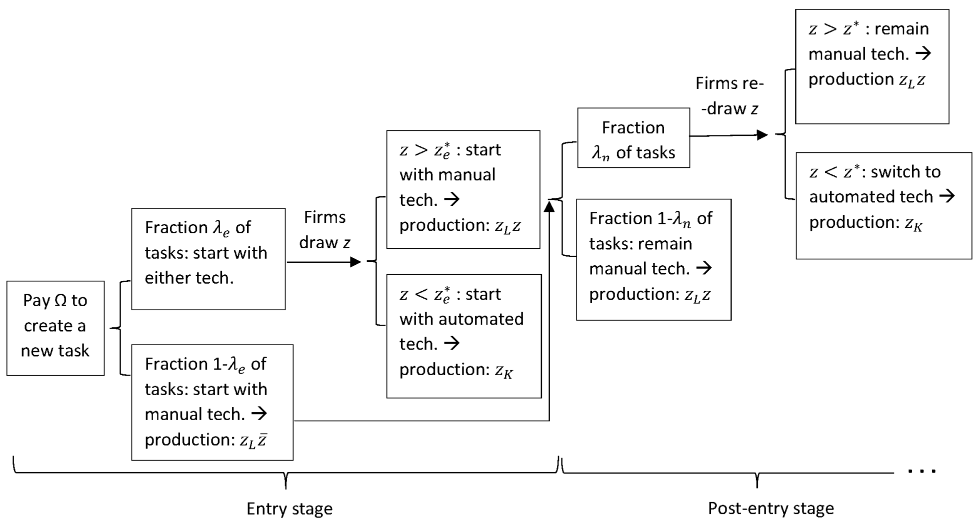

tasks previously performed by labor can be automated as the tasks' (relative) productivity falls due to the expansion of the technological frontier over time and the implied gradual obsolescence of existing manual tasks. We also find a similar mechanism in the model of Hornstein, Krusell and Violante (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 deterministic while, in our model, we assume it to be stochastic.[12] 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

[12]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.

$G(z)$ 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.[13] As a result, λ_n controls for how rapidly workers may lose their comparative advantage, which directly affects job destruction.

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

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

2.1 Firms

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

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

[13]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.

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 $\beta J_L(z) - \frac{\kappa_L}{\mu(\theta)}$.[14] 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.

Given Nash bargaining, job destruction only occurs when the surplus of the match is negative; thus, both workers and firms deem it optimal to destroy the job. The surplus of the match is only negative if it is less profitable for the firm to stay in the manual technology than to move to the automated technology, which occurs when $J_L(z) <$

[14] Importantly, 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.

$\beta J_K - \kappa_K$. In other words, all firms that draw the task's idiosyncratic productivity below the cutoff z^*, fire the worker and move to the automated technology. Simultaneously, when the task's idiosyncratic productivity is too low, workers prefer to move to nonemployment than to stay employed and earn a low wage because $E(z) < U$. Thus, the cutoff z^* satisfies $E(z^*) = U$ or, equivalently, Eq. (2).

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 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

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

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

Looking into the last four decades, recent empirical studies on the effects of new technologies (automation, industrial robots, artificial intelligence, and routine-replacing innovations) point to a net increase in employment (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 the future be different?

Our goals are mainly conceptual: we want to assess (i) the conditions under which an automation-augmenting shock –an increase in z_K– increases and decreases employment and (ii) the conditions that magnify the response of employment to the shock. Yet, because we are not able to obtain analytical results, our approach is to calibrate the model and assess the effects of an automation-augmenting shock under various calibrations. We conclude that our results essentially depend on one key parameter:

λ_n, which determines the frequency at which the productivity of a task is redrawn and, thus, indirectly controls for how rapidly a task can be automated.[15] If it is low (about 0.02; meaning that the productivity of a task changes on average every four years), employment tends to increase; otherwise, it tends to fall. The remaining parameters either magnify the response of employment (changing the λ_n threshold slightly) or have negligible effects.

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.*, $G(z) = \frac{z - z_{min}}{\bar{z} - z_{min}}$, in which $\bar{z} = 0.25$ and $z_{min} = 0.15$.[16] 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.

To calibrate the exogenous probability of manual firm destruction, δ_L, we impose that the steady-state probability that a firm-worker match breaks equals the average job destruction rate in the US from 1948 to 2010 (Shimer, 2012); thus $JD \equiv \delta_L + (1 - \delta_L)\lambda_n G(z^*) = 0.036$. For the automated technology, we assume it is $\delta_K = 0.01$. We do not impose any particular value for λ_e and λ_n; instead we analyze how different values of these two parameters change our results.

Finally, we normalize κ_K and κ_L to unity and set z_L, z_K, χ, and Ω such that our steady-state matches four targets. We target the prime-age (aged 25-54) workers' em-

[15] The per-period probability that a task is automated is given by $(1 - \delta_L)\lambda_n G(z^*)$, which is the endogenous component of the job destruction probability defined as $JD \equiv \delta_L + (1 - \delta_L)\lambda_n G(z^*)$. Higher levels of λ_n imply a higher sensitivity of the job destruction probability to changes in z^*.

[16] This implies that the most productive manual firms are 67% more productive than the least productive manual firms, which is slightly below the empirical estimates in, e.g, Syverson (2011) and OECD (2017) for all firms in manufacturing. Yet, in our sensitivity analysis, we show that the distribution of productivity draws affects the results quantitatively but does not change our main messages.

Table 1: Benchmark Calibration

Discount factor:	$\beta = 0.996$
Matching function elasticity:	$\eta = 0.5$
Workers' bargaining power:	$\phi = 0.5$
Minimum productivity draw:	$z_{min} = 0.15$
Maximum productivity draw:	$\bar{z} = 0.25$
Nonemployment income:	$b = 0.7z_L\left(z_{min} + \frac{\bar{z}+z_{min}}{2}\right)$
Rate of automated-firm destruction:	$\delta_K = 0.01$
Cost of Capital/Robot:	$\kappa_K = 1$
Job-filling Cost:	$\kappa_L = 1$

ployment rate and the labor share in the US from 1977 until 2018;[17] this implies that $n = 0.78$ and $LS = 0.61$. We also target $G(z_e^*) = 0.5$ such that half of the productivity draws exceed the entry cutoff in our various experiments, but run sensitivity analysis on this target.[18] And, following Pissarides (2009), we target labor market tightness in the US so that $\theta = 0.72$.

3.2 Employment: Is the Future like the Past?

Figure 2 summarizes our main results. On the left, this figure plots how an increase of 1% in the productivity of the automated technology, z_K, changes employment, n, when all firms draw productivity at the time of entry ($\lambda_e = 1$) and under different values of λ_n. Clearly, the probability that workers lose their comparative advantage and are endogenously fired – controlled by λ_n – affects the response of employment to an automation-augmenting shock (rise in z_K). In the case of (very) low λ_n, manual firms rarely automate the production of the tasks, and a rise in z_K slightly increases employment.[19] If, however, we assume larger values of λ_n, an automation-augmenting shock may lead to sizable losses in employment: if $\lambda_n = 0.15$, manual firms are more likely to automate the production of the tasks after the shock, and employment falls 2.5%, that

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

[18] Our baseline calibration implies that the frequency at which a task is on average automated ranges between eight and 20 years depending on λ_e and λ_n.

[19] In this case, our model in this paper is close to that in Guimarães and Gil (2019); thus, the results are similar.

Figure 2: The effect of higher z_K under $\lambda_e = 1$ and different values of λ_n

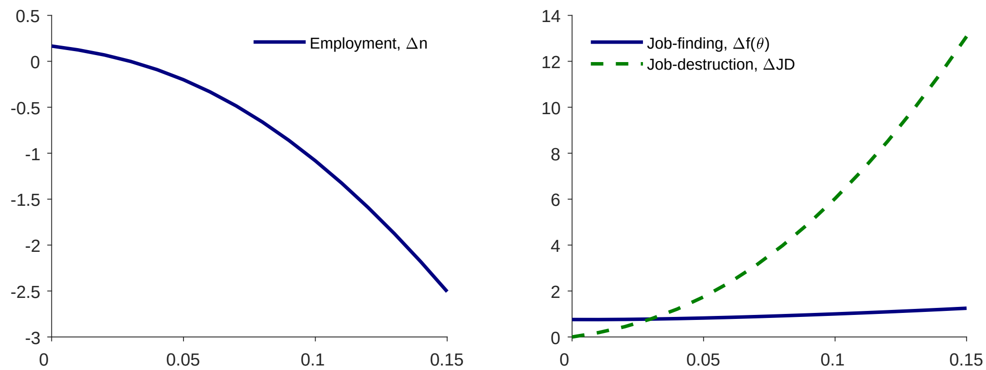

Note: This figure shows the effects of an automation-augmenting shock in the case in which all firms draw the tasks' productivity at the time of entry, $\lambda_e = 1$, and for different probabilities that this productivity changes, λ_n. The left-panel shows the percentage change in employment, n. The right-panel shows the percentage change in the job-finding rate, $f(\theta)$, and in the job-destruction rate, $JD \equiv \delta_L + (1 - \delta_L)\lambda_n G(z^*)$. The shock to z_K is of 1%.

is, two and a half times the magnitude of the shock to z_K.[20] Remarkably, this magnitude is insensitive to changes in most parameters as shown in Section 3.3.

Shocks in the economy affect employment through changes in both job creation and job destruction. Thus, to shed more light on the mechanisms in our model, we decompose the two effects of an automation-augmenting shock of 1% on employment on the right-hand side of Figure 2. In particular, we show how the job-finding probability, $f(\theta)$, (which indicates job creation) and the job-destruction probability, $JD \equiv \delta_L + (1 - \delta_L)\lambda_n G(z^*)$, react to the automation-augmenting shock (also as a function of λ_n and in the case of $\lambda_e = 1$). To understand how a rise in z_K affects employment, let's first consider the extreme case of $\lambda_n = 0$. This case implies that tasks that start as manual are never automated: tasks have constant idiosyncratic productivity, z, meaning that workers never lose their comparative advantage; thus, firms have no incentive to shift from the manual to the automated technology in equilibrium. As a result, $\lambda_n = 0$ also implies that job destruction is constant and unaffected by the automation-augmenting shock.

[20] $\lambda_n = 0.15$ implies that the tasks' productivity is redrawn, on average, approximately every six months. Our benchmark calibration then implies that a task is initially on average automated every 16 years; after the shock, in our experiment, it is automated on average every 14 years.

The same is not true for job creation. 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.[21] The free-entry condition, Eq. (4), is only satisfied if the value of the manual technology drops, which occurs in our model through higher wages and, most importantly, greater labor market tightness. A tighter labor market is synonym of greater job-finding probability and, necessarily, higher job creation. Therefore, if $\lambda_e = 1$, the aggregate effect of greater firm entry exceeds the reallocation effect implied by the increase in z_e^* and, thus, an automation-augmenting shock increases job creation. This, together with the constant job destruction ($\lambda_n = 0$), increases employment.

If $\lambda_n > 0$, a rise in z_K affects both job-finding and job-destruction probabilities.[22] As before, the job-finding probability increases because a rise in z_K boosts entry more than it boosts reallocation at the time of entry. Because firms are forward-looking, they have an even higher incentive to create jobs and invest in the manual technology (when $\lambda_n > 0$ than when $\lambda_n = 0$) in anticipation of the greater profits when they automate production. But the job-destruction probability also increases: as machines are more productive, firms that use the manual technology are motivated to shift to the automated one. This translates into a higher z^*, reducing the average time of a worker-firm match. Because λ_n is the probability that the firm redraws the productivity of the task, a higher λ_n increases the number of manual firms drawing low productivity (for a given z^*), leading to even greater job destruction. If λ_n is large enough, then the increase in job destruction surpasses the increase in job creation, implying less employment.[23]

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

[22] In this case, our model differs substantially from our previous work in Guimarães and Gil (2019) by endogenizing job destruction, which implies remarkably different results for the effects of automation-augmenting shocks on employment under some calibrations.

[23] The increase in z^* exacerbates the rise on the left-hand side of Eq. (4) as firms only destroy jobs if it is

The rise in z_K may lead to greater employment even if we mute the general-equilibrium effect at the time of entry and set $\lambda_e = 0$. The bottom three lines of Table 2 show the effects of higher z_K on the employment, job-finding probability, and job-destruction probability (besides the labor share) when $\lambda_e = 0$ and λ_n equals 0.01, 0.05, or 0.15. To allow for a direct comparison, the top four lines of Table 2 show the same experiments when $\lambda_e = 1$ (and we include the case of $\lambda_n = 0$ for completeness). If $\lambda_e = 0$, all tasks demand labor when created, as in Acemoglu and Restrepo (2018), and firms may only take advantage of the increased productivity if they automate the production of the task. Thus, it is remarkable that an increase in z_K – the productivity of a technology that can only be used after a job is destroyed – is still capable of leading to greater employment under a slightly positive λ_n (see the line regarding $\lambda_e = 0$ and $\lambda_n = 0.01$ in Table 2). Indeed, in the case of $\lambda_e = 0$, an increase in z_K continues to affect both job creation and job destruction. First, it continues to promote greater firm entry and job creation because of the increase in the value of the automated technology, as an outside option of the firms using the manual technology. But different from the case of $\lambda_e > 0$, if $\lambda_e = 0$, workers only benefit from larger firm entry because all firms start as manual and must hire a worker. Second, an automation-augmenting shock implies that firms have a higher opportunity cost of employing the worker and, thus, prefer to shift earlier to the automated technology (z^* increases). This increases job destruction. If λ_n is low, the job-creation effect dominates; but if λ_n is large, the job-destruction effect dominates.[24]

more profitable for them ($J_L(z)$ increases for all z; see Eq. (1)). Thus, a higher increase in job destruction must be accompanied by an even tighter labor market. But, as we will show in Section 4.1, λ_n affects job destruction by more than job creation because the automation-augmenting shock increases wages and the prevalence of matching frictions.

[24]It is not possible to pin down analytically why this result obtains in the case of $\lambda_e = 0$. But there are two aspects that offer a hint on why it happens. First, when λ_n is low, the weight of endogenous job destruction on total job destruction, JD, is very low: a change in z^* barely alters JD if λ_n is low. Yet, λ_n does not change the elasticity of $f(\theta)$ with respect to θ. Second, if we use Eqs. (1) and (6) both measured at $\bar{z}$ and z^* together with the firing cutoff equation, Eq. (2), and free-entry condition, Eq. (4), we obtain

$$\frac{\kappa_L}{\beta\mu(\theta)} = (1-\phi)\left[\beta J_K - \kappa_K + \frac{z_L(\bar{z}-z^*)}{1-\beta(1-\delta_L)(1-\lambda_n)}\right] - \Omega\left(\frac{1}{\beta}-\phi\right).$$

To properly assess assess how z^* and θ affect each other, we need another equation relating them. But

Table 2: The effect of an increase of 1% in z_K

λ_e	λ_n	Δn	$\Delta f(\theta)$	ΔJD	ΔLS
1	0	0.17	0.76	0.00	-0.52
1	0.01	0.12	0.76	0.19	-0.71
1	0.05	-0.20	0.82	1.74	-1.87
1	0.15	-2.51	1.25	13.08	-8.35
0	0.01	0.01	0.19	0.15	-0.06
0	0.05	-0.20	0.56	1.46	-0.58
0	0.15	-2.21	0.97	11.34	-4.37

Note: This table shows the effects of an automation-augmenting shock under various combinations of the probability that the task's productivity is drawn at the time of entry, λ_e, and the probability that it is redrawn afterwards, λ_n. The first two columns show the calibration of these two probabilities. The next four columns show the percentage change in employment, job-finding probability, job-destruction probability, and labor share. The shock to z_K is of 1%.

Our results show how our model may both agree and disagree with the empirical literature (Bessen, 2016; Autor and Salomons, 2018; Gregory, Salomons and Zierahn, 2018; and Bessen et al., 2020). Under some calibrations, job creation increases more than job destruction, agreeing with their findings that employment increased after productivity enhancements in the past; this increase in employment also coincided with the fall in the labor share, which is in line with our results. Therefore, our model suggests that λ_n has been about 0.02 or lower**, implying that labor's comparative advantage in producing a task has been stable for at least four years on average**. But under other calibrations, job destruction increases more than job creation and employment may significantly fall. Thus, this suggests that the future of employment may differ from the past. Our model calls the attention specifically to λ_n, which we interpret as a feature intrinsic to tasks that characterizes how rapidly workers may lose their comparative advantage. In an economy in which workers rapidly lose their comparative advantage (*rapid-changing environments*; high λ_n) and with matching frictions, em-

the equation above shows that the labor market becomes tighter when the productivity of the automated technology goes up (J_K increases). This is a direct effect that takes into account that without a change in z^*, the increase in z_K directly increases the value of the firm in the cases in which the task is already automated. This naturally increases the value of a job and, thus, job creation. This equation also shows that a rise in z^* reduces θ (because jobs last for less periods) and that the elasticity of θ with respect to z^* increases with λ_n (we confirm this numerically given that z_L and δ_L are used to reach our steady-state targets). Thus, for a given change in θ, if λ_n is low, z^* cannot change much to satisfy this equation. Furthermore, any change in z^* has a minor effect on JD. But if λ_n is higher, z^* has to fluctuate more to satisfy this equation and has a larger impact on JD, shifting the ranking of the forces at play.

ployment falls after an automation-augmenting shock. In this economy, jobs last for less periods and the increase in labor market tightness makes it more costly to hire the right worker for the task and allows workers to enjoy greater wages. These effects prevent job creation from keeping pace with job destruction. Therefore, if the nature of the new and current jobs is different from the past – particularly, if tasks feature a higher λ_n in the future than in the past and, thus, tasks rapidly become liable to be automated – the same productivity shock of the past may have dramatically different consequences in the future. This may occur especially if technologies like Artificial Intelligence allow software and robots to rapidly adapt to new tasks once enough data is gathered.

3.3 Sensitivity Analysis

In this section, we assess how different calibrations of our model change the outcomes of an automation-augmenting shock of 1%. We consider seven experiments, and in each experiment we recalibrate one parameter (or target) of the model. We conclude that none of the experiments changes the qualitative predictions of our model. In all cases, both job creation and job destruction increase after an automation-augmenting shock (except in the case of $\lambda_n = 0$, in which case the job-destruction probability is constant by assumption). And the change in the job-destruction probability is still more sensitive to λ_n than the change in the job-finding probability. This implies a negative relation between the change in employment after the rise in z_K and λ_n: if λ_n is low, employment increases; on the contrary, if λ_n is high, employment falls.

Our experiments do, however, change the results quantitatively. And among our seven experiments, two have particularly large quantitative effects that we show in Panels B and C of Table 3. These two panels show how a rise in z_K affects employment, job-finding probability, and job-destruction probability in economies with $\bar{z} = 0.225$ (instead of $\bar{z} = 0.25$) and with a Pareto distribution of productivity draws (instead of a uniform distribution), respectively. As in Table 2, we consider various combinations of λ_n and λ_e. And to ease comparability with the results of our model using the baseline calibration (reported in Table 2), we reproduce those results in Panel A of Table 3.

Table 3: The effect of an increase of 1% in z_K – Sensitivity Analysis

		A: Baseline			B: $\bar{z} = 0.225$			C: Pareto			D: $\eta = 0.4$		
λ_e	λ_n	Δn	$\Delta f(\theta)$	ΔJD	Δn	$\Delta f(\theta)$	ΔJD	Δn	$\Delta f(\theta)$	ΔJD	Δn	$\Delta f(\theta)$	ΔJD
1	0	0.17	0.76	0.00	0.19	0.87	0.00	0.30	1.40	0.00	0.21	0.95	0.00
1	0.01	0.12	0.76	0.19	0.13	0.85	0.25	0.21	1.20	0.24	0.17	0.95	0.19
1	0.05	-0.20	0.82	1.74	-0.30	0.93	2.30	-0.52	1.28	3.69	-0.16	1.04	1.80
1	0.15	-2.51	1.25	13.08	-3.25	1.51	17.00	-6.65	2.82	36.11	-2.51	1.58	13.45
0	0.01	0.01	0.19	0.15	0.03	0.34	0.21	-0.00	0.13	0.14	0.02	0.24	0.15
0	0.05	-0.20	0.56	1.46	-0.27	0.69	1.94	-0.35	0.73	2.34	-0.17	0.70	1.48
0	0.15	-2.21	0.97	11.34	-2.84	1.18	14.61	-5.09	2.05	26.93	-2.20	1.22	11.57

Note: This table shows the effects of an automation-augmenting shock under various combinations of the probability that the task's productivity is drawn at the time of entry, λ_e, and the probability that it is redrawn afterwards, λ_n. The first two columns show the calibration of these two probabilities. The next columns show the percentage change in the employment, job-finding probability, and job-destruction probability under a slightly different calibration in each panel. The shock to z_K is of 1%. Panel A presents the baseline results; Panel B presents the results assuming a lower maximum productivity draw; Panel C presents the results assuming a Pareto distribution of productivity draws; Panel D presents the results assuming a lower elasticity of the matching function.

Economies with a tighter range of productivity draws (low $\bar{z}$ or high z_{min}) experience larger changes in the flows after the rise in z_K and also tend to experience larger changes in employment than in our baseline economy. We also find a similar result in the case of the Pareto distribution. If the cumulative distribution of productivity draws is of the form $G(z) = 1 - \left(\frac{z_{min}}{z}\right)^{\xi}$, a higher ξ (which concentrates productivity draws near the minimum) increases the effects of the shock.[25] The intuition is simple. If we reduce $\bar{z}$ or increase ξ, the distribution of productivity draws becomes more concentrated and, thus, the same change in z^* and z_e^* alters the optimal decision of a larger proportion of firms. In these circumstances, the same rise in z_K amplifies the required change in labor market tightness, θ, to balance the free-entry condition, Eq. (4), and – most importantly – motivates a much larger proportion of manual firms to destroy jobs and automate the production of the tasks. Therefore, these experiments paint an even bleaker picture than our baseline: depending on the calibration, the fall in employment after the shock can be as catastrophic as 6.5-fold the magnitude of the shock.

[25] In Panel C of Table 3, we assume that $\xi = 5$. In all our experiments with the Pareto distribution, we continue assuming that firms that do not draw productivity at the time of entry start with productivity $\bar{z}$.

In Panel D of Table 3, we consider the case of a smaller matching function elasticity, $\eta = 0.4$ (instead of $\eta = 0.5$). We consider this case as it reduces the elasticity of the hiring costs, $\frac{\kappa_L}{\mu(\theta)} = \frac{\kappa_L \theta^\eta}{\chi}$, relative to labor market tightness, θ. As a result, we would expect greater flows in the labor market, particularly for job creation, to balance the free-entry condition, Eq. (4). We show that this does occur but the final impact of reducing η on employment is whimsy because it also magnifies job destruction.[26] Finally, we consider the cases of a lower cost of capital, κ_K, lower workers' bargaining power, ϕ, lower proportion of firms that draw productivity below the entry cutoff, $G(z_e^*)$, and lower job-filling costs, κ_L. The results of these experiments are detailed in Tables A1 and A2, which we relegate to the Appendix A as they barely affect the results of our model.

4 Dissecting the Mechanism

4.1 Ad hoc Function for Wages

Our baseline model shows that after an automation-augmenting shock, employment increases if λ_n is low and falls if λ_n is large. We find that both job creation and job destruction increase after a rise in z_K (unless $\lambda_n = 0$, in which case the job destruction rate is fixed). But, the change in the job-destruction rate increases much more with λ_n than the change in the job-finding rate. One factor that may explain this behavior is the wage response. In all our calibrations, wages increase due to the rise in the job-finding probability and in the value of the manual firm (which increases namely due to 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. The fall in employment is, to some extent, surprising because it coincides with an increase in wages. So we ask: if wages were only a function of the task's productivity, how would the job-creation and job-destruction margins react to an increase in z_K? In other words, if wages would not increase with automation-augmenting

[26] In Section 4.1, we explain that the good effects of a lower η on job creation also promote higher wages, which motivate firms to destroy jobs and automate the production of the tasks.

shocks, 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.

Panel B of Table 4 shows how employment, job-finding probability (indicator of job creation), and job-destruction probability change after an automation-augmenting shock of 1% under various combinations of λ_e and λ_n.[27] For convenience, Panel A of the same table reproduces the results for the same experiments using our baseline model of Section 2. The main takeaway from Panel B is that employment increases for all the combinations we consider of λ_e and λ_n, which is in stark contrast with the results reported in Panel A. Therefore, the rise in wages in our baseline model crucially influences the fall in employment. By further contrasting Panels A and B, we see that the job-finding rate increases much more while the job-destruction rate increases less in this version of the model than in the baseline one. Thus, 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. Furthermore, and by the same token, firms have less incentives to fire the worker and move to the automated technology.

Panel B of Table 4 also shows that the change in the job-destruction rate continues to increase much more with λ_n than the change in the job-finding rate. The net effect is that the change in employment tends to be negatively related with λ_n, which suggests that for sufficiently high λ_n, employment may still drop after an automation-augmenting shock. We confirm this in parallel experiments: employment falls if $\lambda_e = 1$ and $\lambda_n \geq 0.22$, as well as if $\lambda_e = 0$ and $\lambda_n \geq 0.2$, because job destruction increases more

[27] To calibrate the model with the ad hoc wage, we start by determining z_L and b using our baseline model under each calibration. Once determined z_L and b, we obtain ϕ_{nb} together with z_K, Ω, and δ_L to reach our targets for the employment rate, labor share, job-destruction rate, and $G(z_e^*)$.

Table 4: The effect of an increase of 1% in z_K – Model comparison

		A: Baseline			B: Ad hoc wage			C: Low friction			D: CES ($\epsilon = 5$)		
λ_e	λ_n	Δn	$\Delta f(\theta)$	ΔJD	Δn	$\Delta f(\theta)$	ΔJD	Δn	$\Delta f(\theta)$	ΔJD	Δn	$\Delta f(\theta)$	ΔJD
1	0	0.17	0.76	0.00	1.53	7.35	0.00	0.19	0.76	0.00	0.13	0.59	0.00
1	0.01	0.12	0.76	0.19	1.50	7.36	0.17	0.15	0.76	0.18	0.10	0.57	0.11
1	0.05	-0.20	0.82	1.74	1.35	7.93	1.40	-0.16	0.87	1.71	-0.04	0.54	0.72
1	0.15	-2.51	1.25	13.08	0.69	12.83	9.31	-2.30	1.50	12.55	-0.25	0.54	1.66
0	0.01	0.01	0.19	0.15	0.40	2.03	0.18	0.01	0.19	0.15	0.01	0.16	0.10
0	0.05	-0.20	0.56	1.46	1.04	6.46	1.47	-0.17	0.58	1.44	-0.08	0.44	0.78
0	0.15	-2.21	0.97	11.34	0.50	11.99	9.47	-2.06	1.17	10.99	-0.34	0.52	2.07

Note: This table shows the effects of an automation-augmenting shock under various combinations of the probability that the task's productivity is drawn at the time of entry, λ_e, and the probability that it is redrawn afterwards, λ_n. The first two columns show the calibration of these two probabilities. The next columns, divided in four panels, show the percentage change in the employment, job-finding probability, and job-destruction probability. In each panel, we use a different version of our model. The shock to z_K is of 1%.

than job creation after the rise in z_K. It is natural that the job destruction rate increases with λ_n as this rate becomes more sensitive to endogenous factors. Yet, at first sight, it is unclear why the job-finding probability increases less than the job-destruction probability given that there are also greater incentives to create new tasks and jobs if λ_n is high.

We conjecture that matching frictions are behind this pattern. As the 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 the labor market tightness ($\frac{\kappa_L}{\mu(\theta)} = \frac{\kappa_L \theta^\eta}{\chi}$). Thus, matching frictions are less relevant for job creation and we should observe greater job creation after an automation-augmenting shock. Using our baseline model, in Section 3.3, we concluded that η barely affects how employment reacts to the increase in z_K. Yet, Figure 3 shows a different result if we use our model with the ad hoc wage equation; in fact, it confirms our conjecture that matching frictions prevent a greater increase in employment. This figure plots how the job-finding and job-destruction rates change after the automation-augmenting shock for a range of values of λ_n and using our model with the ad hoc wage equation. On both panels,

Figure 3: The effect of higher z_K under $\lambda_e = 1$ and different values of λ_n and η

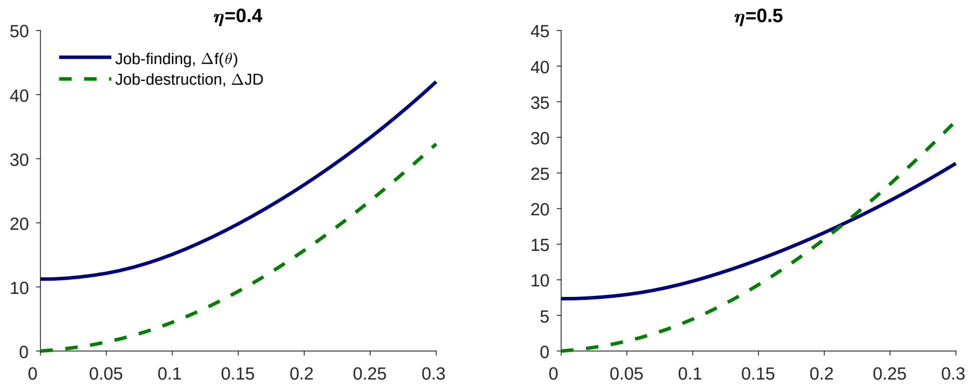

Note: This figure shows the effects of an automation-augmenting shock using our model with the ad hoc wage equation. In both panels, we assume that all firms draw the tasks' productivity at the time of entry, $\lambda_e = 1$, and different probabilities that this productivity is redrawn afterwards, λ_n. Both panels show the percentage change in the job-finding probability and in the job-destruction probability. In Panel A, $\eta = 0.4$; in panel B, $\eta = 0.5$. The shock to z_K is of 1%.

$\lambda_e = 1$. The difference between the panels lies only in the value of η: the left-panel assumes $\eta = 0.4$; the right-panel assumes $\eta = 0.5$. Confirming our conjecture, job creation increases much more after the rise in z_K if $\eta = 0.4$ than if $\eta = 0.5$. Interestingly, η barely affects the change in job destruction. Thus, employment reacts more after an automation-augmenting shock if $\eta = 0.4$. But why are the results so different when we use Nash bargaining and when we use our ad hoc equation? The reason seems to lie in the outside option of workers, U. If the job-filling probability, $\mu(\theta) = \chi\theta^{-\eta}$, is less sensitive to changes in labor market tightness, θ, then the job-finding probability, $f(\theta) = \chi\theta^{1-\eta}$, is more sensitive. Thus, given that U and $f(\theta)$ are positively related (see Eq. 7), *ceteris paribus* a lower η increases the elasticity of U relative to θ, allowing all workers to demand greater wages. Our ad hoc wage, however, prevents the operational profit of manual firms to be affected by U, leading to the different results.

Our experiments with the model assuming the ad hoc equation work as counterfactuals to understand the dynamics in our original model. But these experiments do not seem to be a good account of how an automation-augmenting shock is likely to unfold in the future. Unless the historical positive relationship between labor market tightness and wage increments definitely breaks in the future, the automation-augmenting shock will increase wages, which may promote the sizable negative employment effects that we obtain using our baseline model.

4.2 Lower Frictions

Our baseline model suggests that, as λ_n increases, it becomes easier to fire a worker than to hire a worker due to matching frictions, because the latter increase wages and the costs to find a suitable worker. Matching frictions in our model come from the matching function but also come from our assumption that workers who lose jobs stay unemployed for at least a period (month). Although this is a typical assumption in models with matching frictions and finds support in the evidence (Hall and Kudlyak, 2019), we can argue that in an economy that experiences a surge in labor market flows, this assumption may be too restrictive. In such an economy, it is likely that workers find jobs even within a month from losing them and start production immediately.[28] Relaxing this assumption may be important in our model: in an economy that experiences a surge in job destruction, the pool of available workers to match with firms may become too narrow, raising the relevance of matching frictions. Thus, we ask: what are the predictions of our model if workers can look for jobs and start production immediately after losing their jobs?

Panel C of Table 4 answers this question and, by contrasting the results in this panel with those in Panel A, confirms our prediction. In an economy that experiences an automation-augmenting shock and in which workers who lose jobs can look for other jobs and restart production immediately, matching frictions become less relevant and the job-finding probability increases more with λ_n. The implication of this is that employment becomes less negatively correlated with λ_n; yet, and even though this model also generates a smaller increase in the job destruction rate than the baseline, the change

[28] Christiano, Eichenbaum and Trabandt (2016) make a similar assumption. They build a model with matching frictions but calibrate each period as a quarter, whereas tipically these models are calibrated with monthly data. Because in US data many workers find jobs and start production within a quarter, it would be too restrictive to assume that workers who lose jobs need to wait for the quarter to end to restart production. In our case, the probability to find jobs may increase so much that it can be equally restrictive.

in employment continues to fall significantly with λ_n.

The lack of firepower of this experiment is not completely surprising. First, our sensitivity analysis with η in Section 3.3 shows that our results are not much sensitive to the calibration of the matching function. This suggests that the degree of matching frictions are not much quantitatively relevant in our model.[29] Second, the change in the pool of nonemployed workers imposed by the rise in the job-destruction rate is not so great. Even in the case of $\lambda_e = 1$ and $\lambda_n = 0.15$, the rise in the job-destruction rate displaces only an additional 0.0037 proportion of the workforce per period. Given our steady-state target of nonemployment of $1 - n = 0.22$, the number of workers looking for jobs is not much affected.

4.3 CES Aggregator

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. 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, then, 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 an automation-augmenting shock affects

[29] This follows from the fact that the wage increases with the automation-augmenting shock (see the discussion in Section 4.1) and the remaining parameters adjust to balance the steady-state of our model and reach our steady-state targets.

employment. In particular, can this setup reverse our prediction that economies with high λ_n experience lower employment after an automation-augmenting shock? Or are there any relevant quantitative implications?

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:

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

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).

Panel D of Table 4 shows the effects of an automation-augmenting shock in the model with the CES assuming $\epsilon = 5$ and under the various combinations of λ_e and λ_n. Assuming that the outputs of the two technologies are imperfect substitutes does not change our results qualitatively. In economies with high λ_n, employment still falls. Yet, our setup with a CES affects the results quantitatively: it reduces the elasticities in the model because the total impact of the shock, $p_K z_K$, is lower reflecting the fall in the price of the automated good, p_K, after the rise in z_K.

One interesting outcome reported in Panel D of Table 4 is that our setup with the CES constrains job destruction much more than job creation. To shed light on this, on the left panel of Figure 4, we plot how the job-destruction probability, JD, and job-finding probability, $f(\theta)$, change with the elasticity of substitution, ϵ, under the case of $\lambda_e = 1$ and $\lambda_n = 0.15$. On the right-hand side of the same figure, we plot the prices of the tasks produced by each type of technology also as a function of ϵ. The shock is, as usual, an automation-augmenting shock of 1%. Undoubtedly, the job-destruction margin is much more affected by the elasticity of substitution to the point that the shift of the two margins almost converges if $\epsilon = 2$. (Recall that in the baseline, $\epsilon = \infty$, the job-finding probability increases 1.25% and the job-destruction probability increases 13%). There are two aspects that can explain this. First, an automation-augmenting shock reduces p_K and, thus, curbs down the increase in machines' productivity, $p_K z_K$. This naturally reduces the incentives to destroy jobs and automate tasks after the rise in z_K. It also reduces the incentives to create jobs as the shock has a lower impact on the value of firms. Yet, the same automation-augmenting shock increases p_L and, thus, increases workers' productivity, $p_L z_L z$. This balances the effect (of lower $p_K z_K$) on job creation but further reduces the motivation to destroy jobs and automate tasks. As we increase ϵ, the fall in p_K and the rise in p_L become smaller; thus, the incentives to automate and destroy jobs increase significantly while job creation changes much less as the effects of the two prices tend to almost balance out.

These mechanisms help explain why in calibrations with high λ_n (keeping ϵ fixed),

Figure 4: The effect of higher z_K under $\lambda_e = 1$, $\lambda_n = 0.15$, and different values of ϵ

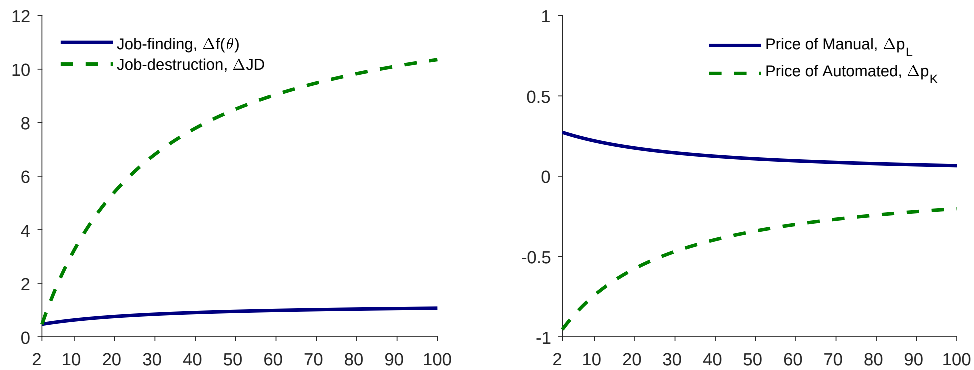

Note: This figure shows the effects of an automation-augmenting shock using our model with the CES aggregator. To produce these results, we assume that all firms draw the tasks' productivity at the time of entry, $\lambda_e = 1$, and that on average about every six months this productivity is redrawn afterwards, $\lambda_n = 0.15$. The left-panel shows the percentage change in the job-finding probability and in the job-destruction probability. The right-panel shows the percentage change in the price of tasks produced using the manual technology and in the price of the tasks produced using the automated technology. The shock to z_K is of 1%.

the assumption of imperfect-substitutability between the two outputs, y_K and y_L, (results reported in Panel D of Table 4) affects job-destruction much more than job creation. Economies with high λ_n experience greater reallocation from the manual to the automated technology after an automation-augmenting shock. Greater reallocation then implies a greater rise in the number of firms using the automated technology, n_K, and, thus, in the output produced using the automated technology, y_K. In this setup with the CES, the greater rise in y_K further drops p_K and further increases p_L, leading to lower incentives to fire workers and, thus, a greater drop in job destruction when contrasted with the baseline results. The two effects of p_K and p_L tend to balance the change in job creation, leading to the smaller relative drop in job creation when compared with the baseline.

These experiments with the CES aggregator show that consumers have an important role in determining the effects of automation-augmenting shocks on employment. If a large proportion of the tasks are directed to consumers, their preference for the *human touch* may severely reduce the negative effects of automation-augmenting shocks on employment.

5 Concluding Remarks

In this paper, we build a model to assess how an automation-augmenting shock – a generalized increase in the productivity of machines/robots – affects employment. This model relies on multiple previous contributions (Mortensen and Pissarides, 1994; Hornstein, Krusell and Violante, 2007; Acemoglu and Restrepo, 2018; and Guimarães and Gil, 2019) to satisfy two criteria. First, it is consistent with the past documented by the empirical literature (e.g., Bessen, 2016; Autor and Salomons, 2018; Gregory, Salomons and Zierahn, 2018; and Bessen et al., 2020): an automation-augmenting shock can simultaneously reduce the labor share and increase employment. Second, our model is flexible enough to offer insights on how the future may differ from the past: depending on the calibration, an automation-augmenting shock may increase or decrease employment.

In our model, an automation-augmenting shock increases 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. Yet, this robust 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). **[novo] Yet, the downward trend in labor market fluidity 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 quar-**

ter; 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.[30] Furthermore, even though aggregate labor market flows fell in all sectors, they fell unevenly across them. Particularly, 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, the calibrations implying a rise in employment as observed in the past coincide with a relatively low elasticity of job separation and job finding rates with respect to automation-augmenting shocks. Thus, the effect of automation on labor market fluidity was likely small in the past.

[**old; apagar**] But a closer look into the changes in labor market flows across US sectors reveals that, even though labor market flows fell in all sectors, they fell unevenly across them. Particularly, 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. We can interpret these patterns in light of two trends: a general trend reducing labor market flows in all sectors (e.g., demographics as argued by Engbom) and a trend increasing labor market flows in some sectors (with greater pervasiveness of automation). Our model abstracts from the general trend and only takes into account the positive contribution of automation-augmenting shocks to labor market flows. Furthermore, the calibrations implying a rise in employment as observed in the past imply tiny changes to labor market flows. These changes are likely too small to clearly affect aggregate dynamics.

Using our model, we sort the cases in which employment increases after an automation-augmenting shock and those in which it falls. In environments in which the comparative advantage of workers in producing a task is relatively stable – *slow-changing* environments – the increase in job creation dominates the increase in job destruction.

[30] 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.

Therefore, in slow-changing environments, employment increases. On the contrary, in environments in which the comparative advantage of workers changes frequently – *rapid-changing* environments – an automation-augmenting shock leads to massive job destruction that clearly offsets the increase in job creation. In these environments, employment can catastrophically fall.

We also find that the fall in employment in rapid-changing environments crucially depends on 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. If that is the case, an automation-augmenting shock (and ensuing spread of usage of machines/robots) increases the price of the tasks produced by workers relative to those produced by the machines/robots. This curtails job destruction, reducing the fall in employment.

Our paper then clarifies how the future may differ from the past. If the comparative advantage of workers in producing new tasks starts to vanish more rapidly than in the past, then automation-augmenting shocks will curb down employment rather than increase it. In this respect, Artificial Intelligence can change the paradigm as these technologies allow software and robots to rapidly adapt in order to perform new tasks once enough data is gathered. Thus, tasks can more rapidly be automated than in the past, increasing the pace at which workers lose their comparative advantage in producing each task. In such a scenario, employment would drop with increases in the productivity of the automated technology. The extent of this fall will also naturally depend on demand and, particularly, consumers' preferences. If many of the tasks produced in an economy are sold directly to consumers and they have a preference for the *human touch*, then the fall in employment will unlikely be catastrophic. But if most of the tasks are part of a vast value chain to produce a final good or if consumers have no preference for the *human touch*, then the fall in employment in the future may be catastrophic.

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.

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*.

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.

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."

Caballero, Ricardo J., and Mohamad L. Hammour. 1998. "Jobless Growth: Appropriability, Factor Substitution, and Unemployment." *Carnegie-Rochester Conference Series on Public Policy*, 48: 51 – 94.

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

Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt. 2016. "Unemployment and Business Cycles." *Econometrica*, 84(4): 1523–1569.

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.*

Engbom, Niklas. 2019. "Firm and Worker Dynamics in an Aging Labor Market."

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."

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.

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

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.

OECD. 2017. *Entrepreneurship at a Glance 2017.*

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. 2011. "What determines Productivity?" *Journal of Economic literature*, 49(2): 326–365.

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

Zeira, Joseph. 2010. "Machines as Engines of Growth." mimeo.

A Further robustness checks

Table A1: The effect of an increase of 1% in z_K – Sensitivity Analysis

		A: Baseline			B: $\phi = 0.4$			C: $G(z_e^*) = 0.4$		
λ_e	λ_n	Δn	$\Delta f(\theta)$	ΔJD	Δn	$\Delta f(\theta)$	ΔJD	Δn	$\Delta f(\theta)$	ΔJD
1	0	0.17	0.76	0.00	0.16	0.71	0.00	0.22	1.02	0.00
1	0.01	0.12	0.76	0.19	0.12	0.72	0.18	0.17	0.95	0.18
1	0.05	-0.20	0.82	1.74	-0.19	0.81	1.67	-0.17	0.94	1.71
1	0.15	-2.51	1.25	13.08	-2.41	1.31	12.66	-2.41	1.42	12.81
0	0.01	0.01	0.19	0.15	0.01	0.17	0.15	0.00	0.15	0.13
0	0.05	-0.20	0.56	1.46	-0.20	0.53	1.42	-0.18	0.59	1.42
0	0.15	-2.21	0.97	11.34	-2.15	0.99	11.07	-2.16	1.11	11.27

Note: This table shows the effects of an automation-augmenting shock under various combinations of the probability that the task's productivity is drawn at the time of entry, λ_e, and the probability that it is redrawn afterwards, λ_n. The first two columns show the calibration of these two probabilities. The next columns show the percentage change in the employment, job-finding probability, and job-destruction probability under a slightly different calibration in each panel. The shock to z_K is of 1%. Panel A presents the baseline results; Panel B presents the results assuming a lower workers' bargaining power; Panel C presents the results assuming a lower proportion of productivity draws below the entry cutoff.

Table A2: The effect of an increase of 1% in z_K – Sensitivity Analysis

		A: Baseline			B: $\kappa_K = 0.5$			C: $\kappa_L = 0.5$		
λ_e	λ_n	Δn	$\Delta f(\theta)$	ΔJD	Δn	$\Delta f(\theta)$	ΔJD	Δn	$\Delta f(\theta)$	ΔJD
1	0	0.17	0.76	0.00	0.17	0.76	0.00	0.17	0.76	0.00
1	0.01	0.12	0.76	0.19	0.12	0.76	0.19	0.12	0.76	0.19
1	0.05	-0.20	0.82	1.74	-0.20	0.82	1.73	-0.21	0.82	1.77
1	0.15	-2.51	1.25	13.08	-2.49	1.25	13.00	-2.54	1.26	13.24
0	0.01	0.01	0.19	0.15	0.01	0.19	0.15	0.01	0.19	0.15
0	0.05	-0.20	0.56	1.46	-0.20	0.55	1.45	-0.20	0.56	1.47
0	0.15	-2.21	0.97	11.34	-2.20	0.97	11.30	-2.22	0.97	11.41

Note: This table shows the effects of an automation-augmenting shock under various combinations of the probability that the task's productivity is drawn at the time of entry, λ_e, and the probability that it is redrawn afterwards, λ_n. The first two columns show the calibration of these two probabilities. The next columns show the percentage change in the employment, job-finding probability, and job-destruction probability under a slightly different calibration in each panel. The shock to z_K is of 1%. Panel A presents the baseline results; Panel B presents the results assuming a lower cost of capital/robot; Panel C presents the results assuming lower job-filling costs.