Endogenous Automation*

Stéphane Auray[†] Aurélien Eyquem[‡]

First version: October 2018.
This version: July 14, 2023

Abstract

We develop an open-economy model of endogenous automation with heterogeneous firms, endogenous business creation and tradability and labor reallocation. We match key trends in the U.S. economy, among which the rise in automation. The decline in the relative price of robots is a key factor leading to automation, and it affects automation non-linearly – the price of robots needs to fall enough to trigger automation. Running counterfactual simulations, other shocks are also found to affect the extent of automation resulting from a given declining path of robots price: negative shocks slow down automation while positive shocks accelerate it. Hence factors such as rising markups, rising entry costs, rising labor productivity or declining trade costs have welfare costs/gains on their own but also affect the economy, labor markets or business dynamism through their effects on the adoption of robots by firms.

Keywords: Robots, Automation, Heterogeneous Firms, Labor Market, Open-economy

JEL Class.: E24, E32, F41.

*We would like to thank Basile Grassi and Sergio Rebelo for suggestions.

†CREST-Ensai and Université du Littoral Côte d'Opale. ENSAI, Campus de Ker-Lann, Rue Blaise Pascal, BP 37203, 35172 BRUZ Cedex, France. stephane.auray@ensai.fr.

‡Department of Economics, University of Lausanne, Internef, CH-1015 Lausanne, Switzerland. Email: aurelien.eyquem@unil.ch.

1 Introduction

"You can see the robots everywhere but in the productivity statistics" could be the modern version of Solow's productivity paradox. Since the 1980's, the price of robots fell very significantly, the number of robots per thousands of industry workers has been rising very fast (Acemoglu and Restrepo (2022)), and yet no major productivity improvement has been seen in TFP or GDP growth statistics, while it should intuitively result in large productivity gains. What if other changes in the environment of firms partly hindered the adoption of robots in the U.S. economy? What if automation was actually fostered by other factors?

We build a model of endogenous automation with heterogeneous firms, endogenous business formation and tradability, and labor reallocation. The model extends Ghironi and Melitz (2005) to include segmented labor markets and endogenous automation, and allows for a comprehensive account of the effects of automation on the labor market, business dynamism, productivity and trade. As automation is endogenous, it does not only depend on the relative price of robots as in many papers but is more generally affected by changes in the environment of firms, namely the marginal production cost, and the relative amounts of labor types used to produce, and both margins are affected by a variety of shocks.

We feed the model with a collection of shocks aimed at capturing different trends of the U.S. economy since the mid-80's: a fall in the price of robots, a rise in firms' markups (De Loecker, Eeckhout, and Unger (2020)), a rise in entry costs to account for the observed decline in business dynamism (Akcigit and Ates (2023)), a fall in trade costs accounting for the greater observed trade openness of the U.S. economy, and a rise in labor productivity. Calibrated to match the data, these shocks allow for counterfactual analyzes that single out the pure effects of the observed fall in the price of robots. More generally, the procedure and model allow for a decomposition of the respective contributions of the various shocks to GDP growth, business dynamism, productivity, or trade openness by shutting or tuning them down one by one.

In the model, when robots become cheaper, they are used as perfect substitutes to labor in routine tasks and complement non-routine tasks. This gives rise to (*i*) a displacement effect that lowers the demand for routine labor and raises the demand for non-routine labor and a rise in the non-routine wage premium, (*ii*) a productivity effect that lowers the production cost and results in more output, and (*iii*) a labor supply effect driven by the non-routine wage premium, as workers reallocate towards the non-routine sector, which further amplifies the displacement effect. In addition, automation fosters the creation of new firms and boosts business dynamism, which further deepens automation and magnifies its effects on output and labor markets. Our model also predicts large aggregate production and wage gains from automation but very unevenly distributed. As in Moll, Rachel, and Restrepo (2022), routine workers loose by all possible measures – lower employment, lower wages and lower consumption – while non-routine workers gain. As the number of non-routine workers rises, the aggregate gains are larger than the average of routine and non-routine per capita welfare gains.

Looking at counterfactual scenarios allows us to highlight important results. The first one is that, at least in our model, the fraction of automated tasks depends non-linearly on the relative price of robots. We find that smaller declines in the price of robots would not have produced the automation process observed in the data. Actually, below a certain threshold, the decline in robots price does not produce automation at all. The reason is that, automation being endogenous and depending on labor supply – and not only labor demand – and labor reallocation responding to the wage premium, households need to experience or expect a substantial enough rise in the non-routine wage premium to reallocate. The reallocation process further deepens the extent of automated tasks. In our baseline experiment, the relative price of robots falls by 2 percents annually. Whenever the decline becomes less than 1.5 percent annually, we find that the fraction of automated tasks remains constant or even declines.

A second key finding is that the potential effects and benefits from the fall in the relative price of robots could be much larger than put forth by other papers, but that other negative shocks, such as the rise in markups or the decline in business dynamism have been offsetting these potential gains. Indeed, absent the rise in markups or entry costs, we find that the fraction of automated tasks would have risen to 60 or 70 percents, instead of 36 percents in the baseline simulation. This would have led to much stronger wage growth for non-routine workers, a more massive reallocation of labor to the non-routine sector and much larger aggregate welfare gains. Of course, part of the reported welfare gains relate to the fact that both shocks – pushing up markups and entry costs – generate welfare losses on their own, so removing them brings welfare gains. But a substantial fraction of the the reported welfare gains stems from a stronger automation process. Singling out the marginal effects of the automation induced by removing these negative shocks, we find that the corresponding aggregate welfare gains could have reached 30 to 35 percents of permanent consumption with constant markups and business dynamism.

By the same token, we find that positive shocks increasing labor productivity and reducing trade costs helped foster the automation process. Had labor productivity or trade costs remained constant, the fraction of automated tasks would have been around 30 percents instead of 36 percents. Labor reallocation and the wage gap would have been less important, and so the welfare consequences. While the absence of changes in trade costs has relatively little additional consequences except for the degree of trade openness, the absence of productivity gains would have also significantly moderated wage growth for both types of workers, generating large welfare losses that would have partly compensated the welfare gains experienced by non-routine workers and deepened the welfare losses suffered by routine workers.

Our paper thus suggests that, for a given path of the relative price of robots, the resulting degree of automation can vary more than substantially depending on how other shocks shape labor markets, aggregate productivity or business dynamism. When automation is endogenous, shocks with negative (positive) effects on output not only impose welfare costs (gains) on their own, but also slow down (accelerate) the process of automation, and dampen (amplify) the associated welfare and labor-market reallocation effects.

Literature. Recent contributions by Acemoglu and Restrepo (2020) and Acemoglu and Restrepo (2022) document the upward trend in robots' adoption and automation over the last 30 years around the world. Building on a task-model of automation, they also show theoretically and empirically that automation gives rises to a displacement effect, lowering the demand for routine labor and increasing the demand for non-routine labor, and to a positive productivity effect. Moll, Rachel, and Restrepo (2022) offer a quantification of automation modeled as a rise in the share of capital in production. As a source of biased technical change, automation has also been considered a potentially important factor of job polarization in advanced economies (see Goos, Manning, and Salomons (2009)). As such, it has been advocated to contribute to slow or jobless recoveries (see Jaimovich and Siu (2020) and Graetz and Michaels (2017) for challenging views). In addition, numerous questions regarding the inefficient nature of automation and its potential taxation have been raised and partly answered (see Guerreiro, Rebelo, and Teles (2021), Beraja and Zorzi (2022)).

However, most contributions on automation like Acemoglu and Restrepo (2020) or Eden and Gaggl (2018) focus on labor demand and production effects, but abstract from labor supply effects, labor-market reallocations or general equilibrium effects. A notable exception is Guerreiro, Rebelo, and Teles (2021), who build a closed-economy general equilibrium model of automation to analyze the optimal taxation of robots. However, we consider a more general CES production function that fits empirical evidence in a more compelling way, as it generates a substantial increase in the non-routine labor share, which a Cobb-Douglas production function can not do. Further, our paper considers the effect of automation in an environment with endogenous entry, which is, to the best of our knowledge, new. As such, the above papers may miss some of the effects of automation running through the relative supply of skills or through changes in business dynamism.

Our second main addition to the literature is the open-economy dimension. The latter received relatively little attention, to the exception of Artuc, Bastos, and Rijkers (2018), who investigate the trade and wage effects of automation. Artuc, Bastos, and Rijkers (2018) consider a Ricardian model of trade but disregard firms' entry in production and export markets, *i.e.* effects at the extensive margin. According to our results, both features – firm creation and entry in export markets – matter to fully assess the macroeconomic effects of automation, especially in light of the above-mentioned general equilibrium linkages through which automation proceeds. Mandelman and Zlate (2022) is more closely related to our work. They propose a model of offshoring, immigration, automation and trade in tasks to disentangle the respective contributions of these factors in explaining the polarization of U.S. labor markets.[1] We adopt a somewhat similar approach in combining various shocks to account for the observed effects of automation and grasp its welfare implications. However, rising markups and declining business dynamism appear to be salient trends of the U.S. economy that are not taken into account by Mandelman and Zlate (2022), while we highlight their critical importance in shaping the effects of automation. In

[1] The paper adopts many of the assumptions featured in Mandelman (2016).

addition, Mandelman and Zlate (2022) do not really discuss how shocks other than automation – immigration, offshoring – affect the extent of automation while this question is central in our paper.

Last, our paper also draws on the literature dealing with rising markups and declining business dynamism. De Loecker, Eeckhout, and Unger (2020) document a very large increase in market concentration and markups in the U.S. over the last 40 years. Although their finding is being discussed extensively in subsequent papers and potentially explained by changed in cost-price pass-through (Conlon et al. (2023)), by the rise of the service sector (Marto (2022)) or challenged on empirical grounds (Foster, Haltiwanger, and Tuttle (2022)), it is still considered a key ingredient to account for the dynamics of the U.S. economy, along with factors weighing on business dynamism (De Loecker, Eeckhout, and Mongey (2021)). The decline in U.S. business dynamism has been documented in details by Decker et al. (2016) and, according to Akcigit and Ates (2021), rationalized by interacting changes in market concentration and innovation, leading to decreasing knowledge diffusion (see also Akcigit and Ates (2023)). Our goal is not to explain the observed rise in markups or the decline in business dynamism but to highlight how these – arguably related – trends affect the extent of automation. Our findings suggest that these interactions are critical.

The paper is structured as follows. Section 2 presents the model. Section 3 calibrates the model using U.S. data and assesses the fit of the model. Section 4 investigates the relation between the relative price of robots and automation, and shows that the latter can be non-linear. Section 5 looks at counterfactual scenarios by shutting down other shocks. It shows that the fraction of automated tasks – automation – varies substantially: shutting down shocks with negative effects on output (markups, entry costs) raises automation while shutting down shocks with positive effects on output (labor productivity, trade costs) lowers automation. This Section also isolates the marginal effects of the degree of automation induced by other shocks to quantify its marginal effects on wage growth, labor-market reallocations and welfare. Section 6 concludes.

2 Model

We build a two-country open-economy model with endogenous entry and export participation, endogenous automation and heterogeneous labor. In each country there are two families of workers supplying differentiated labor: routine workers and non-routine workers. Routine workers are hand-to-mouth while non-routine workers have access to financial markets. The size of each family is endogenous and affected by the wage premium.

2.1 Families

We consider two countries of equal size, Home and Foreign, and denote Foreign variables with an asterisk. The Home economy represents the U.S. while the Foreign economy represents a pool of advanced economies. We focus on the Home economy (U.S.) and assume perfect symmetry in the structure and parameter values.

In the Home economy lives a family of non-routine workers (h) of size $(1+\lambda_t)$ and a family of routine workers (u) of size $(1-\lambda_t)$, in which workers are respectively paid a real wage w_t^h and w_t^u. Variable λ_t reflects structural decisions affecting the skill composition of the labor force and the relative size of families such as training or immigration, and is assumed to respond to the non-routine wage premium:

$$\lambda_t = \varpi \left(\frac{w_t^h}{w_t^u} - \frac{w^h}{w^u} \right), \tag{1}$$

Each family chooses consumption to maximize its lifetime per-capita welfare:

$$\mathcal{W}_t^i = \mathbb{E}_t \left\{ \sum_{s=t}^{\infty} \beta^{s-t} \log \left(c_s^i \right) \right\}, \; i = \{h, u\}, \; \beta < 1. \tag{2}$$

The family of non-routine workers owns firms and uses Home bonds to smooth consumption, so that welfare maximization is subject to the following budget constraint:

$$b_t + (1+\lambda_t)\, c_t^h = r_{t-1} b_{t-1} + w_t^h \chi^h (1+\lambda_t) + \Pi_t, \tag{3}$$

where χ^h is an exogenous measure of hours worked, b_t is the real value of dollar-denominated bonds, returning r_t between period t and $t+1$ and Π_t represents the firms' total profits. Maximizing (2) subject to the budget constraint implies

$$\mathbb{E}_t \left\{ \beta_{t,t+1} r_t \right\} = 1$$

where $\beta_{t,t+1} = \frac{(1+\lambda_{t+1}) u_c \left(c_{t+1}^h \right)}{(1+\lambda_t) u_c \left(c_t^h \right)}$ is a stochastic discount factor.[2] The family of routine workers does not have access to financial markets and the per-capita level of consumption of family members is given by:

$$c_t^u = w_t^u \chi^u. \tag{4}$$

For each family, the aggregate consumption basket is made of Home varieties $\omega \in \Omega$ and Foreign varieties $\omega^* \in \Omega$ where Ω denotes the space of possible varieties:

$$c_t^i = \left((n_t)^{\frac{\varsigma-1}{\theta_t}} \int_{\omega \in \Omega} c_{dt}^i (\omega)^{\frac{\theta_t - 1}{\theta_t}} d\omega + (n_{xt}^*)^{\frac{\varsigma-1}{\theta_t}} \int_{\omega^* \in \Omega} c_{xt}^i (\omega^*)^{\frac{\theta_t - 1}{\theta_t}} d\omega^* \right)^{\frac{\theta_t}{\theta_t - 1}}, \; i = \{h, u\}. \tag{5}$$

Here $\theta_t > 1$ is the elasticity of substitution between goods, n_t and n_{xt}^* respectively denote the number of Home and imported varieties available for consumption in the Home country. Following Bénassy (1996), ς captures the love for variety of consumers. Preferences boil down to

[2]While decisions to switch family type are not explicitly modeled the consequences are internalized through the subjective discount factor.

CES preferences for $\varsigma = 1$, $\varsigma < 1$ ($\varsigma > 1$) indicating a lower (higher) marginal welfare benefit from additional varieties. Defining $\rho_{dt}(\omega)$ and $\rho^*_{xt}(\omega^*)$ as the relative prices of each type of varieties, the demands for Home and imported goods are:

$$c^i_{dt}(\omega) = \rho_{dt}(\omega)^{-\theta_t} n_t^{\varsigma-1} c^i_t, \text{and } c^i_{xt}(\omega^*) = \rho^*_{xt}(\omega^*)^{-\theta_t} n_{xt}^{*\varsigma-1} c^i_t. \tag{6}$$

One difference in the problems solved by the two families in the Foreign economy is that the Foreign family of non-routine workers has access to two types of bonds, the dollar-denominated bond and a local bond. In addition, changes in the amount of dollar-denominated bonds held by Foreign non-routine workers require the payment of an adjustment cost:[3]

$$ac^{b*}_t = (\phi_b/2)\left(q_t^{-1} b^*_t - q_{t-1}^{-1} b^*_{t-1}\right)^2, \tag{7}$$

that depends on deviations of external assets from their past value, where q_t is the consumption-based real exchange rate from the perspective of the Home economy. The resulting Euler equations imply:

$$\mathbb{E}_t\left\{\beta^*_{t,t+1} r^*_t\right\} = 1, \tag{8}$$

$$\mathbb{E}_t\left\{\beta^*_{t,t+1}\left(r_t - r^*_t \frac{q_{t+1}\Gamma_t}{q_t}\right)\right\} = 0, \tag{9}$$

where $\Gamma_t = 1 + \phi_b\left(q_t^{-1} b^*_t - q_{t-1}^{-1} b^*_{t-1}\right)$ and $\beta^*_{t,t+1}$ is the counterpart of $\beta_{t,t+1}$ for Foreign non-routine workers. The first-one is a standard Euler equation on local bonds, and the second is a modified real interest rate parity condition that determines the dynamics of the real exchange rate.

2.2 Firms

Robot producers. Robots are produced in quantity x_t under perfect competition. The production cost (and thus the market price of robots) is ϕ_t, exogenous and identical across tasks. Profits from producing robots are null.

Intermediate goods. In the intermediate sector, a representative firm uses a quantity of non-routine labor ℓ^h_t and a combination of routine labor ℓ^u_t and robots x_t to produce y_t units of intermediate goods. Among automatable tasks, only a time-varying fraction $m_t \in [0, 1]$ of automatable tasks are effectively automated and performed by robots, and the remaining fraction $1 - m_t$ is performed by routine workers. As shown by Acemoglu and Restrepo (2022), the key underlying assumption is that *routine labor and robots are strong substitutes to perform automated*

[3] This assumption is purely technical and pins down the long-run level of net foreign assets as explained by Schmitt-Grohé and Uribe (2003).

tasks.[4] Guerreiro, Rebelo, and Teles (2021) build on this idea for automatable tasks and assume a Cobb-Douglas combination of automatable and non-automatable tasks. We adopt a more general CES representation, allowing both types of tasks to be complement, which results in rising demand for non-routine labor and a rising non-routine labor share in case of shocks fostering automation.[5] Our production function of intermediate goods is thus:

$$y_t = \left(\alpha\left(\psi_t \xi \ell_t^h\right)^\nu + (1-\alpha)\,\mathcal{A}_t^\nu\right)^{\frac{1}{\nu}}, \tag{10}$$

where ξ is the relative productivity of non-routine workers, ψ_t is the productivity of labor, ν determines the substitutability of routine and non-routine tasks, and $\mathcal{A}_t$ is a bundle of routine tasks:

$$\mathcal{A}_t = \left[\int_0^{m_t} (x_{it})^\mu \, di + \int_{m_t}^1 (\psi_t \ell_{it}^u)^\mu \, di\right]^{\frac{1}{\mu}}, \tag{11}$$

where μ determines the elasticity of substitution among automatable tasks. Let φ_t denote the relative price at which the intermediate good is sold to final goods producers, the representative intermediate firm maximizes its profits:

$$\Pi_{yt} = \varphi_t y_t - \phi_t \int_0^{m_t} x_{it} di - w_t^u \int_{m_t}^1 \ell_{it}^u di - w_t^h \ell_t^h. \tag{12}$$

The first-order condition with respect to non-routine labor is independent from the type of task completed and yields:

$$\alpha\left(\frac{y_t}{\psi_t \xi \ell_t^h}\right)^{1-\nu} = \frac{w_t^h}{\psi_t \xi \varphi_t}. \tag{13}$$

This equation shows that the demand for non-routine labor depend positively on the production of intermediate goods y_t and negatively on the non-routine real wage relative to the price of the intermediate good. The first-order conditions for routine labor and robots depend on whether task i is automated or not:

$$(1-\alpha)\left(\frac{\mathcal{A}_t}{x_{it}}\right)^{1-\mu}\left(\frac{y_t}{\mathcal{A}_t}\right)^{1-\nu} = \frac{\phi_t}{\varphi_t}, \text{ for } i \in [0, m_t], \tag{14}$$

$$(1-\alpha)\left(\frac{\mathcal{A}_t}{\psi_t \ell_{it}^u}\right)^{1-\mu}\left(\frac{y_t}{\mathcal{A}_t}\right)^{1-\nu} = \frac{w_t^u}{\psi_t \varphi_t}, \text{ for } i \in (m_t, 1]. \tag{15}$$

We assume $\mu \to 1$, implying perfect substitutability among routine automatable tasks. As discussed in Guerreiro, Rebelo, and Teles (2021), it follows that anything else than $w_t^u/\psi_t = \phi_t$ either yields full automation ($m_t = m = 1$ if $\phi_t < w_t^u/\psi_t$) or null automation ($m_t = m = 0$ if $\phi_t > w_t^u/\psi_t$). If m_t is to be interior, a necessary condition is thus $w_t^u/\psi_t = \phi_t$. Finally,

[4] Acemoglu and Restrepo (2022) show clearly that this type of productive transformation is different from any alternative way of modeling automation (including skill-biased or capital-biased technical change).

[5] A Cobb-Douglas production function keeps the non-routine labor share constant, which does not fit the trend observed in the past decades.

since the production technology has constant returns to scale, $\Pi_{yt} = 0$ and φ_t, the real price of the intermediate good, is also the real marginal production cost. An equilibrium with interior automation thus implies $x_{it} = \psi_t \ell^u_{jt}$ for any $i \in [0, m_t]$ and for any $j \in (m_t, 1]$. As a consequence, the degree of automation writes:

$$m_t = \frac{x_t}{x_t + \psi_t \ell^u_t}, \tag{16}$$

where $x_t = m_t x_{it}$ is the aggregate stock of robots. Accordingly, using $x_{it} = \frac{x_t}{m_t} = \frac{\psi_t \ell^u_t}{1-m_t} = \mathcal{A}_t$, the aggregate production function writes:

$$y_t = \psi_t \left(\alpha \left(\xi \ell^h_t \right)^{\nu} + (1-\alpha) \left(\frac{\ell^u_t}{1-m_t} \right)^{\nu} \right)^{\frac{1}{\nu}}, \tag{17}$$

and the first-order condition on robots implies:

$$m_t = 1 - \left(\frac{1}{\alpha} \left(\frac{\phi_t}{(1-\alpha)\,\varphi_t} \right)^{\frac{\nu}{1-\nu}} - \frac{1-\alpha}{\alpha} \right)^{\frac{1}{\nu}} \frac{\ell^u_t}{\xi \ell^h_t}, \text{ with } 0 \leq m_t \leq 1. \tag{18}$$

Equation (18) is key. It shows that the dynamics of automation critically depends on the relative price of robots ϕ_t. As the latter goes down, a first direct effect stems from the direct substitution (displacement) of routine workers in the production process, pushing m_t up. Further, as ϕ_t falls, so does the routine wage w^u_t/ψ_t in equilibrium – unless ψ_t, the productivity of labor, falls more than ϕ_t, which is not plausible. Since the demand for robots jumps, whether the demand for non-routine labor increases or not depends on whether automatable tasks and non-automatable tasks are substitutes or complements. In the case of complementarity, the demand for non-routine labor rises, which raises the non-routine wage premium. A positive response of the non-routine wage drives the marginal production cost φ_t up, which amplifies automation endogenously. This second effect is completed by a third effect stemming from general equilibrium interactions. Since households reallocate among sectors depending on the non-routine wage premium, the resulting effects on $\ell^u_t / \left(\xi \ell^h_t\right)$ act as reinforcing or dampening automation. If reallocation leads to more labor in the non-routine sector, *i.e.* if $\ell^u_t / \left(\xi \ell^h_t\right)$ falls, then automation is reinforced. This equation thus shows that modeling automation as an endogenous process delivers much richer dynamics than just assuming exogenous robot capital deepening, as the latter also depends on general equilibrium effects. As a consequence, in our framework, any shock affecting the marginal production cost or the relative amounts of labor types used in production has implications for endogenous automation.

Final goods. In the final goods sector, a continuum of heterogeneous firms differentiate intermediate goods into varieties before selling them to consumers at home and abroad. The sector allows for endogenous entry and endogenous tradability. Over the entire space of potential varieties, only a subset will actually be created and commercialized. Each firm produces one variety. Firms have specific random productivity draws z, which remain fixed once firms have been created. Entry implies a once-and-for-all sunk cost f_{et}, paid in units of intermediate goods. At each period t, there are two types of firms: n_t firms that are already productive at the beginning of the period and $n_{e,t}$ firms that are newly created – but non-productive yet. At the end of period t, a fraction $\delta \in [0,1]$ of existing firms is affected by an exogenous exit shock. The total number of varieties/firms thus evolves according to:

$$n_t = (1-\delta)(n_{t-1} + n_{et-1}). \tag{19}$$

Among the firms created, only the most productive address the export market. Entry in the export market is subject to the repeated payment of a fixed cost f_x, also paid in units of intermediate goods, and incurs the payment of iceberg melting costs τ_t.[6] Firm-specific productivity z has a Pareto distribution with lower bound $z_{\min}$ and shape parameter $\varepsilon > \theta_t - 1$. The probability density function of z is $g(z) = \varepsilon z_{\min}^{\varepsilon} / z^{\varepsilon+1}$ and the cumulative density function is $G(z) = 1 - (z_{\min}/z)^{\varepsilon}$. Over the total number of potential firms, the number of existing firms will be determined by a free-entry condition. In addition, out of the total number of firms addressing the local market, the number of exporting firms n_{xt} will be those that are productive enough to cover the additional fixed export costs and trade costs. Their number is:

$$n_{xt} = (1 - G(z_{xt}))\, n_t = (z_{\min}/z_{xt})^{\varepsilon}\, n_t, \tag{20}$$

where z_{xt} is the individual productivity of the cut-off exporting plant. Firm z produces a quantity $y_t(z)$ of variety z using $y_{mt}(z)$ of the intermediate good and the following production function:

$$y_t(z) = z y_{mt}(z). \tag{21}$$

As such, the firm-specific marginal production cost is $\varphi_t(z) = \varphi_t/z$. Let $\kappa_t(z)$ denote the total current real profits of a firm with productivity z. Total current profits comprise domestic and export profits, $\kappa_{dt}(z)$ and $\kappa_{xt}(z)$, respectively defined as:

$$\kappa_{dt}(z) = (\rho_{dt}(z) - \varphi_t/z)\, y_{dt}(z), \text{ and } \kappa_{xt}(z) = (q_t \rho_{xt}(z) - (1+\tau_t)\,\varphi_t/z)\, y_{xt}(z) - f_x \varphi_t, \tag{22}$$

where $\rho_{dt}(z)$ is the relative price of good z when sold in the Home market and $\rho_{xt}(z)$ its price expressed in terms of the Foreign currency when exported, with q_t the real exchange rate. Variables $y_{dt}(z)$ and $y_{xt}(z)$ respectively denote the Home and Foreign total demand of domestic varieties, such that firm z produces $y_t(z) = y_{dt}(z) + y_{xt}(z)$ and therefore demands $y_{mt}(z) = y_t(z)/z = (y_{dt}(z) + y_{xt}(z))/z$ units of intermediate goods. The optimal pricing conditions are derived subject to the goods demand functions and optimal prices imply:

$$\rho_{dt}(z) = \frac{\theta_t}{\theta_t - 1}\frac{\varphi_t}{z}, \text{ and } \rho_{xt}(z) = (1+\tau_t)\frac{\theta_t}{\theta_t - 1}\frac{\varphi_t}{q_t z}, \tag{23}$$

[6] Trade costs are exogenous but time-varying. Out of a quantity y produced and shipped, only $y/(1+\tau_t)$ actually arrives. Firms need to produce $(1+\tau_t)\,y$ to sell y.

Entry occurs one period before production starts and the productivity draw of the last entering firm remains fixed until the corresponding firm exits. Firms do not know their productivity draw prior entry. Hence, the entry condition equates the current entry cost, expressed in units of the intermediate good, to the total (domestic and export) discounted expected profits (starting in $t+1$) made by the average incumbent. The corresponding entry condition writes:

$$\mathbb{E}_t \left\{ \sum_{s=t+1}^{\infty} \left(\beta_{t,s}\left(1-\delta\right)\right)^{s-t} \widetilde{\kappa}_s \right\} = f_{et}\varphi_t. \tag{24}$$

where $\widetilde{\kappa}_t$ denotes average profits. Expressing the condition recursively, we get:

$$\mathbb{E}_t \left\{ \beta_{t,t+1}\left(1-\delta\right)\left(\widetilde{\kappa}_{t+1} + f_{et}\varphi_{t+1}\right) \right\} = f_{et}\varphi_t. \tag{25}$$

This equation shows the determinants of firms' entry. Given the definition of profits, entry is high when the current marginal cost is low, and when domestic and export markets are large. The entry condition also shows that entry is high when current entry costs are low or expected discounted entry costs higher than current entry costs. Among incumbents, only the most productive firms profitably enter the export market given that exporting requires the repeated payment of iceberg costs. The export productivity cut-off is thus $\kappa_{xt}\left(z_{xt}\right) = 0$ or, after using the optimal pricing and demand equations:

$$z_{xt} = \frac{1+\tau_t}{\theta_t - 1} \left(\frac{\theta_t \varphi_t}{q_t} \right)^{\frac{\theta_t}{\theta_t - 1}} \left(\frac{f_x}{n_{xt}^{\varsigma-1} \Phi_t^*} \right)^{\frac{1}{\theta_t - 1}}, \tag{26}$$

where Φ_t^* denotes the aggregate Foreign demand. As in the case of firms' entry, this equation sheds light on the determinants of entry in the export market: low trade costs, a low marginal cost, a low relative price captured by a high value of q_t, low export costs, a large Foreign market and strong love for variety. The Foreign economy is characterized by symmetric conditions that are therefore not detailed.

2.3 Aggregation and Data Consistency

We define the average productivity of Home firms selling on the domestic market as $\widetilde{z} = \nabla z_{\min}$ where $\nabla = \left(\varepsilon / \left(\varepsilon - \left(\theta_t - 1\right)\right)\right)^{\frac{1}{\theta_t - 1}}$ and the average productivity of Home firms addressing both markets as $\widetilde{z}_{xt} = \triangledown z_{xt}$ (see Ghironi and Melitz (2005) for a discussion and Melitz (2003) for proofs; see Hamano (2022) for comparable results with love for variety).

Average prices. Defining the average price of a Home good as $\widetilde{\rho}_{dt} = \rho_{dt}\left(\widetilde{z}\right)$ and the average price of Home exported good as $\widetilde{\rho}_{xt} = \rho_{xt}\left(\widetilde{z}_{xt}\right)$, we obtain real average prices:

$$\widetilde{\rho}_{dt} = \frac{\theta_t}{\theta_t - 1} \frac{\varphi_t}{\triangledown z_{\min}}, \text{ and } \widetilde{\rho}_{xt} = \left(1+\tau_t\right) \frac{\theta_t}{\theta_t - 1} \frac{\varphi_t}{q_t \triangledown z_{xt}}, \tag{27}$$

while conditions for Foreign goods are symmetrically defined and thus not reported.

Average profits and variety effect. Using the profit and pricing equations, average profits are given by:

$$\widetilde{\kappa}_{dt} = \frac{1}{\theta_t}\widetilde{\rho}_{dt}^{1-\theta_t} n_t^{\varsigma-1}\Phi_t, \text{and } \widetilde{\kappa}_{xt} = \frac{(\theta_t - 1)}{\varepsilon - (\theta_t - 1)} f_{xt}\varphi_t, \tag{28}$$

where $\Phi_t = (1+\lambda_t)\, c_t^h + (1-\lambda_t)\, c_t^u + \phi_t x_t$, with $x_t = \int_0^{m_t} x_{it} di$. These equations can be used to obtain the dynamics of average total profits:

$$\widetilde{\kappa}_t = \widetilde{\kappa}_{dt} + (n_{xt}/n_t)\widetilde{\kappa}_{xt}. \tag{29}$$

Based on the expression of the CPI, we uncover the following variety effect:

$$n_t^{\varsigma}\widetilde{\rho}_{dt}^{1-\theta_t} + n_{xt}^{*\varsigma}\widetilde{\rho}_{xt}^{*1-\theta_t} = 1. \tag{30}$$

Aggregate productivity in final goods. Endogenous changes in the number of producers and exporters affect aggregate productivity in the final goods sector. As more (less) varieties are created and/or more (less) varieties are exported, the productivity of the marginal producer (exporter) varies, even for a fixed and given distribution of productivity draws. Whether aggregate productivity in the sector is computed based on the density function $g(z)$ or based on weights reflecting relative output shares, as in Melitz (2003), entry and exit on domestic and export markets affect the aggregate productivity level through (*i*) changes in the average productivity level due to the marginal productivity of entering firms, and (*ii*) the reallocation of market shares among firms. We follow Melitz (2003) and define two measures of aggregate productivity, one for all domestic final goods producers and one for final goods producers that export:

$$Z_t = n_t^{\frac{\varsigma}{\theta_t - 1}} \bigtriangledown, \text{ and } Z_{xt} = n_{xt}^{\frac{\varsigma}{\theta_t - 1}} \bigtriangledown z_{xt}. \tag{31}$$

These equations show that aggregate productivity may increase either because the average productivity level rises, or because of a scale effect by which there are more producers on the markets. For the aggregate productivity of all domestic final goods producers, only the scale effect is present and any rise in the total number of firms will raise aggregate productivity. But as shown also in the above equations, the productivity gains are scaled by the love-for-variety parameter ς. Any mechanism that lowers the marginal production cost φ_t or favors lower entry costs will have positive effects on aggregate productivity through the entry of new firms. The strength of the effect though depends on the substitutability of varieties and on love for variety.

Market clearing. On the labor markets for routine and non-routine workers respectively, the equilibrium is:

$$\ell_t^h = \chi^h (1+\lambda_t), \text{ and } \ell_t^u = \chi^u (1-\lambda_t). \tag{32}$$

On goods markets, the clearing condition for Home intermediate goods reads:

$$y_t = \widetilde{\rho}_{dt}^{-\theta_t} \bigtriangledown^{-1} n_t^{\varsigma}\Phi_t + (1+\tau_t)\, \widetilde{\rho}_{xt}^{-\theta_t} \left(\bigtriangledown z_{xt}\right)^{-1} n_{xt}^{\varsigma}\Phi_t^* + n_{et}f_{et} + n_{xt}f_{xt}, \tag{33}$$

where $\Phi_t^* = (1+\lambda_t^*)\, c_t^{h*} + (1-\lambda_t^*)\, c_t^{u*} + ac_t^{b*} + \phi_t^* x_t^*$ and the market clearing condition for Home final goods is:

$$y_t^c = n_t^{\varsigma}\widetilde{\rho}_{dt}^{1-\theta_t}\Phi_t + q_t n_{xt}^{\varsigma}\widetilde{\rho}_{xt}^{1-\theta_t}\Phi_t^*. \tag{34}$$

On the market for dollar-denominated bonds, the international clearing condition gives:

$$q_t b_t + b_t^* = 0, \tag{35}$$

while the Foreign local bond is assumed to be in zero net supply. The dynamics of Home net foreign assets are obtained by aggregating all budget constraints and combining with market clearing conditions:

$$b_t - r_t b_{t-1} = q_t n_{xt}^{\varsigma}\widetilde{\rho}_{xt}^{1-\theta_t}\Phi_t^* - n_{xt}^{*\varsigma}\widetilde{\rho}_{xt}^{*1-\theta_t}\Phi_t. \tag{36}$$

Data consistency. In principle, as explained by Ghironi and Melitz (2005), our model variables would have to be deflated by a price index capturing the aggregate variety effect. The presence of endogenous varieties with love for variety implies that welfare-based price indices may vary even though individual product prices remain fixed. However, given that we consider a long period of time and focus on structural change, we consider that statistical offices have enough time to update price indices to take account of new varieties. We thus focus on welfare-based macroeconomic aggregates and abstract from deflating key variables.

Shocks. We consider five different driving forces capturing the following trends of the U.S. economy.

- First, we consider a persistent decline in the relative price of robots ϕ_t. In most contributions analyzing automation, this shock is the single driver of the increased use of robots in production. In our set-up however, because automation results from endogenous choices from intermediate goods producers, the relative price of robots interacts with other factors determining the marginal production cost of intermediate goods, and relative amounts of labor. Rising markups, higher entry costs, trade costs or labor productivity all contribute to the endogenous adoption of robots.

- We also include a persistent decline in the elasticity of substitution between goods θ_t, implying rising markups as documented by De Loecker, Eeckhout, and Unger (2020). This shock captures the observed upward trend in market concentration. In our set-up with endogenous entry, markups play a dual role with ambiguous net effects. On the one hand, as in most macro models, they distort the divide of value-added into profits and factor shares in favor of profits, lower net factor payments and thus lead to less production. On the other hand, these effects on the intensive margin can be compensated by changes in the extensive margin of goods production, since higher profits imply higher expected gains for firms just below the entry cut-off, and higher markups might stimulate entry and business dynamism. Hence, to account for the observed decline in business dynamism, this shock is combined with a persistent rise in entry costs. Of course changes in markups also have implications for aggregate productivity through the extensive margin, as explained above and shown by Equation (31).

- A persistent rise in the entry cost f_{et} is introduced to account for the observed decline in business dynamism discussed by Akcigit and Ates (2023). While this shock is reduced-form, it aims at capturing the dynamics of the total number of producers which, as shown again by Equation (31), has implications for aggregate productivity and thus aggregate output. What's more, the decline in business dynamism has implications for the aggregate demand of both types of labor, and thus for the extent of automation.

- In our open-economy model, aggregate productivity and output are not only affected by domestic conditions but also by the dynamics of exports and the number of exported goods. Equation (31) shows it, and we also want to take into account the rise in trade openness of the U.S. economy over the sample, from 18% in the late 80's to 30% in the early 2010's. We thus consider a persistent fall in trade costs τ_t, because it also has implications for aggregate productivity, output and automation.

- Last, we consider a persistent rise in labor productivity ψ_t. All the above shocks, when combined, might result in positive or negative aggregate output dynamics. We thus introduce the possibility of a steady rise in labor productivity to account for any potential additional rise in aggregate GDP. Introducing this shock will also help disentangle the different productivity effects at work: those pertaining to the increased use of robots, those arising from changes in existing and traded varieties of goods, and those driven by increasing labor productivity.

Contrary to the literature, since firm's entry and automation are both endogenous, all these shocks will have implications for the number of existing and traded varieties as well as for the degree of automation, and thus effects on output, productivity and trade openness. We build on existing evidence for the fall in the relative price of robots ϕ_t, and for rising markups $\theta_t/(\theta_t - 1)$. The three other shocks are set based on simulation methods to match the observed dynamics of various aggregate, as described in details below.

Solution method. Given the size of shocks and adjustments, non-linearities are likely to be significant. We thus rule out linearization methods and solve the model using an extended path algorithm, which is fully non-linear (see Fair and Taylor (1983)).[7]

3 Parameter values and data fit

We use various data sources to calibrate our model, all pertaining to the U.S. economy. A subset of parameters is calibrated using empirical evidence about their value. The remaining subset of parameters is adjusted to fit time-series evidence.

3.1 Calibrated parameters

Both countries have similar size and are symmetric in the initial equilibrium, which implies $q = 1$ and $b = 0$. All shocks will be symmetric (*i.e.* world-wide) in the baseline experiment. The model is annual, and the discount factor is $\beta = 0.96$, implying a 4% steady-state interest rate. Our baseline calibration is meant to capture the situation of the U.S. in the late 1980's, when the routine and non-routine labor shares were roughly equal. In 1987, Eden and Gaggl (2018) report a share of ICT capital in output of 2.95%, which implies assuming $m_0 = m_0^* = 0.06$ in our model, and adjusting the initial relative price of robots $\phi_0 = \phi_0^*$ accordingly. Matching the 1987 routine and non-routine shares of income (0.481 and 0.519 respectively) leads us to impose $\chi^h = 0.27$ and $\chi^u = 0.33$.

In the intermediate goods sector, the productivity of labor is initialized to $\psi_0 = 1$. Further, we follow Guerreiro, Rebelo, and Teles (2021), and impose $\alpha = 0.48$. The relative productivity parameter ξ is adjusted to match the 1987 non-routine wage premium $w_0^h / w_0^u = 1.33$, implying $\xi = 1.2983$. As explained below, ν is set along with other parameters to match the dynamics key ratios and variables. In the final goods sector, without loss of generality, we impose $n = 1$ which pins down the value of f_{e0}, the initial value of the entry cost. We also impose the initial share of exporting firms in the steady state at $n_{x0}/n_0 = 0.2$, and adjust the export cost f_x accordingly. Based on U.S. Census Bureau Business Dynamics Statistics, the annual death rate of firms is set to $\delta = 0.125$ to match an observed entry rate in 1987 of $n_{e0}/n_0 = 0.143$. Regarding the initial elasticity of substitution between varieties θ_0, we follow Ghironi and Melitz (2005) and set $\theta_0 = 3.8$, which delivers a 1.35 markup, to be compared with the 1.32 estimate reported by De Loecker, Eeckhout, and Unger (2020) for the U.S. in 1987.

For the calibration of the trade sector, we adjust the iceberg cost parameter to $\tau_0 = 2.2$ to reproduce the 18% U.S. openness ratio in 1987, measured as total trade divided by GDP using World Bank data. This number might seem high but is not far from the estimates reported by Anderson and van Wincoop (2004), and aligns well with the U.S. economy being one of the

[7]Imagine that the model is in its non-stochastic steady state and a one-time shock hits in period 1. After the unexpected shock hits, the algorithm computes the convergence path to the steady state using a fully non-linear Newton-type algorithm. Using the implied equilibrium values for period 2, suppose that a new unexpected shock hits in period 2, and that the same procedure is used to obtain the equilibrium values for period 3. The full equilibrium path can thus be computed as a succession of unexpected shocks by solving a sequence of non-linear convergence paths towards the non-stochastic steady state.

most 'closed' economies world-wide. Our calibration implies that final goods' exporters are 57% more productive than non-exporters, and that final goods' domestic prices are 97% higher than export prices. Finally, the international bond adjustment cost parameter is $\phi_b = 0.0007$ following Schmitt-Grohé and Uribe (2003).

3.2 Adjusted parameters and data fit

On top of shocks' parameters, we are left with a subset of free parameters, which do not have clear data counterpart and/or are unobservable: the labor reallocation parameter ϖ, the elasticity of substitution between factors $\nu/(\nu-1)$, the Pareto parameter ε, and the parameter governing households' love for variety ς. We proceed as follows. We feed the model with a series of repeated and persistent shocks from 1987 to 2013 on $\{\phi_t, \theta_t, f_{et}, \tau_t, \psi_t\}$, that respectively lower the unit cost of robots, raise markups, raise entry costs, lower trade costs and raise labor productivity. The size and persistence of shocks on $\{\phi_t, \theta_t\}$ are first set to match the observed fall in the depreciation-adjusted relative price of ICT capital and the observed the rise in firm's markup.

Then the size and persistence of remaining shocks $\{f_{et}, \tau_t, \psi_t\}$ and parameters $\{\varpi, \nu, \varepsilon, \varsigma\}$ are set to match the following empirical trends: (*i*) the fall in the share of routine labor in output, (*ii*) the rise in non-routine employment relative to routine employment, (*iii*) the rise in the non-routine wage premium, (*iv*) the decline in business dynamism, measured by the fall in the entry rate (n_{et}/n_t), (*v*) the rise in trade openness and (*vi*) the rise in GDP per capita.[8]

We find $\varpi = 0.5361$, $\nu = -69.14$, $\varepsilon = 3.5456$ and $\varsigma = 0.1862$. The value labor reallocation cost implies that a 10 percentage point increase in the wage premium with respect to its initial value (from 1.33 to 1.43) generates a 5.3 percentage point increase in relative non-routine labor, *i.e.* from 1 to 1.053. The production parameter ν implies a strong complementarity between routine and non-routine tasks, as the elasticity of substitution is $1/\left(1-\nu\right) = 0.0143$. In a framework where robots and routine labor are perfect substitutes, this complementarity is needed to account for a substantial rise in the non-routine labor share, consistent with what is observed – although not targeted by our procedure – in the data. The value of $\varepsilon = 3.5456$ aligns well with usual values found in the literature, and implies a bit less dispersion in firm-level productivity than Ghironi and Melitz (2005), who use a very close measure of $\varepsilon = 3.4$. Last, the parameter governing the strength of love for variety implies that households value varieties – roughly four times – less than in the typical CES case since $\varsigma = 0.1862 < 1$, in line with previous findings (see Lewis and Poilly (2012) for instance).

The shock parameters imply a lot of persistence $\rho_\phi = 0.9789$, $\rho_\theta = 0.9854$, $\rho_{f_e} = 0.9949$,

8 The share of routine labor, the relative non-routine employment, the non-routine wage premium and the relative price of ICT capital are taken from Eden and Gaggl (2018). The relative price of robots is adjusted for the observed rise in ICT capital depreciation based on a user cost approach using the depreciation rate provided by Eden and Gaggl (2018). Markups are taken from De Loecker, Eeckhout, and Unger (2020), the entry rate comes from the Census Bureau Business Dynamics Statistics, GDP per capita and trade openness are taken from World Bank data.

$\rho_\psi = 0.9477$ and $\rho_\tau = 0.995$, consistent with the idea that most shocks are close-to-permanent and that reported changes are structural. The average fall in the relative price of robot capital is 2%, which happens to match exactly the annual rate imposed by Guerreiro, Rebelo, and Teles (2021). As reported in Figure 1, between 1987 and 2013, our simulations further imply that markups rise from 1.3571 ($\theta_0 = 3.8$) to 1.5243 ($\theta_T = 2.9075$), as in the data; entry costs are almost multiplied by four – a 300% increase – which helps capture the observed decline in business dynamism; trade costs fall by roughly 50% to match trade openness; and labor productivity increases by 26%, an average 0.86% annual growth rate. The calibration is summarized in Table 1 and Figure 1 reports our baseline simulation.

Table 1: Parameter values.

Parameter	Target / Source	Value
Discount factor	Annual interest rate of 4%	$\beta = 0.96$
Labor reallocation cost	Optimized	$\varpi = 0.5361$
Non-routine labor	Non-routine labor share	$\chi^h = 0.27$
Routine labor	Routine labor share	$\chi^u = 0.33$
Non-routine labor rel. pdty	Wage premium $w_0^h/w_0^u = 1.33$ (data)	$\xi = 1.2983$
Initial cost of robots	Share of ICT capital ($m_0 = m_0^* = 0.06$)	ϕ adjusted
Non-routine labor share	Guerreiro, Rebelo, and Teles (2021)	$\alpha = 0.48$
Factor elast. of subs.	Optimized	$1/(1-\nu) = 0.0143$
Initial entry cost	Normalization ($n_0 = 1$)	f_{e0} adjusted
Export cost	Normalization ($n_{x0}/n_0 = 0.2$)	f_x adjusted
Exogenous exit rate	$n_{e0}/n_0 = 0.143$ (U.S. Census Bureau)	$\delta = 0.125$
Varieties elast. of subs.	Ghironi and Melitz (2005)	$\theta_0 = 3.8$
Pareto curvature	Optimized	$\varepsilon = 3.5456$
Love for variety	Optimized	$\varsigma = 0.1862$
Initial trade costs	Trade/GDP ratio of 18%	$\tau_0 = 2.2$
Portfolio adjustment cost	Schmitt-Grohé and Uribe (2003)	$\phi^b = 0.0007$
Persistence of ϕ_t	Optimized	$\rho_\phi = 0.9789$
Persistence of θ_t	Optimized	$\rho_\theta = 0.9854$
Persistence of f_{et}	Optimized	$\rho_{f_e} = 0.9949$
Persistence of ψ_t	Optimized	$\rho_\psi = 0.9477$
Persistence of τ_t	Optimized	$\rho_\tau = 0.995$
Shock size – robots price	Optimized	$\xi_{\phi t} = -0.0167$
Shock size – markups	Optimized	$\xi_{\theta t} = -0.0123$
Shock size – entry cost	Optimized	$\xi_{f_e t} = 0.0550$
Shock size – labor prod.	Optimized	$\xi_{\psi t} = 0.0160$
Shock size – trade costs	Optimized	$\xi_{\tau t} = -0.0302$

In Figure 1, the two first rows are our targets and their dynamics are all well matched. The third row and the two first panels of the last row report the dynamics of exogenous processes, showing the large drop in the relative price of robots, the rise in markups, the rise in entry costs, the rise in labor productivity and the fall in trade costs. Notice that trade shocks are adjusted to match openness but the resulting dynamics matches independent empirical evidence about trade costs. Finally, the last panel of the last row reports the dynamics of the non-routine labor share, which is not targeted by our procedure. While the model does not entirely capture its rise, our

Figure 1: Data vs. targeted time series.

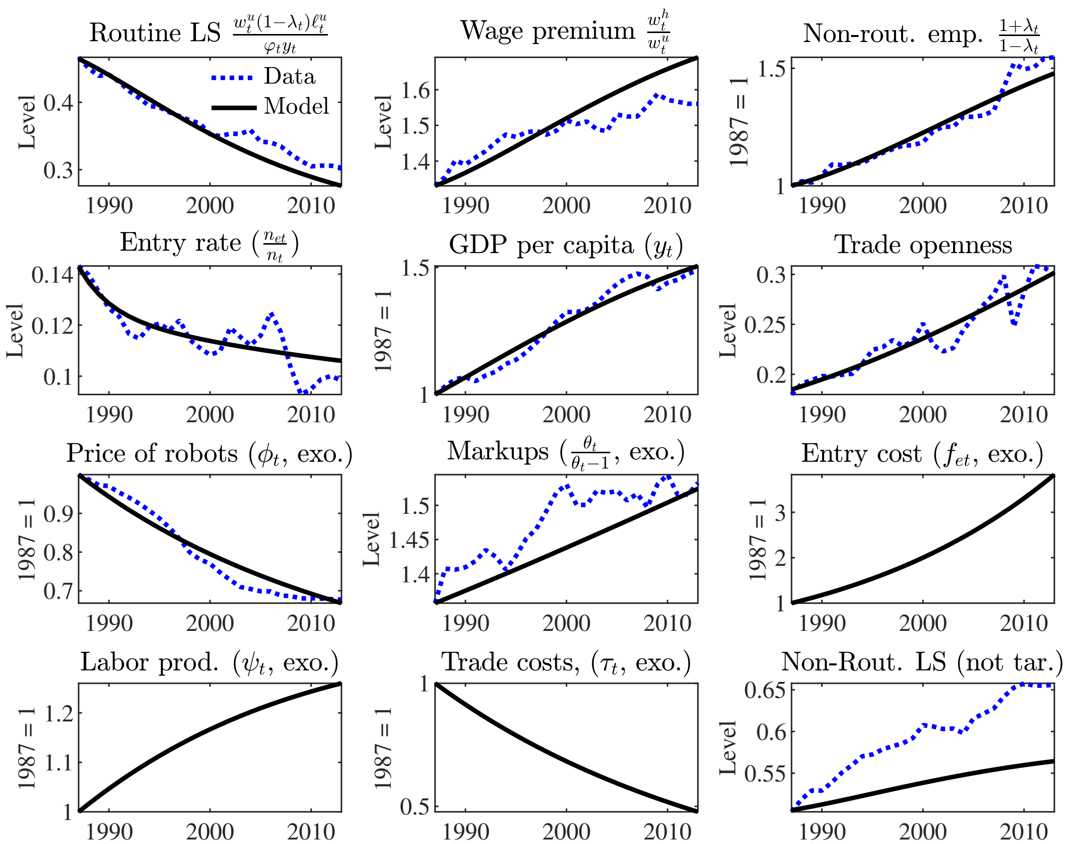

Note: Data taken from De Loecker, Eeckhout, and Unger (2020) (markups), Eden and Gaggl (2018) (routine labor share, non-routine to routine employment and adjusted price of ICT capital), Census Bureau Business Dynamics Statistics (entry rate) and World Bank (GDP per capita, trade openness).

simulation features a 6 percentage point increase.[9]

4 The relative price of robots and automation

Let us now dig further in the effects of the fall in the relative price of robots. As already mentioned, our model of endogenous automation implies that the degree of automation is not only affected by shocks to the relative price of robots. However, this price remains a key driver of automation, and a closer look at its effects illuminates the associated transmission mechanisms, as reported in Figure 2. The Figure reports the baseline simulation, a simulation where the fall in the relative price of robots is 50% less than the baseline, and a simulation with a smaller drop in robots price and with constant labor productivity.

Figure 2: Dynamics – Baseline *vs* slower automation.

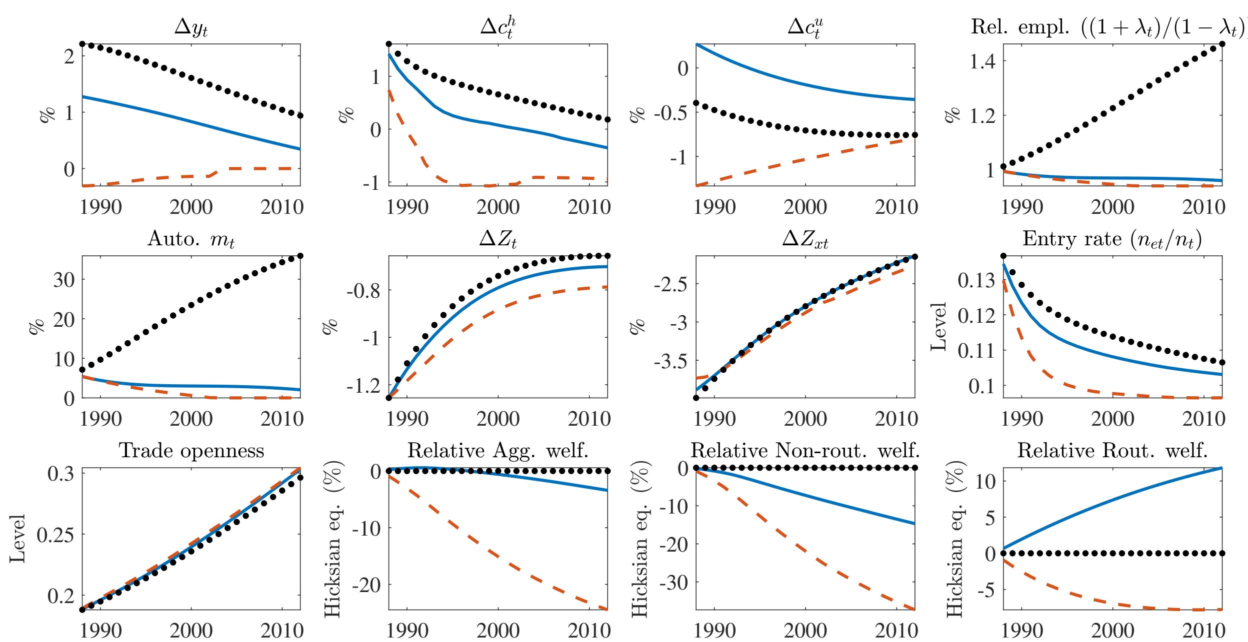

Note: Circled: baseline, Solid: smaller fall in robots price, Dashed: smaller fall in robots price and constant labor productivity. "Relative welfare" denotes welfare gains relative to the baseline case.

Let us start with the baseline case. The marginal effects of a fall in the relative price of robots can be singled out by looking at the difference between the baseline and the case of slow automation. In Figure 2, the fall in the relative price of robots is mirrored by an equivalent fall in the wage of routine workers.[10] Routine wage growth is negative and implies an equivalent

[9] A Cobb-Douglas production function subjected to automation shocks would simply reallocate the income shares between the automated and non-automated routine tasks, that is, between robot capital and routine labor, and leave the non-routine labor share constant and equal to its steady-state level.

[10] Actually the wage received by routine workers is only driven by two exogenous factors namely the relative price of robots ϕ_t and labor productivity ψ_t, since, as shown in the model $w_t^r/\psi_t = \phi_t$ in equilibrium.

negative consumption growth rate per capita for routine workers, who are hand to mouth as shown by Equation (4). Equation (18) also shows that a fall in ϕ_t – combined with other shocks – fosters automation and leads the fraction of routine tasks performed by robots to surge from 6% in 1987 to roughly 35%. Given the strong complementarity between automatable and non-automatable tasks, the more intensive use of robots is also met with a more intensive use of non-routine labor. Non-routine employment rises by more than 45% relative to routine employment in the baseline case. This joint movement in the labor force, along with a constant use of labor – the total number of workers is fixed – is what Acemoglu and Restrepo (2020) coin as the displacement effect. As a result of the increase in non-routine labor demand, non-routine workers experience strong wage growth, and thus strong consumption growth. Welfare increases for non-routine workers and decreases for routine workers. The aggregate welfare effects combine the effects stemming from per capita consumption growth with a composition effect – the increasing number of non-routine workers relative to routine workers.

Beyond the displacement effect, the fall in the relative price of robots also lowers the production cost of intermediate goods φ_t, which then results in positive output growth. In addition to the direct increase in production, the fall in the relative price of robots generates indirect productivity gains at the extensive margin by inducing more firms to enter the market. First, households' income jumps, driven by the increasing number of non-routine workers and by their rising wage. Further, the demand for robots increases. Both movements are reflected in the large increase in output. Since aggregate demand expands, expected profits rise and new firms are created, which further fuels aggregate demand, and even more final goods producers enter. Our baseline simulation suggests that the entry rate is higher in this case than with a more modest fall in robots price, resulting in more firms. Regarding the open-economy dimension, the fall in the relative price of robots does not appear to play a key role for trade openness or the export threshold *per se*. But as shown below, the open-economy dimension – the fact that trade costs fall as well – plays a key role for the transmission of the robots price shock, and for the resulting degree of automation it produces.

Inspecting the case of a smaller fall in the relative price of robots further confirms the above results about its marginal effects. Reduced robots price dynamics dampen the wage-growth rate of non-routine workers and make the wage growth rate of routine workers less negative. Compared to the baseline case, this case produces welfare losses for non-routine workers and welfare gains for routine workers. A smaller fall in the relative price of robots is also associated with a lower growth rate of output and a less dynamic entry rate, so less new firms. But this case also add a dimension to the question at hand, which is *non-linearity*. Clearly, when the relative price does not fall enough, the endogenous degree of automation m_t *falls* instead of *rising* as in the baseline case. One could think that it stems from the fall in the relative price of robots not compensating the rise in labor productivity, but actually the case of a smaller price decrease with constant labor productivity shows that it is not the case: the degree of automation remains constant and displacement effects are absent in both alternative cases. So our model predicts that the fall in the relative price of robots has to be large to trigger an endogenous adoption of robots by firms. Looking at Equation (18) reproduced here:

$$m_t = 1 - \left(\frac{1}{\alpha} \left(\frac{\phi_t}{(1-\alpha)\,\varphi_t} \right)^{\frac{\nu}{1-\nu}} - \frac{1-\alpha}{\alpha} \right)^{\frac{1}{\nu}} \frac{\ell_t^u}{\xi \ell_t^h}, \tag{37}$$

the reason can be two-fold. Given our chosen values for α and ν, either the marginal production cost φ_t falls as much as – or more than – ϕ_t, or the labor reallocation triggered by the widening of the wage premium not strong enough. Figure 3 offers an answer running an additional sensitivity exercise in which different size of the average annual decline in robots price are simulated.

Figure 3: Automation and the size of the robots price shock.

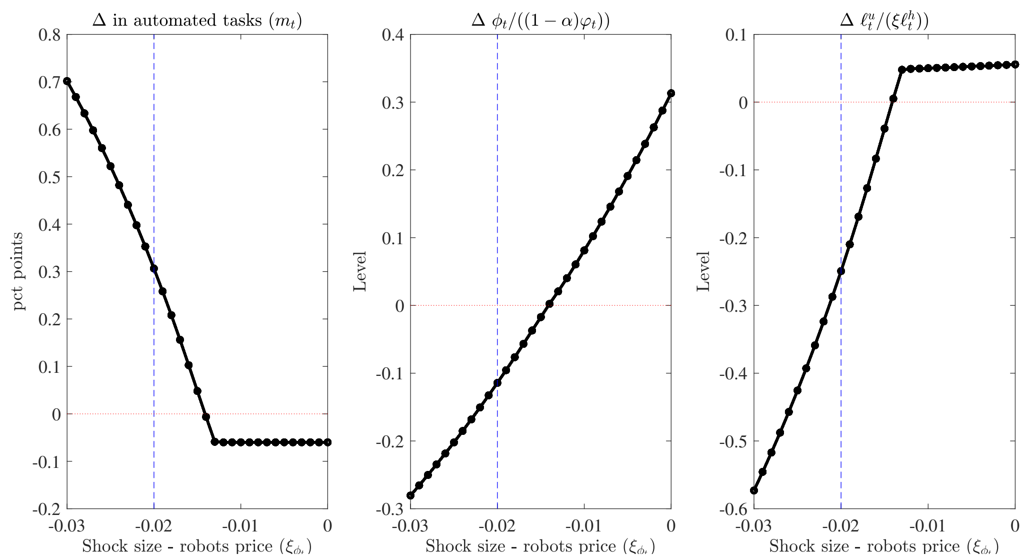

Note: The dashed blue line marks the baseline calibration.

The kinks in labor reallocation and automation coincide, confirming the role of the second factor. When the price of robots does not fall enough, the rise in the non-routine wage premium is not large enough to trigger labor reallocation. Since the latter further fuels the automation process and widens the wage premium, a feedback loop between the two phenomena arises, that requires a big push to trigger actual automation.

The above results suggest that the falling relative price of robots is a key driver of automation. It generates displacement effects leading to a rising non-routine wage premium and pushes workers to reallocate from the routine to the non-routine sector. Further, it brings positive productivity effects that lower the production cost and boost output at the intensive margin. But a decline in the price of robots also adds positive effects at the extensive margin and raises business dynamism. The resulting welfare gains fall on non-routine workers while routine workers experience welfare losses, that are basically proportional to the fall in the relative price of robots.

Since routine workers reallocate towards the non-routine sector however, the aggregate welfare gains are larger than the average of the per capita gains of both types of workers. Finally, we find that the effects of the price of robots on automation are non-linear: only large shocks actually trigger an endogenous automation process. So far our discussion has focused on the effects of robot price shocks but, as already mentioned, other shocks can affect automation and/or interact with the fall in robots price to alter the way it triggers automation.

5 Counterfactual analyzes

Let us now assess the role of other shocks in driving, amplifying or reducing the strength of automation. To do so, we simulate counterfactual trajectories of our model shutting down each shock in turn. We start our discussion with the role of shocks affecting markups and business dynamism, before turning to the specific roles played by labor productivity and trade cost shocks.

5.1 Markups, business dynamism and automation

Markups and automation Our first counterfactual simulation is one where markups remain constant instead of rising as observed in the data. The resulting effects are reported in Figure 4 below.

Figure 4: Dynamics – Baseline *vs* constant markups and entry costs.

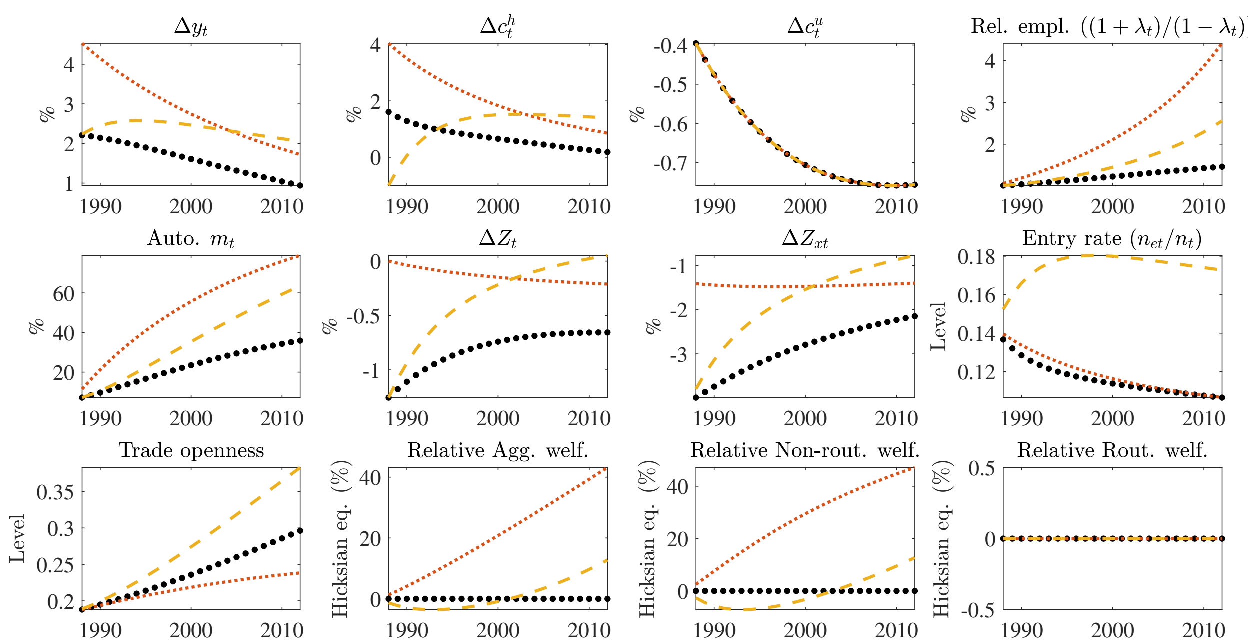

Note: Circled: baseline, Dotted: constant markups, Dashed: constant entry costs. "Relative welfare" denotes welfare gains relative to the baseline case.

In models with endogenous entry and export status, markups play a dual role. On the one hand, higher markups distort real wages and lower output at the intensive margin, which has large welfare costs. On the other hand, larger markups also imply higher expected profits and facilitate firms entry and the payment of export costs. As such, higher markups may stimulate the economy at the extensive margin, raising the number of producers and exporters. The overall welfare balance clearly depends on the relative weight placed on the number of varieties in the consumption bundle, the parameter controlling the extent of love for variety. Our estimation favors a relatively small parameter $\varsigma = 0.1862$ compared to the CES case ($\varsigma = 1$). Unsurprisingly then, our counterfactual simulation with constant markups displays massive positive effects compared to the baseline. With constant markups, output growth would have been between 1% and 2% larger every year. Driven by all the other shocks, our simulation also suggests large additional gains through enhanced automation. Since the productivity gains from automation and labor productivity are not absorbed by rising markups, production costs further fall, which reinforces the automation process. The proportion of automated tasks would have risen much faster and reached almost 70% – against 35% in the baseline, labor reallocation would have been much stronger, which is reflected in a much larger counterfactual increase in the non-routine wage premium and non-routine wage growth. In the mean time, the wage of routine workers would have remained unaffected, since the latter is determined by ϕ_t and ψ_t, which are both unchanged in this counterfactual scenario. The larger displacement and productivity effects would also have resulted in much larger welfare gains for non-routine workers, and in larger aggregate gains since labor reallocation would have further reduced the proportion of routine workers. Last, had markups not risen as in the baseline case, exporting would have become a bit less interesting – markups not only affect overall profits but also export profits – and the number of exporters would have risen less than in the baseline case, resulting in an overall dampened trade openness ratio.

Business dynamism and automation Figure 4 also contrasts the counterfactual trajectory resulting from constant – instead of rising – entry costs. In this case business dynamism would have *risen* instead of declining slowly over time. Indeed, the combined effects of the robot and labor productivity shocks with rising markups would have increased firm's entry and the number of producers and exporters. As in the case of constant markups but with a different time profile and magnitude, this counterfactual would have amplified the endogenous automation process, leading to larger reallocations on the labor market and boosting output. However, because of time-to-build in the process of firm creation, the full extent of these effects would have taken some time to fully materialize. Regarding the magnitude of the effects, annual output growth would have been up to one percentage point larger on average than in the baseline case, relative non-routine labor would have increased but less than under constant markups, and the extent of automated tasks would have reached 60%, which is above the baseline number (35%) but below the counterfactual number obtained with constant markups (70%). Finally, as the number of exporters rises along with the number of producers, trade openness would have been greater than in the baseline (35% against 30%). As in the case of constant markups, the dynamics of routine wages would not have been affected, and so the negative per capita consumption growth is identical to the baseline case, just as the welfare losses experienced by routine workers. The entirety of the additional output and productivity gains would have fallen on non-routine workers which would have resulted in much larger welfare gains for them, and larger aggregate welfare gains given their increased weight in the household sector compared to the baseline.

The marginal effects of induced automation The above counterfactual simulations shed light on the impact of factors affecting firm's environment and business dynamics in our model. Had markups or entry costs remained constant, we find that automation – and both the associated displacement and productivity effects – would have been much stronger *for a given decline in robots price*. The size and distribution of the corresponding welfare gains and losses would have been even more polarized, and the positive impact on U.S. growth would have been much larger. But how much exactly remains unclear at this stage because, with endogenous automation, the marginal effects of automation triggered by the counterfactual stability of markups and entry costs are entangled with the positive effects of constant – instead of rising – markups and entry costs. To single out the contribution of induced automation, we compute the same counterfactual trajectories with an adjusted relative price of robots that yields the same final value for automation m_t as in the baseline. The difference between the counterfactual and the adjusted counterfactual then captures the marginal effects implied by the automation induced by the absence of the shocks, and reported in Figure 5.

Figure 5 shows the marginal effects of induced automation in each case. The overall picture is already known: enhanced automation triggers displacement effects, wage growth and welfare gains for non-routine workers, negative wage growth and welfare losses for routine workers, and overall aggregate welfare gains. The positive productivity effects are reflected in the positive marginal effect on output growth, and the amplification of business dynamism translating in more producers and more exporters. Quantitatively speaking, had markups remained constant, automation would have been more than twice the baseline degree of automation at the end of the period, inducing an average annual 1.3 percentage point gain in output growth and aggregate welfare gains above 20% of consumption equivalent. Had entry costs remained constant, the marginal effects of induced automation would have been roughly equivalent to those arising under constant markups.

This exercise shows that, for a given path of the relative price of robots, the resulting degree of automation can vary more than substantially, depending on how other shocks shape labor markets, productivity or business dynamism. When automation is endogenous, negative shocks hitting the economy not only impose welfare costs on their own but also slow down the process of automation, and dampen the associated welfare and reallocation effects. We expect positive shocks, investigated in the next subsection, to have opposing – *i.e.* accelerating – effects on the endogenous extent of automation.

Figure 5: Marginal effect of automation induced by counterfactuals.

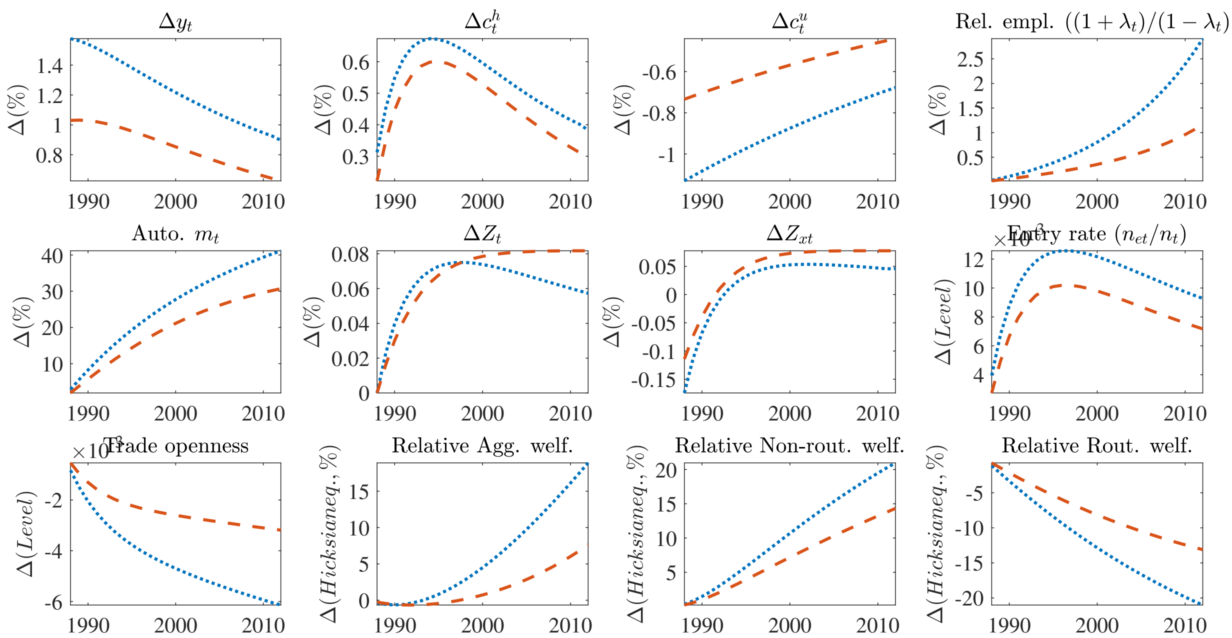

Note: Dotted: constant markups, Dashed: constant entry costs.

5.2 Labor productivity, trade costs and automation

Counterfactuals Let us now discuss counterfactual dynamics arising when shutting shocks with positive effects. We start with a counterfactual that mutes positive shocks to the productivity of labor in Figure 6.

Figure 6 reveals that the rise in labor productivity and the fall in trade costs have relatively similar and mild consequences for automation – milder than markup and entry costs. Absent one of these shocks, the proportion of automated tasks would have reached 30% instead of 35% in the baseline. The rise in the relative use of non-routine labor would have been reduced, reaching 1.3 at the end of the simulation, against 1.45 in the baseline case.

Focusing more specifically on labor productivity, Figure 6 shows that wage growth would have been much lower – and even negative – for both types of workers, resulting in large welfare losses relative to the baseline case. With constant – instead of growing – labor productivity, welfare losses are larger for non-routine workers than for routine workers. This suggests radically different welfare implications in comparison of shocks to robots price. Rising labor productivity brings welfare gains for both types and larger gains for routine workers. A declining relative price of robots produces welfare gains for non-routine workers and losses for routine workers. Last, rising labor productivity also has benefits in terms of business dynamism: absent this positive trend, the entry rate would have been substantially lower, depressing the number of producers and exporters. Regarding welfare, wage and output growth or business dynamism,

Figure 6: Dynamics – Baseline *vs* constant labor productivity and trade costs.

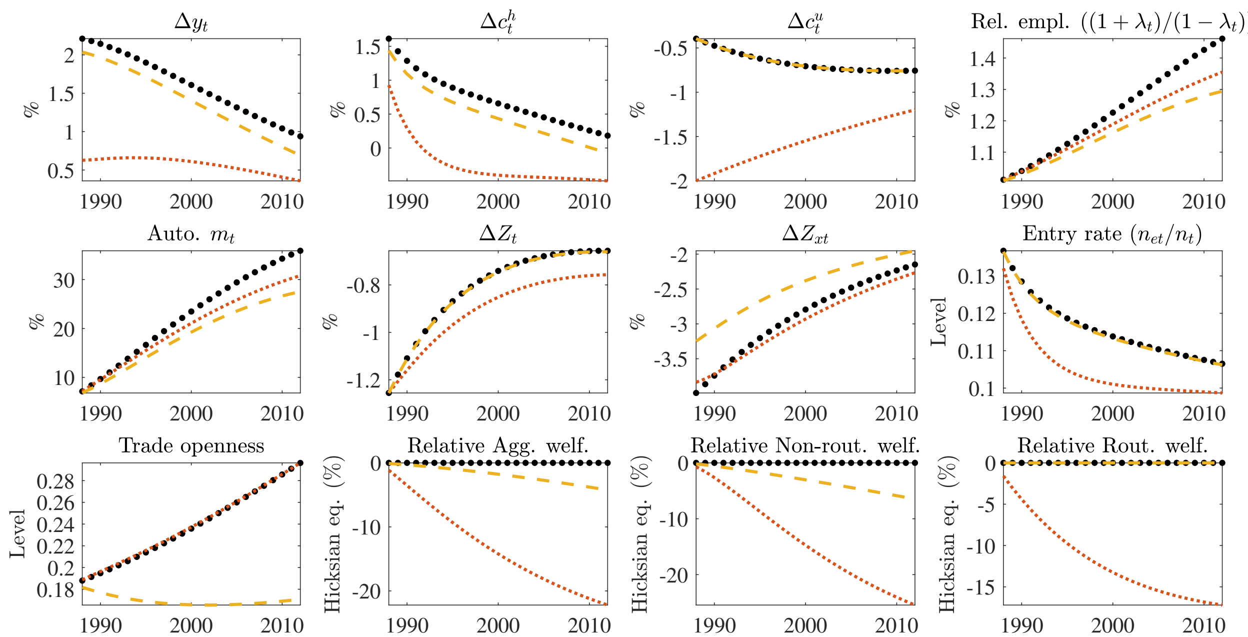

Note: Circled: baseline, Dotted: constant labor productivity, Dashed: constant trade costs. "Relative welfare" denotes welfare gains relative to the baseline case.

the counterfactual case of constant trade costs has relatively minor consequences.

Our counterfactual simulations thus highlights the positive contribution of rising productivity and falling trade costs to the adoption of robots in production: absent these shocks, automation – the fraction of automated tasks – and labor-market shifts would have been less marked by a few percentage points. It also contrasts the importance of labor productivity for aggregate welfare gains, and its key difference compared to robots price shocks, especially regarding the distribution of welfare gains/losses.

The marginal effects of induced automation As in Section 5.1, we now single out the contribution of induced automation, or, in the case of negative shocks producing lesser fractions of automated tasks, the lack thereof. Figure 7 mirrors Figure 5 in the counterfactual cases with constant labor productivity and trade costs. That is, we adjust the relative price of robots in each counterfactual case – increasing the fall in the relative price of robots – so that the end-of-period level of automation is the same as in the baseline. The difference between the counterfactual and the adjusted counterfactual gives a sense of what would have been gained or lost through endogenous automation if labor productivity or trade costs remained constant.

Had labor productivity or trade costs remained constant, the fraction of automated tasks would have ended-up below the baseline value, which highlights the positive contribution of these shocks to robots adoption. Less automation would have produced less displacement on the

Figure 7: Marginal effect of automation induced by counterfactuals (2).

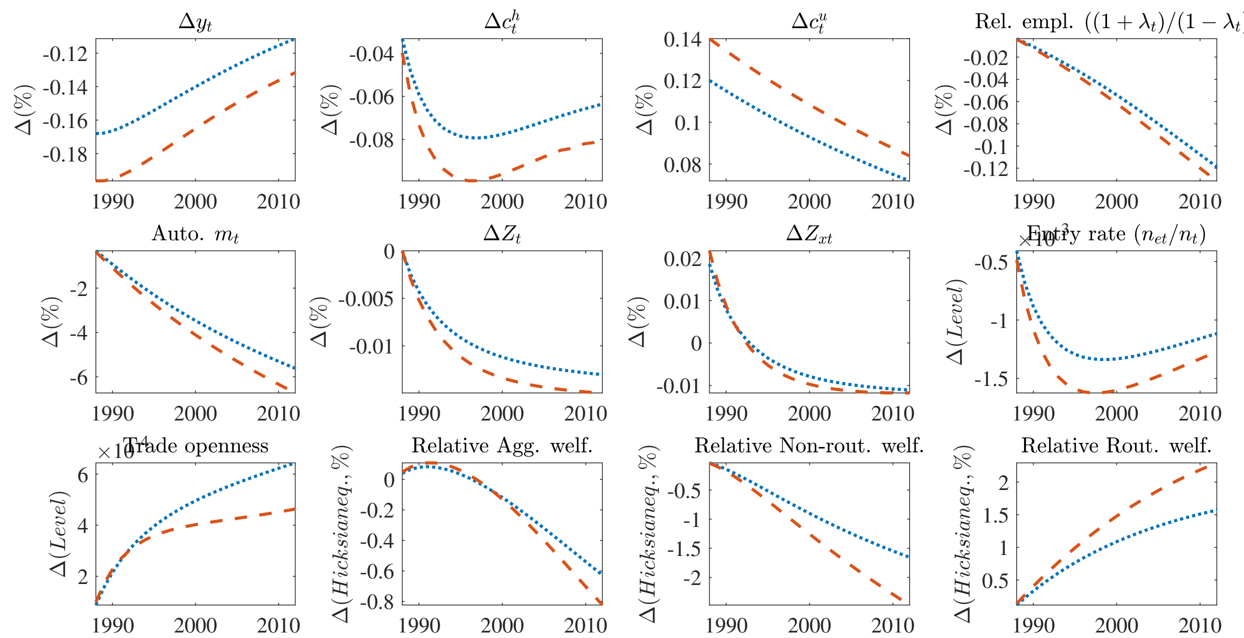

Note Dotted: constant labor productivity, Dashed: constant trade costs.

labor market, less productivity gains, and less output growth. Labor reallocation, wage growth and welfare losses for non-routine workers would have been dampened, and symmetrically, wage growth would have been higher for routine workers, generating welfare gains. Business dynamism would have been less stimulated, resulting in a lesser increase in the number of producers and exporters, which would have contributed to less output growth. The size of the effects are, however, potentially much smaller than those implied by the absence of markup and entry costs shocks.

6 Conclusion

Since the 1980's, the number of robots per thousands of industry workers has been rising very fast and it should intuitively result in large productivity gains. However, no major productivity improvement has been seen in TFP or GDP growth statistics. To match key trends in the U.S. economy, among which the rise in automation, we build an open-economy model of endogenous automation with heterogeneous firms, endogenous business creation and tradability and labor reallocation.

We show that the decline in the relative price of robots is a key factor leading to automation, and that it affects the economy non-linearly: given labor reallocation and its endogenous effects on automation, the fall in the price of robots has to be large to trigger automation. Further, factors such as rising markups, rising entry costs, rising labor productivity or declining trade costs have welfare gains/costs on their own but also affect the economy, labor markets or business dynamism through their effects on the endogenous adoption of robots by firms.

Depending on the context, the environment of firms and the other shocks affecting the economy, a given decline in robots price can lead to a very strong, moderate or small automation, resulting in correspondingly large or small productivity gains, displacement effects on the labor market and welfare outcomes. Our results thus highlight the fundamentally endogenous nature of automation, and provide a tentative answer to the apparent paradox of robots: automation and its positive effects on productivity were partly hindered by other shocks slowing the process, among which rising markup and slowing business dynamism.

References

Acemoglu, Daron and Pascual Restrepo. 2020. "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy* 128 (6):2188–2244.

———. 2022. "Demographics and Automation." *Review of Economic Studies* 89 (1):1–44.

Akcigit, Ufuk and Sina T. Ates. 2021. "Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory." *American Economic Journal: Macroeconomics* 13 (1):257–298.

———. 2023. "What Happened to U.S. Business Dynamism?" *Journal of Political Economy* Forthcoming.

Anderson, James E. and Eric van Wincoop. 2004. "Trade Costs." *Journal of Economic Literature* 42 (3):691–751.

Artuc, Erhan, Paulo S. R. Bastos, and Bob Rijkers. 2018. "Robots, Tasks and Trade." World Bank Policy Research Working Paper 8674.

Bénassy, Jean-Pascal. 1996. "Taste for Variety and Optimum Production Patterns in Monopolistic Competition." *Economics Letters* 52 (1):41–47.

Beraja, Martin and Nathan Zorzi. 2022. "Inefficient Automation." NBER Working Paper 30154.

Conlon, Christopher, Nathan H. Miller, Tsolmon Otgon, and Yi Yao. 2023. "Rising Markups, Rising Prices?" *AEA Papers and Proceedings* 113:279–83.

De Loecker, Jan, Jan Eeckhout, and Simon Mongey. 2021. "Quantifying Market Power and Business Dynamism in the Macroeconomy." NBER Working Paper 28761.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. "The Rise of Market Power and the Macroeconomic Implications." *The Quarterly Journal of Economics* 135 (2):561–644.

Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2016. "Declining Business Dynamism: What We Know and the Way Forward." *American Economic Review* 106 (5):203–07.

Eden, Maya and Paul Gaggl. 2018. "On the Welfare Implications of Automation." *Review of Economic Dynamics* 29:15–43.

Fair, Ray and John Taylor. 1983. "Solution and Maximum Likelihood Estimation of Dynamic Nonlinear Rational Expectations Models." *Econometrica* 51 (4):1169–85.

Foster, Lucia S., John C. Haltiwanger, and Cody Tuttle. 2022. "Rising Markups or Changing Technology?" NBER Working Paper 30491.

Ghironi, Fabio and Marc J. Melitz. 2005. "International Trade and Macroeconomic Dynamics with Heterogeneous Firms." *Quarterly Journal of Economics* 120 (3):865–915.

Goos, Maarten, Alan Manning, and Anna Salomons. 2009. "Job Polarization in Europe." *The American Economic Review* 99 (2):58–63.

Graetz, Georg and Guy Michaels. 2017. "Is Modern Technology Responsible for Jobless Recoveries?" *American Economic Review* 107 (5):168–73.

Guerreiro, Joao, Sergio Rebelo, and Pedro Teles. 2021. "Should Robots Be Taxed?" *The Review of Economic Studies* 89 (1):279–311.

Hamano, Masashige. 2022. "International Risk Sharing with Heterogeneous Firms." *Journal of International Money and Finance* 120:102503.

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

Lewis, Vivien and Céline Poilly. 2012. "Firm Entry, Markups and the Monetary Transmission Mechanism." *Journal of Monetary Economics* 59 (7):670–685.

Mandelman, Federico S. 2016. "Labor Market Polarization and International Macroeconomic Dynamics." *Journal of Monetary Economics* 79 (C):1–16.

Mandelman, Federico S. and Andrei Zlate. 2022. "Offshoring, Automation, Low-skilled Immigration, and Labor Market Polarization." *American Economic Journal: Macroeconomics* 14:355–389.

Marto, Ricardo. 2022. "Structural Change and the Rise in Markups." University of Pennsylvania Manuscript.

Melitz, Marc J. 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71 (6):1695–1725.

Moll, Benjamin, Lukasz Rachel, and Pascual Restrepo. 2022. "Uneven Growth: Automation's Impact on Income and Wealth Inequality." *Econometrica* 90 (6):2645–2683.

Schmitt-Grohé, Stephanie and Martín Uribe. 2003. "Closing Small Open Economy Models." *Journal of International Economics* 61 (1):163–185.