Should Robots Be Taxed?*

Joao Guerreiro† Sergio Rebelo‡ Pedro Teles§

December 2020

Abstract

Using a quantitative model that features technical progress in automation and endogenous skill choice, we show that, given the current U.S. tax system, a sustained fall in automation costs can lead to a massive rise in income inequality. We characterize the optimal tax system in this model. We find that it is optimal to tax robots while the current generations of routine workers, who can no longer move to non-routine occupations, are active in the labor force. Once these workers retire, optimal robot taxes are zero.

J.E.L. Classification: H21, O33
Keywords: inequality, optimal taxation, automation, robots.

*We thank Bence Bardóczy, Gadi Barlevy, V.V. Chari, Bas Jacobs, and Nir Jaimovich, as well as Veronica Guerrieri and three anonymous referees for their comments. Teles is thankful for the support of the FCT as well as the ADEMU project, "A Dynamic Economic and Monetary Union," funded by the European Union's Horizon 2020 Program under grant agreement 649396.

†Northwestern University.
‡Northwestern University, NBER and CEPR.
§Catolica-Lisbon School of Business & Economics, Banco de Portugal and CEPR.

1 Introduction

The American writer Kurt Vonnegut began his career in the public relations division of General Electric. One day, he saw a new milling machine operated by a punch-card computer outperform the company's best machinists. This experience inspired his novel *Player Piano*. It describes a world where children take a test that determines their fate. Those who pass become engineers and design robots used in production. Those who fail have no jobs and are supported by the government. Are we converging to this dystopian world? How should public policy respond to the impact of automation on the demand for labor?

These questions have been debated ever since 19th-century textile workers in the U.K. smashed the machines that eliminated their jobs. As the pace of automation quickens and affects a wider range of economic activities, Bill Gates reignited this debate by proposing the introduction of a robot tax.[1] Policies that address the impact of automation on the labor force have been widely discussed—for example, by the European Parliament—and have been implemented in countries such as South Korea.

In this paper, we use a model of automation to study whether it is optimal to tax robots. Our model has two types of occupations, which we call routine and non-routine. We use the word *robots* to refer to all production inputs that are complements to non-routine workers and substitutes for routine workers. So, our conclusions apply to all forms of routine-biased technical progress.[2]

To build our intuition, we first consider a simple static model in which workers have fixed occupations. In this model, a fall in the cost of automation increases

[1] Kevin J. Delaney, "The robot that takes your job should pay taxes, says Bill Gates, "Quartz, February 17, 2017, https://qz.com/911968/bill-gates-the-robot-that-takes-your-job-should-pay-taxes/.

[2] Autor, Katz and Krueger (1998), Autor, Levy and Murnane (2003), Bresnahan, Brynjolfsson and Hitt (2002), Acemoglu and Autor (2011), Goos *et al.* (2014), Cortes, Jaimovich and Siu (2017), and Acemoglu and Restrepo (2018, 2019, 2020) discuss the impact of various forms of routine-biased technical change on the labor market.

income inequality by increasing the non-routine wage premium.

If the tax system allowed for different lump-sum taxes on different workers, then technical progress would always be welfare improving since the gains could be redistributed. But these discriminatory taxes cannot be levied when the government does not observe the worker type.

For this reason, we solve for the optimal tax system imposing, as in MIRRLEES (1971), the constraint that the government does not observe the worker type or the worker's labor input. The government observes the worker's income and taxes it with a nonlinear schedule. In addition, robot purchases are also observed and taxed with a proportional tax.

In this Mirrleesian tax system, it is optimal to tax robots if the planner wants to redistribute income toward routine workers. To redistribute, the planner seeks to give positive net transfers to routine workers. However, because the tax system is the same for all workers, the non-routine workers can choose the income-consumption bundle of routine workers. This bundle can be particularly attractive for non-routine, high-wage workers because they can earn the same level of income as routine workers in just a few hours. Taxing robots reduces the non-routine wage premium, which makes the routine bundle relatively less attractive to non-routine workers. As a result, the planner can provide a better bundle to routine workers. The optimal robot tax balances these benefits of wage compression with the efficiency losses from distorting production decisions.

This rationale for positive robot taxes differs from the one proposed by Bill Gates. Gates argued that robots should be taxed to replace the tax revenue from the routine jobs lost to automation. In our model, automation increases output and overall tax revenue, so there's no need to replace taxes on routine wages.

The benchmark model that we use in our quantitative work is a dynamic model with endogenous skill acquisition. This model has an overlapping-generations structure that incorporates life-cycle aspects of labor supply. Workers have heterogeneous costs of skill acquisition and choose either a routine or non-routine occupation before they enter the labor market.[3] Once they enter the labor force, they cannot change their skill choice. They work and then retire.

The cost of producing robots falls over time as a result of technical progress. We choose parameters so that the status quo of the dynamic model is consistent with the time series for the non-routine wage premium and the fraction of the population with routine occupations in the U.S. economy.

We show that, under the current tax system, a sustained fall in the cost of automation generates a large rise in income inequality and a fall in the welfare of those who work in routine occupations.

We solve for the optimal Mirrleesian tax policy under perfect commitment. In this model, tax policy affects the skill choices made by the current newborn generation as well as future generations. For this reason, the question of whether robots should be taxed is more complex than in the static model. Initially, it is optimal for the planner to tax robots to help redistribute income toward routine workers of the initial older generations who are still in the labor force. These workers made their skill choices in the past, so those choices are not affected by the planner's generosity. In contrast, the planner gives future routine workers a less generous allocation to give them incentives to acquire non-routine skills.

Implementing this policy requires commitment. The planner treats the initial generations, which can no longer change their skill choices, differently from the generations that will be making skill choices in the future. This time dependence of the optimal commitment solution is a source of time inconsistency. At every future date, the planner would benefit from revising the optimal commitment solution. This revision would involve taxing robots to redistribute more income toward routine

[3] Our model is related to a large literature on the importance of technology-specific human capital for the diffusion of new technologies; see, for example, Chari and Hopenhayn (1991), Caselli (1999), and Adão, Beraja and Pandalai-Nayar (2018).

workers.

We find that robot taxes should be positive in the first three decades of the optimal plan. During this period, the labor force still includes older workers that chose their occupation in the past. The optimal robot tax is 5.1, 2.2, and 0.6 percent in the decades that start in 2018, 2028, and 2038, respectively. The robot tax is initially higher than the estimated effective tax rate of 1.8 percent in the status quo tax system after the 2017 tax reform. Once the initial generations retire, the optimal robot tax is zero.

The paper is organized as follows. In Section 2, we discuss the related literature. In Section 3, we describe a simple static model of automation. In Section 4, we analyze the benchmark dynamic model of automation with endogenous skill acquisition. Section 5 develops the quantitative analysis of this dynamic model. Section 6 concludes. To streamline the main text, we relegate the more technical proofs to the appendix.

2 Related literature

Our results on optimal robot taxes follow from well-known principles of optimal taxation in the public finance literature. The classic result in this literature is the production efficiency theorem of DIAMOND and MIRRLEES (1971). According to this theorem, taxing intermediate goods is not optimal even when the planner has to use distortionary taxes. Since robots are an intermediate good, our result that it is optimal to tax robots represents a failure of the production efficiency theorem.

Why does this theorem fail in our setting? The theorem requires the ability to tax net trades of different goods at different linear rates. In other words, the planner must have enough independent tax instruments to affect every relative price in the economy. In our model, this restriction means that the labor income of different types of workers can be taxed at different rates, even when those workers earn the same income. We do not allow for this form of tax discrimination. Instead, as in MIR-

RLEES (1971), we require that all worker types face the same nonlinear tax schedule. Workers can be taxed at different rates only when they earn different incomes.

Given the restriction that all workers face the same tax schedule, it can be optimal to deviate from production efficiency. But this restriction is not sufficient to justify deviating from production efficiency. ATKINSON and STIGLITZ (1976) show that production efficiency is still optimal in a MIRRLEES (1971)-type model in which labor types are perfect substitutes. In that setting, pretax relative wages are exogenous, so even if the planner does not have instruments to affect every relative price, distorting production does not help in affecting those prices to improve redistribution outcomes.

The result that, when labor types are imperfect substitutes, production efficiency may no longer be optimal was first shown by NAITO (1999), building on the work of STIGLITZ (1982) (see also SCHEUER, 2014 and JACOBS, 2015).[4] This result applies directly to our static model. Routine and non-routine workers are imperfect substitutes. Robots are substitutes for routine labor and complements to non-routine labor. By taxing robots, the planner can raise the pretax relative wage of routine workers through a general equilibrium effect.[5]

We find that taxing robots can also be optimal in the benchmark, dynamic version of our model in which workers choose whether to be routine or non-routine. In

[4]There is also a large literature that studies how the general equilibrium effects on prices and wages first emphasized by STIGLITZ (1982) affect the optimal shape of labor income taxes. This literature includes, among others, ROTHSCHILD and SCHEUER (2013), SCHEUER (2014), ALES, KURNAZ and SLEET (2015), and SACHS, TSYVINSKI and WERQUIN (2016).

[5]SCHEUER and WERNING (2016) clarify these results. Given that different levels of income in a Mirrleesian setup can be interpreted as different goods in the Diamond and Mirrlees setup, there is an equivalence between the two approaches. Since the Mirrleesian tax schedule is nonlinear, different labor incomes can be taxed at different rates. When there is a single occupation (i.e., when workers are perfect substitutes), the different goods (labor income levels) are taxed at different rates and production efficiency is optimal. Instead, with multiple occupations (i.e., when workers are imperfect substitutes), different occupations that pay the same labor income are different goods. But these different goods have to be taxed at the same rate if there is a single nonlinear income tax function. For this reason, production efficiency may cease to be optimal.

allowing for endogenous skill choice, our approach is closely related to SAEZ (2004), ROTHSCHILD and SCHEUER (2013), SCHEUER (2014), and GOMES, LOZACHMEUR and PAVAN (2018), among others. These authors characterize Mirrlees-style optimal tax plans in static models with endogenous occupation choice.

SAEZ (2004) shows that the production efficiency theorem holds in a model in which the worker chooses the occupation but labor supply is exogenous. SCHEUER (2014), instead, considers a model with an endogenous labor supply in which agents choose whether to become workers or entrepreneurs. He finds that, in the absence of differential taxation for these two occupations, the optimal plan may feature production distortions, much like the ones we have in our model.

In our setup, since workers choose their labor hours as well as their skills, both the intensive and extensive margins are potentially relevant. The robot tax is positive as long as the intensive-margin choice for the worker constrains the design of the optimal policy. If the planner needs to provide incentives only along the extensive margin, then production efficiency is optimal. In our calibrated economy, it is optimal to tax robots for the first three decades because the intensive margin is the only relevant margin for the initial old generations who cannot acquire new skills. Once these old workers retire, the optimal robot tax is zero because the only relevant margin for future young generations is skill choice.

Our results are related to the extensive literature on optimal capital taxation. This literature dates back to the seminal Chamley-Judd result that capital should not be taxed in the steady state (CHAMLEY, 1986; JUDD, 1985). WERNING (2007) extends the Chamley-Judd result to a model in which workers are heterogeneous but perfect substitutes in production. He shows that it is optimal to not distort capital accumulation both in the transition and in the steady state.

Our analysis is closest to that of SLAVÍK and YAZICI (2014), who consider optimal Mirrleesian taxation in an infinite-horizon model with low- and high-skill workers and capital-skill complementarity. They find that it is optimal to tax equipment capital in the steady state because it is a complement to high-skill workers and a substitute for low-skill workers.[6] Optimal capital taxes are high initially and rise over time. The highest capital tax rate occurs in the steady state.

Despite our different applications, the reasons for taxing equipment capital in Slavík and Yazici (2014) are very similar to the reasons why we find that robots should be taxed: the imperfect substitutability of labor types and the skill complementarity with either capital or robots. Our model differs from the one in Slavík and Yazici (2014) along two key dimensions: our analysis takes into account technical progress and endogenous skill acquisition. Because of these two elements, the reasons to deviate from production efficiency in our model cease to be relevant in the long run, so robot taxes eventually become zero.

In our model, robots are an intermediate good. We do not model robots as capital because a period represents a decade. So, there is no time to build, and robots depreciate fully. Time to build and partial depreciation are relevant for the optimal taxation of capital in ways that are not present in our model. However, if robots were modeled as a capital good, production efficiency would fail in our model for the same reason that the accumulation of robots would be distorted.[7]

In recent work, Thuemmel (2018) and Costinot and Werning (2018) also study optimal robot taxation.[8] The reason why it is optimal to tax robots in these pa-

[6] Imperfect substitutability of labor types is also why the optimal capital tax is positive in Jones, Manuelli and Rossi (1997) when the Ramsey tax system is the same for all workers.

[7] The literature on capital taxation has emphasized other motives for capital taxation that are not relevant for our analysis for the following reasons. First, Werning (2007) shows that, in a Mirrleesian setting, there is no confiscation motive for future capital taxes. Second, our preference structure and assumptions about available instruments are such that the uniform taxation results of Atkinson and Stiglitz (1972, 1976) apply. As a result, there is no reason to use capital taxes to introduce intertemporal distortions (see Chari and Kehoe, 1999, and Chari, Nicolini and Teles, 2019). Third, we do not consider idiosyncratic income risk, so the reasons to tax capital discussed by Golosov, Kocherlakota and Tsyvinski (2003) are not present (see also da Costa and Werning, 2002).

[8] Another related recent paper is Tsyvinski and Werquin (2017). These authors generalize the idea of a compensating variation to an economy with general equilibrium effects and distortionary taxation. They use their formulas to describe the optimal changes to the tax system required to compensate the effects of automation, but abstract from the possibility of taxing automation directly.

pers is essentially the same as in our work. THUEMMEL (2018) considers a static Mirrleesian economy with three occupations: non-routine cognitive, non-routine manual, and routine workers. This model generates a richer set of implications for the impact of automation on income inequality than a model with only one type of non-routine worker. THUEMMEL (2018) also considers within-occupation wage heterogeneity, which is not present in our analysis. Despite these differences, the quantitative findings in THUEMMEL (2018) are broadly consistent with ours. COSTINOT and WERNING (2018) consider a general static framework with a continuum of worker types. They derive optimal tax formulas that depend on a small set of sufficient statistics that require relatively few structural assumptions. Using empirical estimates of these statistics, they find that small, positive robot taxes are optimal. They also characterize a set of conditions under which the optimal robot tax decreases as automation progresses.

Our motivation for studying a dynamic overlapping-generations economy with skill acquisition comes in part from the work of CORTES *et al.* (2017) and ADÃO *et al.* (2018). These authors show that younger generations are more responsive to routine-biased technical progress than older generations. Using a structural model, ADÃO *et al.* (2018) find weak responses of employment shares to changes in relative wages for old generations, but very strong responses for the newer generations. This empirical result suggests that the incentives of new generations to acquire skills are important in understanding how to optimally tax robots.

3 A static model

We first consider a static model of automation to address our optimal policy questions. The model has two types of workers that draw utility from consumption of private and public goods and disutility from labor.[9] One worker type supplies routine labor and the other non-routine labor. The consumption good is produced combining both types of labor with robots. Robots and routine labor are used in a continuum of tasks.[10]

Workers There is a continuum of unit measure of workers. The index j denotes either non-routine workers, $j = n$, or routine workers, $j = r$. The fractions π_n and π_r of workers are non-routine and routine, respectively. A worker derives utility from consumption, c_j, and from the provision of a public good, G, and derives disutility from the hours of labor, l_j. The worker's utility function is

$$U_j = u(c_j, l_j) + v(G). \tag{1}$$

We assume that the first and second derivatives satisfy $u_c > 0$, $u_l < 0$, u_{cc}, $u_{ll} < 0$. We also assume that consumption and leisure are normal goods, so that $u_{lc}/u_l - u_{cc}/u_c \geq 0$, and $u_{ll}/u_l - u_{cl}/u_c \geq 0$, with one of these conditions as a strict inequality. Finally, we assume that $v_G > 0$, $v_{GG} < 0$ and that $u(c, l)$ satisfies standard Inada conditions.

Worker j chooses consumption and labor to maximize utility (1) subject to the budget constraint

$$c_j \leq w_j l_j - T(w_j l_j),$$

where w_j denotes the wage rate received by worker type j and $T(\cdot)$ denotes the income tax schedule.

Robot producers Robots are produced by competitive firms. It costs ϕ units of output to produce a robot. This cost is the same across all tasks. A representative

[9]See THUEMMEL (2018) for a static Mirrleesian economy with three worker types (non-routine cognitive, non-routine manual, and routine) and within-occupation wage heterogeneity.

[10]See AUTOR, LEVY and MURNANE (2003) for a study of the importance of tasks performed by routine workers in different industries and a discussion of the impact of automating these tasks on the demand for routine labor.

robot-producing firm chooses robot supply, X, to maximize profits: $p_X X - \phi X$. It follows that in equilibrium, $p_X = \phi$ and profits are zero.

Final good producers The representative producer of final goods hires non-routine labor (N_n) and routine labor and buys intermediate goods, which we refer to as robots. Aggregate production follows a task-based framework which has become standard in the automation literature (Acemoglu and Restrepo, 2019, 2020). There is a unit interval of tasks that can be performed by either routine labor or robots. The services produced by these tasks are denoted by y_i for each $i \in [0,1]$. The production function is given by

$$Y = A \left[\int_0^1 y_i^{\frac{\rho-1}{\rho}} di \right]^{\frac{\rho}{\rho-1}(1-\alpha)} N_n^{\alpha}, \quad \alpha \in (0,1), \rho \in [0,\infty). \tag{2}$$

Each task can be produced with n_i workers or x_i robots,

$$y_i = \begin{cases} \kappa_i x_i, & \text{if } i \text{ is automated,} \\ \ell_i n_i, & \text{if } i \text{ is not automated.} \end{cases} \tag{3}$$

The parameters κ_i and ℓ_i represent the efficiency of robots and routine labor, respectively, in task i. Without loss of generality, let κ_i/ℓ_i be weakly decreasing in i. This property implies that tasks are ordered such that routine workers are relatively more efficient in tasks indexed by higher values of i. Given this assumption, firms choose to automate the first tasks in the unit interval. We write the production function as:

$$Y = A \left[\int_0^m (\kappa_i x_i)^{\frac{\rho-1}{\rho}} di + \int_m^1 (\ell_i n_i)^{\frac{\rho-1}{\rho}} di \right]^{\frac{\rho}{\rho-1}(1-\alpha)} N_n^{\alpha}, \tag{4}$$

where m denotes the level of automation (i.e., the fraction of tasks executed by robots).

The firm's problem is to maximize profits,

$$Y - w_n N_n - w_r \int_m^1 n_i di - (1+\tau^x)\phi \int_0^m x_i di,$$

where Y is given by equation (4). The variable τ^x is the proportional tax rate on robots.

The optimal choices of N_n, x_i for $i \in [0, m]$, and n_i for $i \in (m, 1]$ require that the following first-order conditions be satisfied:

$$w_n = \frac{\alpha Y}{N_n}, \tag{5}$$

$$(1+\tau^x)\phi = \frac{(1-\alpha)Y}{x_s} \frac{(\kappa_s x_s)^{\frac{\rho-1}{\rho}}}{\int_0^m (\kappa_i x_i)^{\frac{\rho-1}{\rho}} di + \int_m^1 (\ell_i n_i)^{\frac{\rho-1}{\rho}} di}, \tag{6}$$

$$w_r = \frac{(1-\alpha)Y}{n_s} \frac{(\ell_s n_s)^{\frac{\rho-1}{\rho}}}{\int_0^m (\kappa_i x_i)^{\frac{\rho-1}{\rho}} di + \int_m^1 (\ell_i n_i)^{\frac{\rho-1}{\rho}} di}. \tag{7}$$

To simplify, we assume that $\kappa_i = \ell_i = 1$ for all i (i.e., robots and routine workers are equally productive for all tasks). This assumption lends tractability and clarity to the exposition of our results.[11] Section 4 relaxes this assumption in the context of the dynamic model.

Under this assumption, it is optimal to use the same level of routine labor, n_i, in the $1-m$ tasks that have not been automated and use the same number of robots in the m automated tasks:

$$mx_i = X, \text{ for } i \in [0, m], \text{ and } (1-m)n_i = N_r, \text{ for } i \in (m, 1], \tag{8}$$

where N_r denotes total routine hours and X denotes the total number of robots.

The optimal level of automation is zero, $m = 0$, if $w_r < (1+\tau^x)p_x$. The firm chooses to fully automate, $m = 1$, and to employ no routine workers if $w_r > (1+\tau^x)p_x$. If $w_r = (1+\tau^x)p_x$, the firm is indifferent between any level of automation $m \in [0, 1]$. In the latter case, equations (6) and (7) imply that the levels of routine labor and robots are the same across tasks.

[11] Under the assumption that $\ell_i = \kappa_i = 1$, our task-based production function coincides with the aggregate production function considered by Autor *et al.* (2003).

In the case of an interior solution for the level of automation, we find that the optimal level of automation is $m = X/(N_r + X)$. This result allows us to write the production function as $Y = A\left(X + N_r\right)^{1-\alpha} N_n^{\alpha}$.

Government The government chooses taxes and the optimal level of government spending in order to satisfy the budget constraint

$$G \leq \pi_r T(w_r l_r) + \pi_n T(w_n l_n) + \tau^x p_x X. \tag{9}$$

Equilibrium An equilibrium is a set of allocations $\{c_r, l_r, c_n, l_n, G, N_r, X, x_i, n_i, m\}$, prices $\{w_r, w_n, p_x\}$, and a tax system $\{T(\cdot), \tau^x\}$ that: (i) solves the workers' problem given prices and taxes; (ii) solves the firms' problem given prices and taxes; (iii) satisfies the government budget constraint; and (iv) satisfies market clearing.

The market-clearing conditions for routine and non-routine labor are

$$N_j = \pi_j l_j, \quad j = n, r, \tag{10}$$

and the market-clearing condition for output is

$$\pi_r c_r + \pi_n c_n + G \leq Y - \phi X. \tag{11}$$

The equilibrium with interior automation In an equilibrium with automation, the wage rate of routine workers equals the cost of robot use: $w_r = (1 + \tau^x)\phi$. This condition implies that the number of robots used in each automated task equals the number of routine workers used in each non-automated task:

$$\frac{X}{m} = \frac{\pi_r l_r}{1 - m}.$$

Combining this equation with the firm's first-order condition (6), we obtain

$$(1 + \tau^x)\phi = (1 - \alpha)(X + \pi_r l_r)^{-\alpha}(\pi_n l_n)^{\alpha}. \tag{12}$$

Finally, replacing $X = m\pi_r l_r/(1-m)$ in equation (12), we find that the equilibrium level of automation satisfies

$$m = 1 - \left[\frac{(1+\tau^x)\phi}{(1-\alpha)A}\right]^{1/\alpha} \frac{\pi_r l_r}{\pi_n l_n}. \tag{13}$$

Furthermore, using equations (5) and (6), we find that the wages of both non-routine and routine labor are given by technological parameters and τ^x:

$$w_n = \alpha A^{1/\alpha}\left[\frac{1-\alpha}{(1+\tau^x)\phi}\right]^{\frac{1-\alpha}{\alpha}}, \tag{14}$$

$$w_r = (1+\tau^x)\phi. \tag{15}$$

The wage of routine workers is determined by the after-tax cost of robots. Because of constant returns to scale, the ratio of inputs is pinned down, as is the wage of the non-routine worker. An increase in τ^x raises the wage of routine workers and lowers the wage of non-routine agents.

Production net of the cost of robots is given by

$$Y - \phi X = \pi_n w_n l_n \frac{\tau^x + \alpha}{\alpha(1+\tau^x)} + \frac{\pi_r w_r l_r}{1+\tau^x}. \tag{16}$$

It is useful to note that the shares of routine and non-routine income in total production are

$$\frac{w_r \pi_r l_r}{Y} = (1-\alpha)(1-m) \quad \text{and} \quad \frac{w_n \pi_n l_n}{Y} = \alpha.$$

An increase in automation reduces the income share of routine workers in total production and leaves the share of non-routine workers unchanged. As the economy approaches full automation, non-routine workers earn all labor income. In this sense, an increase in automation leads to an increase in pretax income inequality.

3.1 Status quo equilibrium in the static model

We now compute the status quo equilibrium in the static model. For simplicity, we assume that robot taxes are zero ($\tau^x = 0$).[12] We model the income tax system using the functional form for U.S. after-tax income proposed by Feldstein (1969), Persson (1983), and Benabou (2000) and estimated by Heathcote *et al.* (2017). In this specification, the income tax paid by worker j is given by

$$T(w_j l_j) = w_j l_j - \lambda (w_j l_j)^{1-\gamma}, \tag{17}$$

where $\gamma < 1$. The parameter λ controls the level of taxation—higher values of λ imply lower average taxes. The parameter γ controls the progressivity of the tax code. When γ is positive, the average tax rate rises with income, so the tax system is progressive.

To illustrate the properties of the status quo equilibrium in closed form, we assume that the utility function is given by

$$u(c_j, l_j) + v(G) = \log(c_j) - \zeta \frac{l_j^{1+\nu}}{1+\nu} + \chi \log(G). \tag{18}$$

These preferences, which are also used in Ales, Kurnaz and Sleet (2015) and Heathcote *et al.* (2017), have two desirable properties: they are consistent with balanced growth and with the empirical evidence reviewed in Chetty (2006).

For these preferences and the status quo tax specification, the equilibrium is easily computed. Worker optimality implies that hours worked are constant and depend on the preference parameters ζ and ν and the progressivity parameter γ:

$$l_j = \left(\frac{1-\gamma}{\zeta}\right)^{\frac{1}{1+\nu}} \equiv \ell. \tag{19}$$

[12] In Section 5, we calibrate the baseline dynamic model with positive robot taxes equal to 3.8 percent before the 2017 tax reform and 1.8 percent after this reform.

Consumption of worker type j is equal to

$$c_j = \lambda(w_j \ell)^{1-\gamma}. \tag{20}$$

This property implies that the ratio of consumption of routine and non-routine workers is

$$\frac{c_r}{c_n} = \left(\frac{w_r}{w_n}\right)^{1-\gamma} = \frac{\phi^{\frac{1-\gamma}{\alpha}}}{\left[\alpha A^{1/\alpha}(1-\alpha)^{\frac{1-\alpha}{\alpha}}\right]^{1-\gamma}} \tag{21}$$

and that the equilibrium level of automation is

$$m = 1 - \left[\frac{\phi}{(1-\alpha)A}\right]^{1/\alpha} \frac{\pi_r}{\pi_n}. \tag{22}$$

We assume that government spending is a fraction χ of aggregate consumption, so government spending grows over time. This assumption is natural since, given the form of the utility function, the optimal ratio of government spending to consumption is χ. We also assume that tax progressivity, γ, is constant and that the government adjusts λ to maintain a balanced budget. The resulting value of λ is

$$\lambda = \frac{1}{1+\chi} \frac{\sum_{j=n,r} \pi_j w_j \ell}{\sum_{j=n,r} \pi_j (w_j \ell)^{1-\gamma}}. \tag{23}$$

To investigate the impact of technical progress, we compute the equilibrium effects of a marginal increase in ϕ^{-1}, corresponding to a fall in the robot production cost, ϕ.

As robots become cheaper, pretax labor income rises for non-routine workers and falls for routine workers:

$$\frac{d\log(w_n \ell)}{d\log \phi^{-1}} = \frac{1-\alpha}{\alpha} \quad \text{and} \quad \frac{d\log(w_r \ell)}{d\log \phi^{-1}} = -1. \tag{24}$$

This divergence is associated with an increase in the number of tasks that are automated by replacing routine workers with robots:

$$\frac{d\log(1-m)}{d\log \phi^{-1}} = -\frac{1}{\alpha}. \tag{25}$$

Higher pretax income inequality leads to higher consumption inequality:

$$\frac{d \log c_n / c_r}{d \log \phi^{-1}} = \frac{1 - \gamma}{\alpha}. \tag{26}$$

When income taxes are progressive ($\gamma > 0$), consumption inequality rises by less than pretax income inequality.

The impact of technical progress on individual consumption depends on the response of pretax income and also on how the parameter that controls the level of taxation, λ, adjusts. Interestingly, λ rises as technical progress rises.

To further illustrate the properties of the model, we parameterize this model using the calibration of the dynamic model in Section 5.

Figure 1: The status quo equilibrium in the static economy

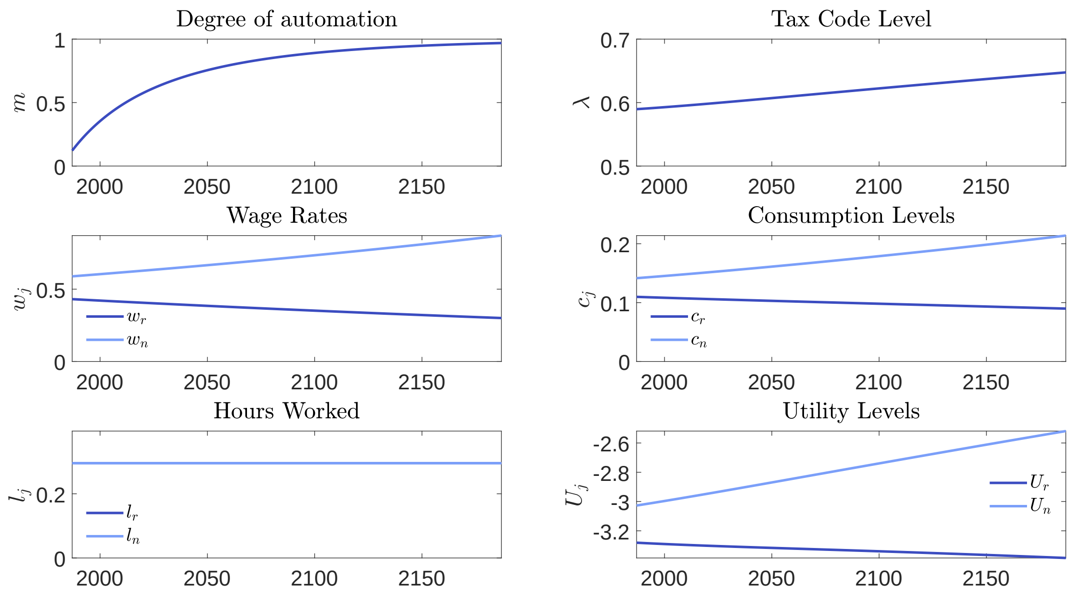

Notes: This figure illustrates the properties for the status quo equilibrium for a sequence of economies in the static model. The parameters are chosen to be consistent with the calibration of the quantitative model in Section 5. The first and second panels show the levels of automation, m_t, and the tax code level, λ_t, respectively. The third and fourth panels plot the wages and consumption levels of non-routine and routine workers. Finally, the fifth and sixth panels plot the level of labor supply and utility for each worker type.

In Figure 1, we consider a sequence of static economies in which the cost of producing robots falls according to $\phi_t = \widetilde{\phi} e^{-g_\phi (t-1)}$ ($t = 1$ corresponds to 1988).[13] As the cost of robots falls over time, the consumption of non-routine workers rises and the consumption of routine workers falls. The cost of robots converges to zero asymptotically, driving the consumption of routine workers toward zero.

In sum, our analysis suggests that, under the current U.S. tax system, a fall in automation costs will lead to massive income and welfare inequality.

3.2 Optimal taxation in the static model

It is useful to briefly consider the allocation that maximizes welfare subject only to technological constraints. Implementing this *first-best* allocation requires setting agent-specific lump-sum taxes.

We assume that the social welfare function is a weighted average of individual workers' utilities. The weights on the social welfare function, ω_n and ω_r for non-routine and routine agents, respectively, are normalized so that $\pi_r \omega_r + \pi_n \omega_n = 1$. The planner's problem is to choose allocations to maximize social welfare,

$$\mathcal{W} \equiv \pi_r \omega_r \left[u(c_r, l_r) + v(G) \right] + \pi_n \omega_n \left[u(c_n, l_n) + v(G) \right], \tag{27}$$

subject only to the economy's resource constraint.

The first-best allocation always features production efficiency. This property implies that the marginal productivity of robots equals their marginal cost, ϕ, so the robot tax is zero. As we have seen, without taxes on robots, a fall in ϕ leads to an increase in pretax wage inequality. However, since the first best features unrestricted taxes/transfers, it is always possible to redistribute income without creating distortions. As a result, pretax wage inequality does not constrain redistribution, and both types of workers benefit from technical progress.

[13] We start our analysis in 1988 and assume that there was zero automation prior to this period. This assumption is consistent with the analysis in Acemoglu and Restrepo (2019, 2020).

In general, the first-best solution cannot be implemented if the planner cannot discriminate between worker types. To see the intuition for this result, consider the case in which $\omega_n = \omega_r$ and the workers' utility function is separable in consumption and leisure. In this case, routine and non-routine workers have the same level of consumption, but non-routine workers work longer hours than routine workers. As a consequence, non-routine workers would have an incentive to act as routine to obtain a more generous consumption and leisure bundle.

In what follows, we consider a restricted planning problem. We show that if the planner cannot discriminate across worker types, then pretax wage inequality becomes relevant in order to determine how much redistribution can be done.

Mirrleesian optimal taxation In this section, we characterize the nonlinear income tax schedule that maximizes social welfare when the planner observes a worker's total income but does not observe the worker's type or labor supply, as in MIRRLEES (1971).

We focus on the case in which the level of automation is interior, $m > 0$.[14] We also assume that $\phi \leq \alpha^{\alpha}(1-\alpha)^{1-\alpha}A$, so that if $\tau^x \leq 0$, non-routine workers earn a higher wage than routine workers, $w_n \geq w_r$ (see equations (14) and (15)).

The Mirrleesian planning problem is to choose the allocations $\{c_j, l_j\}_{j=n,r}$, G, and the robot tax τ^x to maximize social welfare, (27), subject to the resource constraint,

$$\pi_r c_r + \pi_n c_n + G \leq \pi_n w_n l_n \frac{\tau^x + \alpha}{\alpha(1+\tau^x)} + \frac{\pi_r w_r l_r}{1+\tau^x}, \tag{28}$$

and two incentive constraints (IC),

$$u(c_n, l_n) \geq u\left(c_r, \frac{w_r}{w_n} l_r\right), \tag{29}$$

$$u(c_r, l_r) \geq u\left(c_n, \frac{w_n}{w_r} l_n\right). \tag{30}$$

[14] When $m = 0$, this simple model is a special case of the one considered in STIGLITZ (1982).

The wages of the two types of workers are given by equations (14) and (15). The conditions (28), (29), and (30) are necessary and sufficient to describe a competitive equilibrium. We discuss these properties in the appendix.

In MIRRLEES (1971)'s model, the productivities of the different agents are exogenous. ATKINSON and STIGLITZ (1976) show that production efficiency is optimal in that environment. Our model instead features endogenous productivities that depend on τ^x. This property turns out to be central to the question we are interested in studying: whether it is optimal to tax robot use, distorting production, in order to redistribute income from non-routine to routine workers. Based on the work of STIGLITZ (1982) and NAITO (1999), who first considered the impact of endogenous productivities in the design of the optimal tax system, we should expect production efficiency to no longer be optimal. That is indeed the case in our model. As long as automation is interior, robot taxes are positive in our model, as stated in Proposition 1.

The expression for net output on the right-hand side of equation (28) can be written as

$$\frac{\tau^x + \alpha}{\alpha(1+\tau^x)^{1/\alpha}} \frac{\alpha A^{1/\alpha}(1-\alpha)^{\frac{1-\alpha}{\alpha}}}{\phi^{\frac{1-\alpha}{\alpha}}} \pi_n l_n + \phi \pi_r l_r. \tag{31}$$

The term $(\tau^x + \alpha) / \left[\alpha\left(1+\tau_x\right)^{1/\alpha}\right]$ is equal to one for $\tau^x = 0$ and strictly less than one for $\tau^x \neq 0$. This term is a measure of the production inefficiency created by the robot tax. Fixing labor supplies, l_n and l_r, a zero robot tax maximizes the level of production.

Proposition 1 shows that, when automation is incomplete, the planner is willing to bear a resource cost, in terms of production inefficiency, to loosen the incentive constraint. In this proposition, we characterize the optimal allocation under the assumption that the planner wants to redistribute to routine workers to an extent such that the incentive constraint of the non-routine worker binds and the incentive constraint of the routine worker is slack. This approach is standard in the literature.

Proposition 1. *Suppose the optimal allocation is such that the incentive constraint binds for non-routine workers and does not bind for routine workers. Then, if automation is incomplete ($m < 1$), optimal robot taxes are strictly positive ($\tau^x > 0$).*

This proposition is proved in the appendix. The intuition for this result is that, starting from any robot tax that is less than or equal to zero, $\tau^x \leq 0$, there are welfare gains from increasing this tax rate. First, suppose that the tax on robots is strictly negative, $\tau^x < 0$. In this case, a marginal increase in τ^x has two benefits. First, for given levels of the labor supplies, it strictly increases output and hence the amount of goods available for consumption. Second, it reduces the non-routine wage premium, w_n/w_r, and makes the non-routine worker less inclined to mimic the routine worker. This property can easily be seen from the incentive constraint of the non-routine worker, (29).

Suppose instead that the robot tax is zero, $\tau^x = 0$. Since, for a given level of the labor supplies, the value of τ^x maximizes output, a marginal increase in that tax produces only second-order output losses. On the other hand, increasing τ^x generates a first-order gain from loosening the incentive constraint. Therefore, starting from $\tau^x = 0$, the planner can always improve welfare with a marginal increase in τ^x.

Robot taxes are optimal only when automation is incomplete ($m < 1$), so that routine workers are employed in production ($l_r > 0$). When full automation is optimal ($m = 1$ and $l_r = 0$), there are no informational gains from taxing robots. Since the robot tax distorts production and does not help to loosen the incentive constraint of the non-routine agent, the optimal value of τ^x is zero (see the appendix for a proof).

4 A dynamic model

In this section, we study the optimal tax policy in a model with endogenous skill acquisition. We consider an overlapping-generations model in which workers choose their occupation when they enter the workforce, work in the following periods, and then retire.

For computational reasons, we assume that each period represents a decade. Agents live for six decades, working in the first four and retiring in the last two. We assume that robots produced at time t can immediately be used in production, so there is no time to build, and robots depreciate fully within the period.[15]

As in the static model, technical change reduces the cost of producing robots over time. Because robots are better substitutes for routine than for non-routine workers, technical change is biased toward non-routine skills and increases the non-routine wage premium. This effect is analogous to the impact on the skill premium of technical change with capital-skill complementarity discussed in KRUSELL, OHANIAN, RÍOS-RULL and VIOLANTE (2000).

Workers and preferences Time is discrete with an infinite horizon $t = 1, 2, \ldots$. For simplicity, we assume that each generation is composed of a unit measure of workers. Workers live for L periods and work for $L_w \leq L$ periods. We use $a \in \{0, \ldots, L-1\}$ to denote a worker's age: $a = 0$ denotes the first period of life and $a = L - 1$ the final period.

Workers born before the initial date, $t = 1$, enter the economy at age $\widetilde{a}$, where $1 \leq \widetilde{a} \leq L - 1$. In the next period, $t = 2$, their age is $a = \widetilde{a} + 1$, then $a = \widetilde{a} + 2$, and so on. The initial older generations cannot acquire new skills. A share $\pi_{r,1-a}$ of these workers are routine and $\pi_{n,1-a}$ are non-routine. We denote the consumption and labor supply of those workers in occupation j and age a at time t by $c^a_{j,t}$ and $l^a_{j,t}$, respectively. Workers value streams of consumption, government spending, and leisure according to the utility function

$$U_{j,1-\widetilde{a}} \equiv \sum_{a=\widetilde{a}}^{L-1} \beta^{a-\widetilde{a}} \left[u\left(c^a_{j,1+a-\widetilde{a}}\right) + v(G_{1+a-\widetilde{a}}) \right] - \sum_{a=\widetilde{a}}^{L_w-1} \beta^{a-\widetilde{a}} \psi\left(l^a_{j,1+a-\widetilde{a}}\right), \tag{32}$$

[15]The standard modeling of capital accumulation embodies a one-period time to build: capital goods used in production at time t are produced at time $t-1$. This formulation would introduce a time-to-build period of ten years in robot production, which seems unreasonably long.

where $U_{j,1-\tilde{a}}$ denotes their utility level and β is the subjective discount factor. We assume that the utility function is separable in consumption and labor and satisfies the standard assumptions about monotonicity, concavity, and Inada conditions.

Workers born in period $t \geq 1$ have heterogeneous utility costs of skill acquisition, $\theta \in \Theta$. We assume that these costs follow a distribution H with continuous probability density function h. We denote by $c^a_{\theta,t}$ the consumption of a worker with skill cost θ and age a at time t, and by $l^a_{\theta,t}$ the worker's labor supply. The lifetime utility of a worker born in period t, without including skill acquisition costs, is

$$U_{\theta,t} \equiv \sum_{a=0}^{L-1} \beta^a \left[u(c^a_{\theta,t+a}) + v(G_{t+a})\right] - \sum_{a=0}^{L_w-1} \beta^a \psi\left(l^a_{\theta,t+a}\right). \tag{33}$$

This worker's overall utility is equal to lifetime utility net of the costs of skill acquisition, $U_{\theta,t} - \theta s_{\theta,t}$. The indicator function $s_{\theta,t} \in \{0,1\}$ denotes the worker's skill choice, where $s_{\theta,t} = 0$ denotes routine skills and $s_{\theta,t} = 1$ denotes non-routine skills.

Workers with positive values of θ face a positive cost of acquiring non-routine skills, which means that, all else equal, they would prefer to acquire routine skills. Workers with negative values of θ prefer, all else equal, to acquire non-routine skills. We denote by $\Theta_{r,t}$ and $\Theta_{n,t}$ the subsets of Θ that correspond to the choice of routine and non-routine occupations, respectively; that is, $\Theta_{n,t} \equiv \{\theta : s_t(\theta) = 1\}$ and $\Theta_{r,t} \equiv \Theta - \Theta_{n,t}$.

Throughout, we use $\pi_{n,t} \equiv \int_{\Theta_{n,t}} h(\theta) d\theta$ to denote the share of non-routine workers in the newborn population at time t and $\pi_{r,t} \equiv 1 - \pi_{n,t}$ to denote the share of routine workers in the newborn population at time t.

In the competitive equilibrium, workers choose their consumption, labor supply and savings on risk-free bonds to maximize their utility. Workers pay taxes and receive government transfers.

Firms and technology Robots and final output are produced by competitive firms using the same production technology as in the static model. Robots cost ϕ_t units of output to produce at time t. Final output is produced according to (4), so the elasticity of substitution between total tasks and non-routine labor is equal to one. This property is important in order to ensure the existence of a balanced-growth path, which is reached asymptotically.

A representative final goods firm maximizes per-period profits by choosing how much to produce, how much labor to hire, and how many robots to buy. The firm hires non-routine labor at the wage rate $w_{n,t}$, hires routine labor at the wage rate $w_{r,t}$, and pays the robot cost gross of taxes, $(1+\tau_t^x)\phi_t$. The first-order conditions for this profit maximization problem for each period t are the analog of the first-order conditions (5)-(7).

Without loss of generality, suppose that tasks are ordered so that κ_i/ℓ_i is weakly decreasing in $i \in [0,1]$. This property implies that routine workers are relatively more efficient in tasks indexed by higher values of i. Given this assumption, the firm uses robots in the first m_t tasks and routine workers in the final $1-m_t$ tasks. The optimal allocation of routine workers and robots to each of those tasks is described by the same first-order conditions as in the static model. These conditions imply that

$$x_{i,t} = \frac{\kappa_i^{\rho-1}}{\int_0^{m_t} \kappa_j^{\rho-1} dj} X_t, \; i \in [0, m_t] \quad \text{and} \quad n_{i,t} = \frac{\ell_i^{\rho-1}}{\int_{m_t}^1 \ell_j^{\rho-1} dj} N_{r,t}, \; i \in (m_t, 1].$$

Following ACEMOGLU and RESTREPO (2019, 2020), we replace these expressions in the production function and obtain

$$Y_t = A \left[\left(\int_0^{m_t} \kappa_i^{\rho-1} di \right)^{\frac{1}{\rho}} X_t^{\frac{\rho-1}{\rho}} + \left(\int_{m_t}^1 \ell_i^{\rho-1} di \right)^{\frac{1}{\rho}} N_{r,t}^{\frac{\rho-1}{\rho}} \right]^{\frac{\rho}{\rho-1}(1-\alpha)} N_{n,t}^{\alpha}$$

The firm's optimal choice of the level of automation implies that

$$\frac{X_t}{N_{r,t}} = \frac{\int_0^{m_t} \kappa_i^{\rho-1} di}{\int_{m_t}^1 \ell_i^{\rho-1} di} \frac{\ell_{m_t}^{\rho}}{\kappa_{m_t}^{\rho}}.$$

At this level of generality, we cannot solve for m_t in closed form. As in CHEN (2019), we make the analysis more tractable by introducing the following assumption.

Assumption 1. *$\kappa_i = \varsigma i^{\frac{\varepsilon-1}{\varepsilon}}$ and $\ell_i = \varsigma(1-i)^{\frac{\varepsilon-1}{\varepsilon}}$ where $\varsigma = \left[1 + \frac{(\varepsilon-1)(\rho-1)}{\varepsilon}\right]^{\frac{1}{\rho-1}}$ and $(1-\varepsilon)(\rho-1)/\varepsilon < 1$.*

Under this assumption, the optimal value of m_t is given by

$$m_t = \frac{X_t^\varepsilon}{X_t^\varepsilon + N_{r,t}^\varepsilon}. \tag{34}$$

The routine tasks aggregator becomes a contant elasticity of substitution aggregator of total robots and routine labor, where $\varepsilon \leq 1$. The elasticity of substitution between robots and routine workers is given by $1/(1-\varepsilon)$. The production function becomes

$$Y_t = A\left(X_t^\varepsilon + N_{r,t}^\varepsilon\right)^{\frac{1-\alpha}{\varepsilon}} N_{n,t}^\alpha. \tag{35}$$

We write this production function as $F(X_t, N_{r,t}, N_{n,t})$ and denote its partial derivatives at time t as $F_{X,t}$ and $F_{j,t}$ for $j = n, r$. The wage rate per efficiency unit of labor is equal to the worker's marginal productivity

$$w_{j,t} = F_{j,t}.$$

We want to focus on environments in which robots have a higher degree of complementarity with non-routine workers than with routine workers. For this reason, we assume that $\varepsilon > 0$ so that routine workers and robots are substitutes. This assumption implies that the elasticity of the non-routine wage premium with respect to robot use is

$$\mathcal{E}_t \equiv \frac{d\log\left(F_{n,t}/F_{r,t}\right)}{d\log X_t} = \varepsilon m_t \geq 0.$$

To make the model consistent with the life-cycle profile of labor earnings, we assume that workers of age a supply e_a units of labor in efficiency units per hour worked. Total labor supply for occupation $j = n, r$ at time t must satisfy the market-clearing condition:

$$N_{j,t} = \begin{cases} \sum_{a=0}^{t-1} \int_{\Theta_{j,t-a}} e_a l_{\theta,t}^a dH(\theta) + \sum_{a=t}^{L_w-1} \pi_{j,t-a} e_a l_{j,t}^a, & \text{if } t < L_w \\ \sum_{a=0}^{L_w-1} \int_{\Theta_{j,t-a}} e_a l_{\theta,t}^a dH(\theta), & \text{if } t \geq L_w. \end{cases} \tag{36}$$

Aggregate consumption at time t, C_t, is given by

$$C_t = \begin{cases} \sum_{a=0}^{t-1} \int_{\Theta} c_{\theta,t}^a dH(\theta) + \sum_{j=n,r} \sum_{a=t}^{L-1} \pi_{j,t-a} c_{j,t}^a, & \text{if } t < L \\ \sum_{a=0}^{L-1} \int_{\Theta} c_{\theta,t}^a dH(\theta), & \text{if } t \geq L. \end{cases} \tag{37}$$

Using these definitions, the resource constraint in period t can be written as

$$C_t + G_t \leq F\left(X_t, N_{r,t}, N_{n,t}\right) - \phi_t X_t. \tag{38}$$

We define net output as $NY_t = Y_t - \phi_t X_t$.

4.1 First-best allocation

We assume that the planner assigns Pareto weights $\omega_{j,1-a}$ to workers born before the initial period and $\beta^{t-1}\omega_{\theta,t}$ to agents of type (θ, t). To ensure that the first best has a well-defined balanced-growth allocation, we assume that the *current-value* weights converge in the long run; that is, for all θ, $\omega_{\theta,t} \to \omega_\theta \geq 0$ as $t \to \infty$. The planner's objective function is

$$\mathcal{W} \equiv \sum_{\tilde{a}=1}^{L-1} \sum_{j=n,r} \pi_{j,1-\tilde{a}} \omega_{j,1-\tilde{a}} U_{j,1-\tilde{a}} + \sum_{t=1}^{\infty} \int_{\Theta} \beta^{t-1} \omega_{\theta,t} [U_{\theta,t} - \theta s_{\theta,t}] dH(\theta). \tag{39}$$

The first-best allocation maximizes this welfare function subject to the resource constraints, (38). The solution to this problem implies the following efficiency conditions:

$$\frac{\psi'(l_{\theta,t}^a)}{u'(c_{\theta,t}^a)} = e_a F_{s_{\theta,t-a}}(t),$$

$$\omega_{\theta,t-a} u'(c_{\theta,t}^a) = \omega_{\theta,t-a'} u'(c_{\theta,t}^{a'}),$$

for all θ, θ', a, a' and t, and

$$F_X(t) = \phi_t.$$

4.2 Mirrleesian taxation

As in the dynamic Mirrleesian taxation literature, we characterize the second-best problem for a planner who can design allocations that are functions of observable histories of income and consumption, but not of each worker's type, wage, or skill choice. We assume that the planner can tax the different generations differently. This assumption is common in the optimal taxation literature.

We consider a direct revelation mechanism in which the planner elicits information on the worker's type and assigns the worker a profile of consumption, labor supply, and skill choice. In line with Scheuer (2014), we write the implementability constraints as follows.

The first incentive constraint is the same as in the static model. For the workers who are born before the initial date and have $\widetilde{a} < L_w$ at $t = 1$, this is the only relevant constraint:[16]

$$U_{n,1-\widetilde{a}} \geq U_{r,1-\widetilde{a}} + \sum_{a=\widetilde{a}}^{L_w-1} \beta^{a-\widetilde{a}} \left[\psi\left(l^a_{r,1+a-\widetilde{a}}\right) - \psi\left(\frac{F_{r,1+a-\widetilde{a}}}{F_{n,1+a-\widetilde{a}}} l^a_{r,1+a-\widetilde{a}}\right)\right], \tag{40}$$

$$U_{r,1-\widetilde{a}} \geq U_{n,1-\widetilde{a}} + \sum_{a=\widetilde{a}}^{L_w-1} \beta^{a-\widetilde{a}} \left[\psi\left(l^a_{n,1+a-\widetilde{a}}\right) - \psi\left(\frac{F_{n,1+a-\widetilde{a}}}{F_{r,1+a-\widetilde{a}}} l^a_{n,1+a-\widetilde{a}}\right)\right]. \tag{41}$$

for $a = 1, ..., L_w - 1$. For the workers born after the initial date, this constraint is

$$U_{\theta,t} \geq U_{\theta',t} + \sum_{a=0}^{L_w-1} \beta^{a} \left[\psi\left(l^a_{\theta',t+a}\right) - \psi\left(\frac{F_{s_{\theta',t},t+a}}{F_{s_{\theta,t},t+a}} l^a_{\theta',t+a}\right)\right], \tag{42}$$

for all $\theta, \theta' \in \Theta$ and $t \geq 1$. This *intensive-margin incentive constraint* guarantees that the worker chooses the assigned allocation, given the worker's occupation choice.

The second condition is the incentive constraint for the choice of occupation of an individual of type θ:

$$U_{\theta,t} - \theta s_{\theta,t} \geq U_{\theta',t} - \theta s_{\theta',t}, \tag{43}$$

[16] There is no incentive problem for agents who are retired.

for all $\theta, \theta' \in \Theta$ and $t = 1, 2, 3...$. This *extensive-margin incentive constraint* ensures that the worker chooses the assigned occupation.[17] The planning problem is to maximize (39) subject to these incentive constraints and the resource constraints, (38).

As in Scheuer (2014), we now state two results that allow us to simplify the analysis.[18]

Lemma 1. *An allocation satisfies the extensive-margin incentive constraints if and only if for all t there exists $U_{n,t}, U_{r,t} \in \mathbb{R}$ and $\theta_t^* = U_{n,t} - U_{r,t}$ such that:*

1. *If $\theta < \theta_t^*$, then $s_{\theta,t} = 1$ and $U_{\theta,t} = U_{n,t}$;*

2. *If $\theta > \theta_t^*$, then $s_{\theta,t} = 0$ and $U_{\theta,t} = U_{r,t}$.*

This lemma allows us to simplify the incentive constraints. For an allocation to be incentive compatible, all workers that choose the same skill should have the same utility gross of skill-acquisition costs. This property allows us to express the incentive constraints as a cutoff rule: workers with $\theta < \theta_t^*$ acquire non-routine skills, whereas those with high $\theta > \theta_t^*$ acquire routine skills.

The next lemma allows us to further simplify the problem. This lemma shows that all workers who have the same skills should have the same allocation in terms of consumption and labor.

Lemma 2. *At the optimum, if $s_{\theta,t} = s_{\theta',t}$, then these two workers have the same consumption at each age $c^a_{\theta,t+a} = c^a_{\theta',t+a}$, for $a = 0, ..., L-1$, and the same labor supply $l^a_{\theta,t+a} = l^a_{\theta',t+a}$, for $a = 0, ..., L_w - 1$.*

To find the allocations for routine and non-routine workers, it is useful to define

$$U_{j,t} \equiv \sum_{a=0}^{L-1} \beta^a \left[u\left(c^a_{j,t+a}\right) + v(G_{t+a}) \right] - \sum_{a=0}^{L_w-1} \beta^a \psi\left(l^a_{j,t+a}\right),$$

[17] These constraints do not explicitly take into account the possibility that agent θ might choose an allocation that corresponds to an occupational choice that is different from $s_{\theta',t}$. However, those additional constraints are redundant.

[18] See the appendix for the proofs of these results.

for $j = n, r$. The number of incentive constraints can be simplified to just two per generation born before time $t = 1$, (41) and (40), and three constraints per generation born after $t = 1$,

$$\theta_t^* = U_{n,t} - U_{r,t}, \tag{44}$$

and

$$U_{n,t} \geq U_{r,t} + \sum_{a=0}^{L_w-1} \beta^a \left[\psi\left(l_{r,t+a}^a\right) - \psi\left(\frac{F_{r,t+a}}{F_{n,t+a}} l_{r,t+a}^a\right)\right], \tag{45}$$

$$U_{r,t} \geq U_{n,t} + \sum_{a=0}^{L_w-1} \beta^{t-1} \left[\psi\left(l_{r,t+a}^a\right) - \psi\left(\frac{F_{n,t+a}}{F_{r,t+a}} l_{n,t+a}^a\right)\right]. \tag{46}$$

The next proposition states results that are analogous to those we obtained for the static model. As long as automation is incomplete and at least one intensive-margin incentive constraint binds, (40) or (45), it is optimal to tax robots in a given period.

Proposition 2. *At the optimal plan, suppose that at time t, there is an age a such that: (i) the intensive-margin constraint for a non-routine worker of age a at time t is binding and (ii)* $l_{r,t}^a > 0$*; and no intensive-margin constraint of routine workers working at time t is binding. Then, robot usage should be distorted,* $F_X(t) > \phi_t$*; that is, robots should be taxed.*

Since this model features endogenous skill acquisition, the intensive-margin incentive constraint of non-routine workers might not bind. This is because the government can redistribute income in two ways. The first, which we call the *direct redistribution mechanism*, is to redistribute income from non-routine to routine workers. This mechanism is the one used in our static model. But in a model with endogenous skill choice, this mechanism reduces the incentive for workers to acquire non-routine skills.

The second, which we call the *indirect redistribution mechanism*, involves little income redistribution in order to provide an incentive for workers to acquire non-routine skills. When this mechanism is the most relevant, the intensive-margin incentive constraint no longer binds. Because robot taxes are desirable only insofar as they help to provide incentives along the intensive margin, then, if the government redistributes indirectly, robot taxes should be zero. Which mechanism turns out to be optimal is a quantitative question.

Asymptotic balanced growth We assume that the cost of robots declines geometrically over time as a result of exogenous technical progress, $\phi_t = \widetilde{\phi} e^{-g_\phi (t-1)}$. In addition, we assume that $u(\cdot)$ and $v(\cdot)$ are logarithmic functions so that preferences are consistent with balanced growth.

Assumption 2 (Preferences)**.** *The utility function takes the form* $u(c) = \log(c)$ *and* $v(G) = \chi \log(G)$*, with* $\chi > 0$*.*

These preferences have been used in different public finance applications, especially the ones featuring technical change; see, for example, ALES *et al.* (2015) and HEATHCOTE *et al.* (2017). Recall that these preferences are also compatible with the empirical evidence reviewed in CHETTY (2006).

The variables in the model can be normalized to remove trends (see appendix A.2.4). We call the version of our model expressed in terms of these normalized variables the *normalized economy*. We say that the economy is on a *balanced-growth path* if the allocations of the normalized economy are constant over time; that is, the normalized economy is in a steady state.

In dynamic optimal taxation problems, the steady-state allocations generally depend on initial conditions; see, for example, CHAMLEY (1986) or SLAVÍK and YAZICI (2014). This dependence requires solving for the balanced-growth path and transition jointly, which is often challenging from a computational standpoint. In our model, the steady state of the normalized economy is independent of initial conditions because, in our overlapping-generations structure, workers have finite horizons and the government can treat different generations differently.

We show in appendix A.2.5 that if aggregate consumption, government spending, aggregate labor supply, robot use, and the cutoff θ_t^* converge to an interior balanced-growth path, then all other variables including individual allocations and Lagrange multipliers also converge to constant values. In the appendix, we show the necessary and sufficient conditions to compute this balanced-growth path.

Proposition 3. *Suppose that the optimal plan is such that the allocations converge to a balanced-growth path with interior automation. Then, the optimal tax on robots converges asymptotically to zero.*

This proposition holds irrespective of the distribution of skill acquisition costs. As a result, it it remains valid even if costs are arbitrarily high, that is, in an economy with exogenous skills.

Proposition 3 stands in contrast with the optimal tax scheme in Slavík and Yazici (2014). These authors consider optimal Mirrleesian taxation in an infinite-horizon model with low- and high-skill workers and capital-skill complementarity. They find that in this setting, optimal asymptotic production distortions are high.

Slavík and Yazici (2014) abstract from technical progress. In the presence of technical progress, the asymptotic balanced-growth path would be such that the workers for whom wages fall no longer supply any labor. As a result, there is no reason to affect pretax wages to provide intensive-margin incentives, and thus production efficiency is optimal. In the same way that in our model robot taxes converge to zero, capital taxes would also converge to zero in a version of the Slavík and Yazici (2014) model with technical progress.

5 Quantitative analysis

In this section, we describe our calibration and solve the planning problem in order to quantify the effects of advances in automation on optimal tax policy.

5.1 Parameter calibration

We calibrate the parameters so that the status quo economy matches salient features of the U.S. economy for the period 1987-2017. Table 1 summarizes the calibrated parameters.

Our calibration targets the non-routine wage premium and the occupation shares of each skill type. We obtain time-series data for these variables from the Current Population Survey (CPS) March Annual Social and Economic Supplement (ASEC), compiled in FLOOD *et al.* (2018). Our classification of occupations into routine and non-routine is the one proposed by CORTES *et al.* (2014). We compute the share of employment of routine and non-routine workers using the ASEC-CPS weights and the average weekly wage of routine and non-routine workers using personal earnings weights.

The utility function is assumed to be isoelastic in consumption, labor, and government spending:

$$u(c) = \log c\,, \qquad \psi(l) = \zeta \frac{l^{1+\nu}}{1+\nu}\,, \qquad v(G) = \chi \log G. \tag{47}$$

This utility specification is consistent with balanced growth. The cross-sectional distribution of θ, $h(\theta)$, follows a logistic distribution with location parameter μ and scale parameter σ.[19]

In the status quo, conditional on a given skill choice, all workers solve the same problem. Workers of the same occupation choose the same consumption, labor supply, and savings and obtain the same utility,

$$U_{\theta,t} = \begin{cases} U_{n,t}, & \text{if } s_{\theta,t} = 1 \\ U_{r,t}, & \text{if } s_{\theta,t} = 0, \end{cases} \tag{48}$$

[19] This assumption is equivalent to postulating that the worker has occupation-specific utility costs of acquiring skills, θ_n and θ_r, and that these costs follow a Gumbel distribution. This distributional assumption has been widely used in the literature on discrete choice following MCFADDEN (1974); see, for example, JOHNSON and KEANE (2013) and ROYS and TABER (2019).

for the equilibrium levels of $U_{j,t}$. As a result, the skill choice can be described by a threshold rule $\theta_t^* = U_{n,t} - U_{r,t}$, such that all newborns at time t with $\theta < \theta_t^*$ choose non-routine skills, and those with $\theta > \theta_t^*$ choose routine skills.

The budget constraint for workers at time t, age a, in a particular occupation is given by

$$c_{j,t}^a + \frac{b_{j,t}^a}{R_t} = b_{j,t-1}^{a-1} + w_{j,t} e_a l_{j,t}^a - T_t(w_{j,t} e_a l_{j,t}^a), \qquad \text{for } a = 0, ..., L_w - 1, \tag{49}$$

$$c_{j,t}^a + \frac{b_{j,t}^a}{R_t} = b_{j,t-1}^{a-1}, \qquad \text{for } a = L_w, ..., L - 1, \tag{50}$$

where the final wealth holdings are zero, $b_{j,t}^{L-1} = 0$. The initial wealth holdings are zero for the newborn populations, $b_{j,t-1}^{-1} = 0$, and the initial bond holdings, $b_{j,1}^a$ for $a \geq 0$, are given. Here, R_t denotes the gross real interest rate between t and $t+1$. The taxation of labor earnings is the same as in Heathcote *et al.* (2017), which means that $T(y) = y - \lambda_t y^{1-\gamma}$.

The government's time-t flow constraint is given by

$$G_t + B_{t-1} = \sum_{a=0}^{L_w - 1} \sum_{j=n,r} \pi_{j,t-a} T_t \left(w_{j,t} e_a l_{j,t}^a \right) + \tau_t^x \phi_t X_t + \frac{B_t}{R_t}. \tag{51}$$

In our calibration exercises, we set the ratios of government spending to consumption, G_t/C_t, and assets to consumption, $\frac{B_t}{R_t}/C_t$. We set λ_t so that the government budget constraint is satisfied.

We now describe how the parameters are calibrated.

Externally calibrated parameters We assume that a time period corresponds to ten years. Workers live for six periods and work for four of these; that is, $L = 6$ and $L_w = 4$. Following Chetty *et al.* (2011), we set the Frisch elasticity to 0.75, $\nu = 1/0.75$. We calibrate e_a to match the life-cycle earnings profile in Guvenen *et al.* (2015). On the production side, we normalize $A = 1$.

Table 1: Calibration

Parameter		*Value*	*Source/ Target*
Externally calibrated			
L	Life	6	
L_w	Working life	4	
ν	Inverse Frisch elasticity	1.33	Chetty *et al.* (2011)
e	Life-cycle earnings	$\{1, 1.81, 2.19, 2.18\}$	Guvenen *et al.* (2015)
A	TFP	1	Normalization
γ	Tax progressivity	0.18	NBER TAXSIM
B/C	Debt-to-cons. ratio	0.07	World Bank
G/C	Spending-to-cons. ratio	0.19	World Bank
χ	Spending utility parameter	0.19	Average G/C ratio
τ_x	Tax on robots	0.038	See text
Internally calibrated			
β	Discount factor	0.47	$\beta R = 1$ Gourinchas and Parker (2002)
ζ	Labor disutility parameter	14.03	Average $l = 0.33$
α	Income share of non-routine	0.48	Wage premium in 1987
$\widetilde{\phi}$	Robot tech.: level	0.43	Wage premium in 1988 – 2017
g_ϕ	Robot tech.: growth rate	0.02	Wage premium in 1988 – 2017
ε	Routine-robots elasticity	1	Wage premium in 1988 – 2017
μ	Skill-acq. preference: mean	0.34	Occupation share in 1988 – 2017
σ	Skill-acq. preference: variance	0.35	Occupation share in 1988 – 2017

Note: This table summarizes the calibration for the baseline model. Externally calibrated parameters are shown with the source. Internally calibrated parameters are shown next to their most informative targets. See text for details on the calibration.

When utility functions take the form (47), the optimal ratio of government spending to private consumption is equal to χ. We set the value of χ to 0.19 to match the average ratio of government spending to consumption in the data.

Acemoglu *et al.* (2020) estimate that, during this period, the effective tax rate on equipment and software capital is around 10 percent. In our model, this estimate translates into a tax on robot purchases of 3.8 percent.[20] We use the method in Ferriere and Navarro (2014) to calibrate γ using NBER TAXSIM data and find $\gamma = 0.18$. This value is in line with the estimates in Heathcote *et al.* (2017). We use World Bank data to compute the ratio of government spending to private consumption and the ratio of government bonds to private consumption. During this time period, the average values for these ratios are 0.19 percent and 0.07, respectively.[21] We compute the values of G_t and B_t so as to be consistent with these ratios and adjust the level parameter in the tax function, λ_t, to satisfy the government flow budget constraint, (51).

We then proceed with our calibration in two steps: we first find a steady state for the normalized economy with fixed occupations and zero automation for 1987. We use this steady state to calibrate the subjective discount factor, β, the labor disutility parameter, ζ, and the share of non-routine workers in production, α. Next, we compute a perfect foresight transition to the new balanced-growth path and use the data on the non-routine wage premium and the occupation shares between 1988 and 2017 to calibrate the level parameter of the cost of robots, $\widetilde{\phi}$, the rate of technical progress in robot production, g_ϕ, the elasticity of routine workers and robots, $1/(1-\varepsilon)$, and the parameters of the distribution of skill acquisition costs, μ and σ.

[20] We compute the equivalent ad-valorem robot tax as $\tau^x = [r/(1-\tau) + \overline{\delta}]/(r + \overline{\delta}) - 1$, where r denotes the annual real interest rate, $\overline{\delta} = 1 - (1-\delta)e^{-g_\phi}$ denotes the annual depreciation rate adjusted for the rate of technical progress, and δ is the annual depreciation rate. The variable τ denotes the effective tax rate on equipment and software capital income estimated by Acemoglu *et al.* (2020). We use the annual real interest rate implied by the model $r = \beta^{-1/10} - 1$, $\overline{\delta} = 0.15$, and $\tau = 0.1$.

[21] The debt-to-consumption ratio is adjusted for the fact that one period in the model corresponds to 10 years. The annual debt-to-consumption ratio is 70 percent.

Pre-automation steady-state equilibrium Our first step is to calibrate the steady state before automation so as to match the ratio of government spending to private consumption and the ratio of government debt to private consumption for 1987. We also calibrate the shares of routine and non-routine workers to match their 1987 shares: $\pi_n = 0.44$ and $\pi_r = 0.56$. In the pre-automation steady state, these occupational shares are constant across generations.

Following GOURINCHAS and PARKER (2002), we impose $\beta R = 1$. We choose the labor disutility parameter, ζ, so that on average the labor supply is equal to $1/3$. Finally, we calibrate α so that the non-routine wage premium is $w_n/w_r = 1.19$. This calibration leads to $\beta = 0.47$, $\zeta = 14.03$, and $\alpha = 0.48$.

Figure 2: Calibration fit

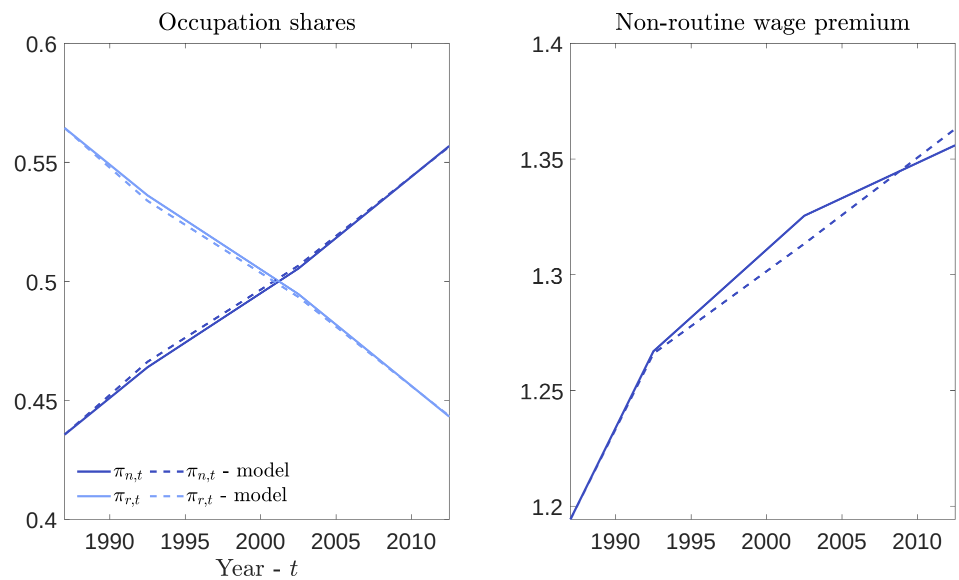

Notes: The left panel shows the labor force shares of routine and non-routine workers, and the right panel shows the non-routine wage premium, $w_{n,t}/w_{r,t}$. In each figure, solid lines represent these quantities computed from ASEC-CPS data, and dashed lines show their counterparts in the status quo equilibrium for the calibrated model. See text for details.

Transitional dynamics and steady state We solve for the perfect foresight transition between the initial and the final steady state of the normalized economy.

In the asymptotic final steady state, there is full automation. Labor hours and consumption are zero for routine workers, so all workers choose non-routine occupations.[22]

Recall that the cost of robots declines geometrically over time as a result of exogenous technical progress, $\phi_t = \widetilde{\phi} e^{-g_\phi (t-1)}$. We choose the technological parameters $\widetilde{\phi}, g_\phi$, and ε, plus the skill acquisition parameters μ and σ to match the time series of observed occupation shares and the non-routine wage premium. Using a least squares procedure, we find $\widetilde{\phi} = 0.43$, $g_\phi = 0.02$, and $\varepsilon = 1$, plus $\mu = 0.34$ and $\sigma = 0.35$. Figure 2 shows that our model fits the ten-year average trends quite well in both the occupation shares and the non-routine wage premium.

5.2 Status quo equilibrium

Acemoglu *et al.* (2020) estimate that the effective tax rate on income from equipment and software capital declined from 10 to 5 percent after the 2017 tax reform. Assuming these tax rates apply to robots, they correspond to ad-valorem tax rates of 3.8 and 1.8 percent, respectively. We recompute the perfect-foresight equilibrium starting in 2017 assuming that the tax reform was unanticipated and permanent.

Figures 3 and 4 plot prices and allocations for the status quo economy. The solid and dashed lines correspond to the paths with and without the 2017 tax reform, respectively.

The top left panel in figure 3 shows the path for the cost of robots. The top right panel shows that as these costs fall, automation rises. The bottom left panel shows that net output rises. The gains from technical progress are unevenly distributed. The bottom right panel shows that wages fall for those who are in routine occupa-

[22] We describe the full set of equilibrium conditions and the numerical method in the appendix.

Figure 3: Status quo equilibrium A

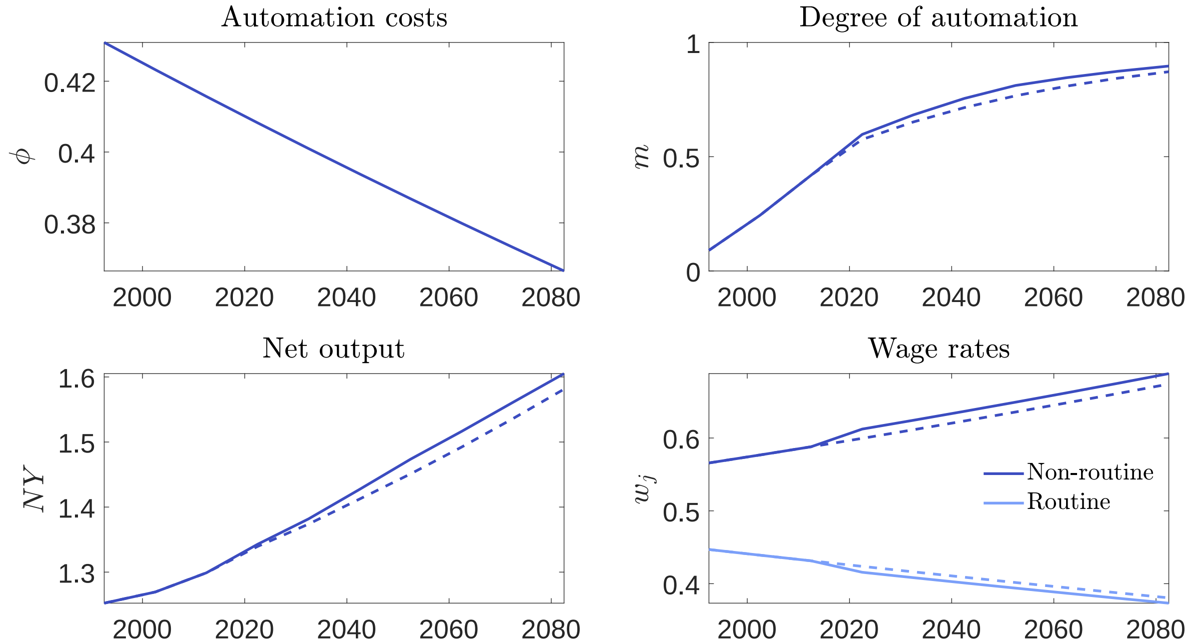

Notes: This figure plots the status quo equilibrium for the calibrated economy starting from 1987. The first and second panels show the evolution of robot costs, ϕ_t, and automation, m_t, respectively. The third and fourth panel plot the equilibrium levels of net output and wages for the two worker types, respectively. In each panel, solid lines represent the equilibrium with the 2017 tax change which lowered capital income taxes, while dashed lines represent the counterfactual equilibrium without the 2017 tax change. See text for details.

tions and rise for those in non-routine occupations.

Turning to the effects of the 2017 tax reform, we see that the reduction in capital income taxes increases automation, and, as a result, net output is higher than without the tax reform. The reduction in robot taxes leads to a fall in routine wages and a rise in non-routine wages.

Figure 4: Status quo equilibrium B

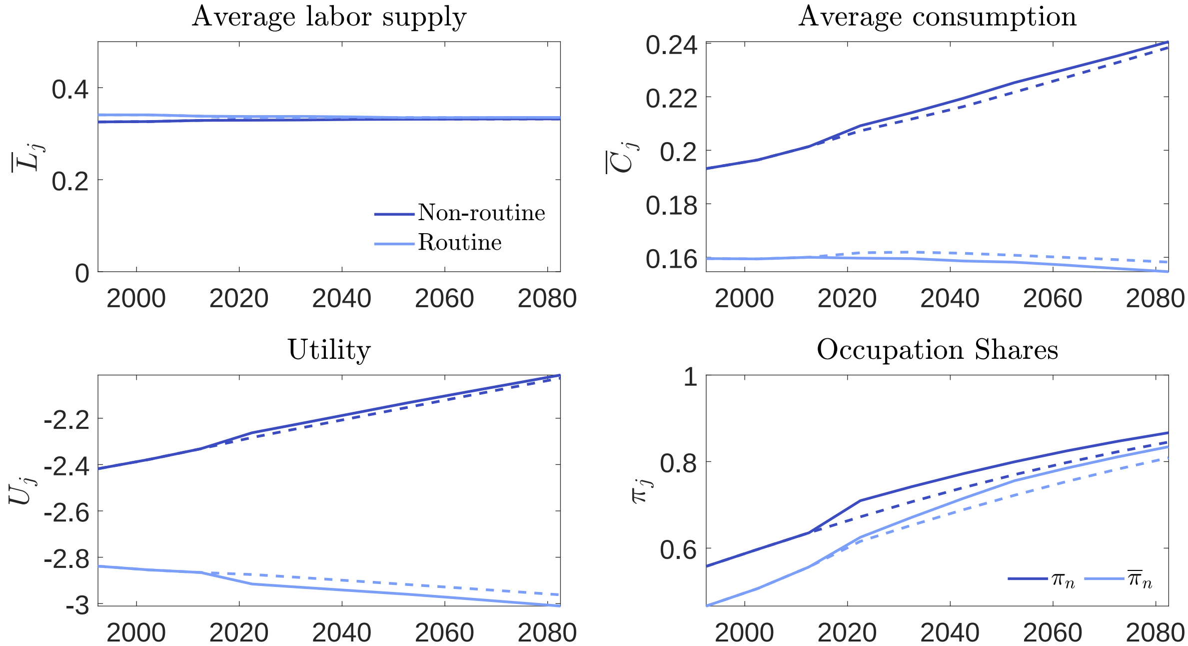

Notes: This figure plots the status quo equilibrium for the calibrated economy starting from 1987. The first and second panels show the equilibrium levels of average labor supply and consumption in the cross section, respectively. The third panel plots the equilibrium levels of utility, gross of skill-acquisition costs, for the generation born in the period. Finally, the fourth panel shows the share of newborns who choose non-routine skills, $\pi_{n,t}$, and the share of non-routine workers in the workforce, $\overline{\pi}_{n,t}$. In each panel, solid lines represent the equilibrium with the 2017 tax change which lowered capital income taxes, while dashed lines represent the counterfactual equilibrium without the 2017 tax change. See text for details.

The top two panels in Figure 4 show the cross-sectional average of labor supply and consumption for routine and non-routine workers, computed as follows: $\overline{C}_{j,t} \equiv \sum_{a=0}^{L-1} \pi_{j,t-a} c_{j,t}^{a} / \sum_{a=0}^{L-1} \pi_{j,t-a}$ and $\overline{L}_{j,t} \equiv \sum_{a=0}^{L_w-1} \pi_{j,t-a} l_{j,t}^{a} / \sum_{a=0}^{L_w-1} \pi_{j,t-a}$. Because our model has preferences that are consistent with balanced growth, labor supply is fairly constant and close to 0.33. The top right panel shows that consumption rises for non-routine workers and falls for routine workers. The third panel plots the level of utility, gross of skill-acquisition costs, for newborn workers in non-routine and routine occupations. Technical progress increases the utility of non-routine workers and reduces that of routine workers. The 2017 tax reform widens consumption and utility inequality in this model by reducing the implied tax on robots.

The lower right panel of Figure 4 plots the time series for the share of newborn workers who are non-routine, $\pi_{n,t}$, and the share of non-routine workers in the labor force, $\overline{\pi}_{j,t} \equiv \sum_{a=0}^{L_w-1} \pi_{j,t-a}/L_w$. As routine wages decline, more agents decide to become non-routine workers . However, the share of non-routine workers in the labor force, $\overline{\pi}$, responds sluggishly. This inertia reflects the inability of older generations to reoptimize their skill choices.

5.3 Mirrleesian optimal taxation

In this section, we discuss the properties of the Mirrleesian optimal plan for the calibrated economy. We solve numerically for the optimal tax policy starting in 2018, assuming that the planner assigns equal current-value weights to all workers. The solid and dashed lines in Figures 5 and 6 display the optimal Mirrleesian allocation and the status quo equilibrium with the 2017 tax reform, respectively.

Optimal robot taxes in the Mirrleesian plan are slightly higher than in the status quo equilibrium for the first two decades. They are 5.15, 2.18, and 0.62 percent for the decades that start in 2018, 2028, and 2038, respectively. The optimal tax rate is zero after three decades.

The fourth panel of Figure 5 shows that these higher robot taxes lower the non-routine wage premium relative to the status quo. This wage compression loosens the incentive constraint of non-routine workers, which makes it easier for the government to redistribute income from non-routine to routine workers in the initial older generations (i.e., those with $\widetilde{a} \geq 1$). As discussed before, the direct redistribution mechanism is always optimal for the initial old generations because they cannot change their skills.

After these initial periods, the tax on robots falls permanently to zero. The reason

Figure 5: Mirrleesian optimal taxation A

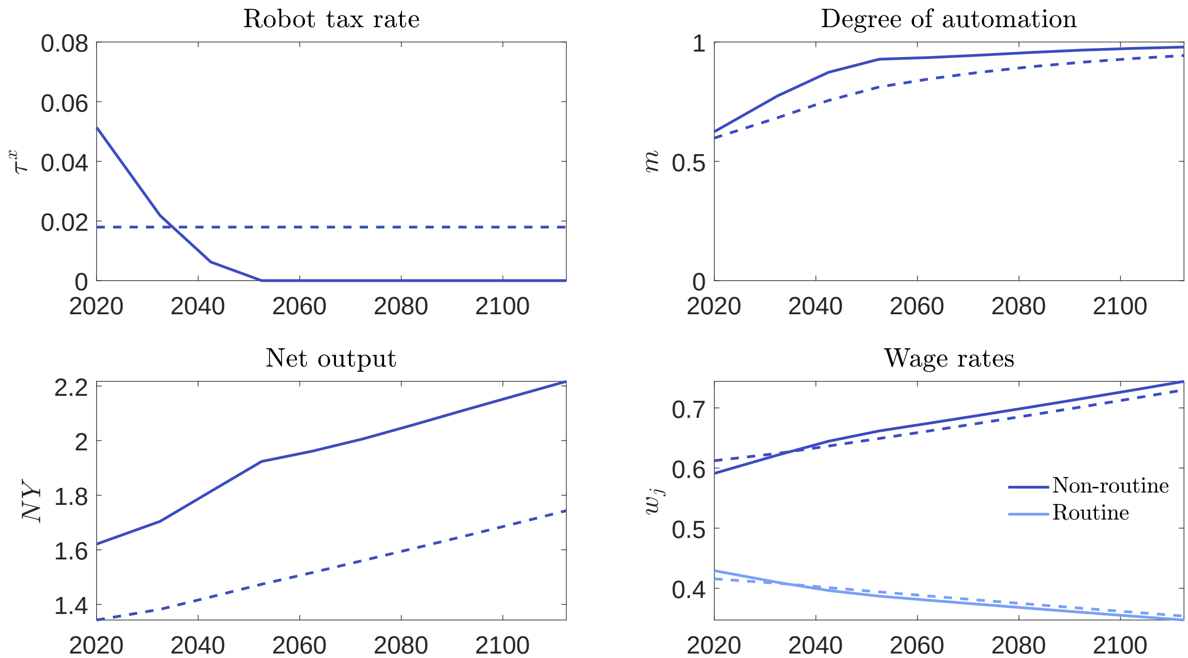

Notes: This figure plots the allocations in the Mirrleesian optimal plan for the calibrated economy starting from 2018. These quantities are shown in solid lines. For comparison, we also plot the status quo equilibrium in dashed lines. The first and second panels show the evolution of robot taxes, τ_t^x, and automation, m_t, respectively. The third and fourth panels plot the equilibrium levels of net output and wages for the two worker types, respectively.

for this result is twofold. First, after the first three decades, the initial old generations are no longer active in the labor force. Second, after three decades, automation costs become relatively low. As a result, the Mirrleesian planner finds it optimal to use the indirect redistribution mechanism exclusively. Given that robot taxes play no role in this mechanism, they are set to zero.

Compared to the status quo equilibrium, the Mirrleesian optimal plan induces more workers to choose non-routine skills. Table 2 compares the composition of the labor force in the status quo equilibrium and the Mirrleesian optimal plan for the initial periods. It shows the share of workers who become non-routine workers in the newborn population, $\pi_{n,t}$, and the share of workers who are non-routine in the

Figure 6: Mirrleesian optimal taxation B

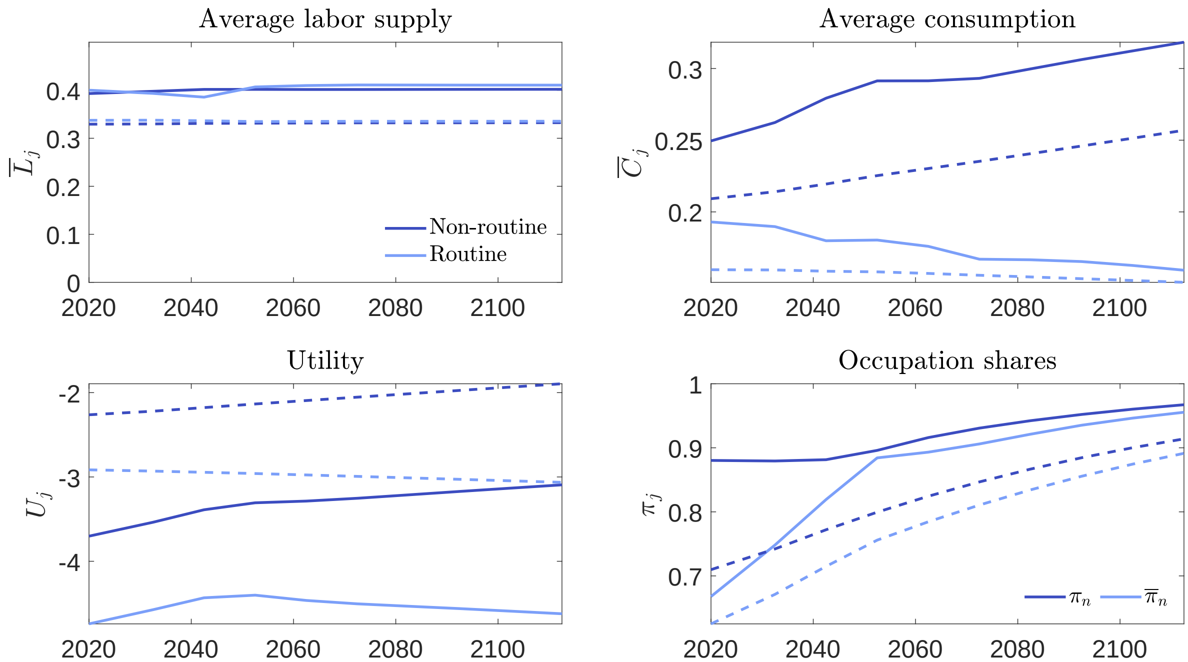

Notes: This figure shows the allocations in the Mirrleesian optimal plan for the calibrated economy starting from 2018. These quantities are shown in solid lines. For comparison, we also plot the status quo equilibrium in dashed lines. The first and second panels show the equilibrium levels of average labor supply and consumption in the cross section, respectively. The third panel plots the equilibrium levels of utility for both agents. Finally, the fourth panel shows the share of newborns who choose non-routine skills, $\pi_{n,t}$, and the share of non-routine workers in the workforce, $\overline{\pi}_{n,t}$.

labor force, $\overline{\pi}_{n,t}$.

We can see that the share of non-routine workers is more than 10 percentage points higher in the Mirrleesian optimal plan relative to the status quo equilibrium. This higher share reflects the proclivity of new generations for non-routine occupations. The model is consistent with the finding in Adão *et al.* (2018) of a weak response of old generations and a strong response of new generations to changes in wages across occupations. Redistribution through occupation choice, which we call the indirect redistribution mechanism, plays an important role in these results. The

Table 2: Labor force composition

Non-routine share		2018–2027	2028–2037	2038–2047	2048–2057
Status quo	Newborn	0.71	0.74	0.77	0.80
	Labor force	0.63	0.67	0.71	0.76
Mirrleesian plan	Newborn	0.88	0.88	0.88	0.90
	Labor force	0.67	0.75	0.82	0.88

Note: This table compares the share of non-routine workers in the newborn population, $\pi_{n,t}$, and in the labor force, $\overline{\pi}_{n,t}$ for the first four periods of the status quo equilibrium and Mirrleesian optimal taxation in the calibrated economy.

planner designs allocations with little direct redistribution between worker types so that a higher share of workers in the new generations acquires non-routine skills. Consequently, the intensive-margin incentive constraint no longer binds. Because robot taxes should be used only insofar as they help to loosen the intensive-margin incentive constraint, robot taxes fall to zero.

Table 2 also shows that, despite large changes in the composition of the newborn population, the composition of the labor force changes slowly. This inertia reflects the inability of older workers to change occupations. Despite large differences in the composition of the newborn population, the share of non-routine workers in the labor force in the period 2018–2027 is only 4 percentage points higher in the Mirrleesian optimal plan relative to the status quo economy. This gap finally reaches 12 percentage points in the fourth period, 2048–2057, at which point all workers in the labor force were born after the initial date of 2018.

Because workers born before 2018 cannot readjust their skill choices, the government can only use the direct redistribution mechanism to improve their welfare. As a result, there is a reason to tax robots initially in order to loosen the intensive-margin constraint of those non-routine workers. This is why robot taxes are positive initially. As time goes by, the share of workers who did not readjust their skill choices decreases, which implies that there is less of a reason to distort robot use. As a result, the tax on robots declines over these initial periods. After these workers leave the labor force, there's no longer any reason to tax robots.

Robustness We calibrate the model to be consistent with the past evolution of the non-routine wage premium. This calibration implies a relatively slow rate of decline in robot costs, as compared to the price of other forms of capital. Our baseline policy assumes that future technical progress occurs at the same pace.

We now consider a scenario in which future robot prices decline on average at the same rate as the producer price index for computer and peripheral equipment between 1995 and 2015, computed by the Bureau of Labor Statistics. We call this alternative calibration the "Moore's law scenario." Table 3 reports our results. We focus on the implications for the optimal robot tax, the share of non-routine workers in the newborn population, and the non-routine wage premium. The faster pace of technical progress leads to a faster decline in the wages received by the initial generations of routine workers, creating more inequality. For this reason, the planner implements a higher initial robot tax. In the optimal plan, almost all newborn workers become non routine.

In the second decade, the optimal robot tax is close to zero. Two main forces drive this result. First, since robot costs are lower and there is more non-routine labor supply, the costs of distorting production are higher. Equation (31) shows that the marginal loss in production associated with an increase in τ^x is proportional to total labor supply by non-routine workers.[23] It is also inversely proportional to $\phi^{\frac{1-\alpha}{\alpha}}$. Both the increase in total non-routine labor supply and the fall in automation costs imply that optimal robot taxes should be lower.

Second, as robots become cheaper, it is efficient to have routine workers supply

[23]Equation (31) pertains to the static model, but an analogous version applies to the calibrated dynamic model. In the latter version, $\pi_j l_j$ is replaced by total labor supply in each occupation $N_{j,t}$.

Table 3: Robustness

			2018–2027	2028–2037	2038–2047
Baseline $g_\phi = 0.02$	Robot tax (%)	τ_t^x	5.1	2.2	0.6
	Newborn non-routine	$\pi_{n,t}$	0.88	0.88	0.88
	Wage premium	$w_{n,t}/w_{r,t}$	1.38	1.52	1.63
Moore's law scenario $g_\phi = 0.29$	Robot tax (%)	τ_t^x	5.7	0.2	0.0
	Newborn non-routine	$\pi_{n,t}$	0.99	1.00	1.00
	Wage premium	$w_{n,t}/w_{r,t}$	1.36	2.79	5.12

Note: This table reports features of the Mirrleesian optimal plan for different configurations of parameters. The baseline parameters are reported in Table 1. For comparison purposes, we start by reporting the baseline results.

fewer hours. Lower labor supply by the routine worker reduces the informational gains from distorting relative wages. It follows that robot taxes are less useful as a tool for improving income redistribution.

We perform other robustness exercises, which are reported in Appendix A.4. In particular, we consider alternative parameter values for the elasticity of skill acquisition, σ, and the levels of robot costs, $\widetilde{\phi}$. We also consider smaller perturbations in the pace of technical progress, g_ϕ.

We vary the level of robot costs by increasing/decreasing $\widetilde{\phi}$ by a factor corresponding to a decade's worth of technical progress. In general, lower robot costs lead to slightly lower tax rates and a higher share of workers in the newborn generations who choose to acquire non-routine skills. Finally, we consider a case with lower elasticity of skill acquisition by increasing the variance of the θ distribution. Because the costs of skill acquisition become higher, the optimal Mirrleesian policy features a lower share of newborn non-routine workers and slightly higher robot taxes. Optimal robot taxes are only moderately affected by small perturbations in the pace of technical progress. As above, higher g_ϕ implies higher robot taxes in the initial period, but lower robot taxes from the second period on.

Finally, we compute the optimal Mirrleesian plan starting in 1988, at the onset of the automation era. Optimal robot taxes are 9.1, 2.9, and 0.7 percent in the first, second, and third decades, respectively. The optimal robot tax is zero from 2018 on. The two main drivers of the higher initial robot taxes are the higher robot costs and the higher share of routine workers in the initial old generations in 1988, which leads to more direct redistribution.

How large are the welfare gains from the robot tax? To answer this question, we compare the utility levels obtained in the Mirrleesian optimal plan above with those that would be obtained with an optimal Mirrleesian income tax system in which the robot tax is constrained to be zero. By comparing these two allocations, we isolate the effect of the robot tax from the welfare gains of moving to a more efficient income tax system.

Figure 7: Consumption-equivalent welfare gains

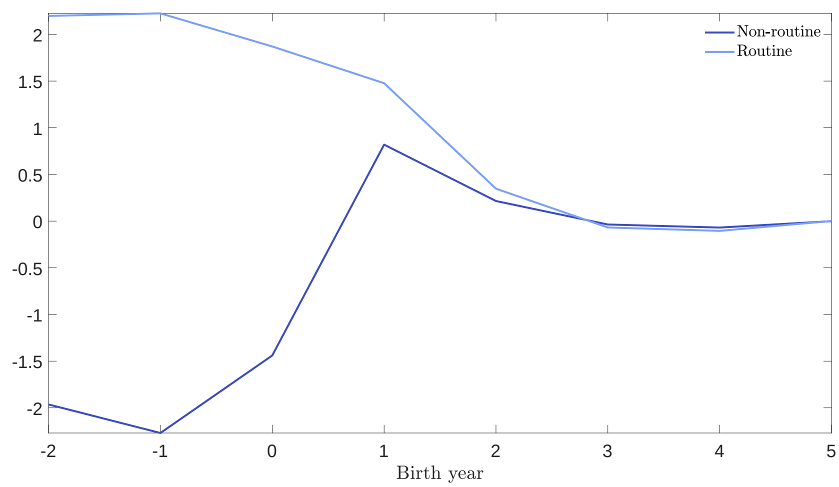

Notes: This figure shows the consumption-equivalent welfare gains of the optimal robot tax. On the x-axis we vary the agent's birth year, with 0 corresponding to the agent being born one period before the initial date. These welfare gains are computed gross of skill-acquisition costs.

We compute the welfare gains in terms of a consumption equivalent — that is, the proportional increase in consumption that compensates the agents from not moving to the optimal robot tax plan. In terms of utilitarian social welfare, the optimal robot tax leads to a consumption-equivalent gain of 0.21 percent.

Figure 7 shows that the main beneficiaries of the robot tax are the routine workers in the initial old generations, and the main losers are the initial non-routine workers. The welfare gains of the initial old generations of routine workers are as high as 2.2 percent of consumption. The losses for the initial non-routine workers are of the same order of magnitude. The non-routine workers born in the initial period, who are able to choose their skills, benefit from the robot tax, albeit less than the routine workers born in that same period. Once robot taxes are no longer used, from period 4 onward, the welfare differences between the two plans are essentially zero.

The role of endogenous skill acquisition To isolate and clarify the effects of the endogenous nature of skill acquisition on the design of the tax system, we consider a model in which the number of routine and non-routine workers follows an exogenous path. This path coincides with the equilibrium evolution of the number of routine and non-routine workers in our benchmark dynamic model.

Figure 8 compares the paths for the economies with endogenous and exogenous skill acquisition, which we call *ES* and *XS*, respectively. This figure shows that routine workers receive a better allocation of consumption and hours worked in the exogenous skills economy. The reason for this property is that, in the endogenous skills economy, improving the allocation of routine workers reduces the incentive of the new generations to acquire non-routine skills.[24]

When skills are exogenous, the planner taxes robots at higher rates than when skills are endogenous. With endogenous skills, robot taxes become zero after 2048. In contrast, with exogenous skills, the tax rate on robots converges to zero only asymp-

[24] Indeed, if we were to fix the allocations of the *XS* economy but allow skill supply to be endogenous, the non-routine share in the young generations would drop by almost 40 percentage points.

tically. These properties follow from the fact that only the direct redistribution mechanism is relevant in the economy with exogenous skills.

Figure 8: Mirrleesian optimal taxation: endogenous (ES) and exogenous skills (XS)

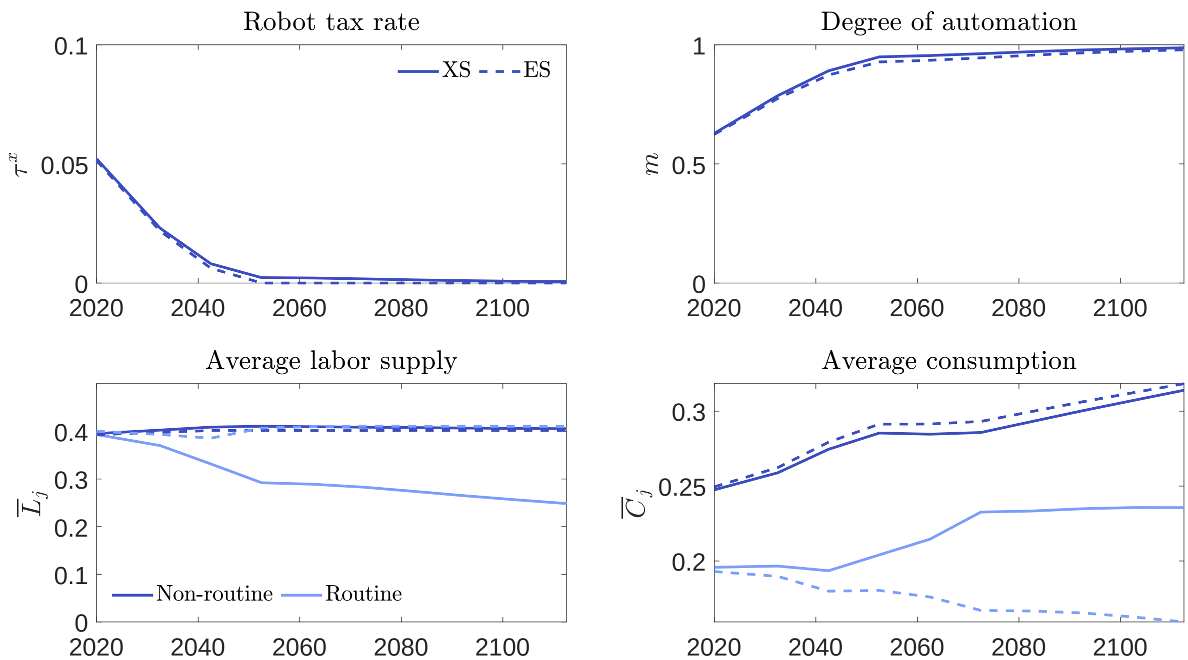

Notes: This figure compares the Mirrleesian optimal allocations in the baseline economy with endogenous skill acquisition (*ES*) with those of a counterfactual economy in which skills evolve exogenously (*XS*), but such that the shares of routine and non-routine workers are the same in the two economies. The first and second panels show the optimal robot tax and the level of automation in each economy, respectively. The third and fourth panels show the equilibrium levels of average labor supply and consumption in the cross section, respectively.

Time inconsistency In our model, optimal policy must take into account the effect of redistribution on workers' skill choices. In our calibration, we find that optimal direct redistribution is limited in order to induce more agents to invest in non-routine skills.

This optimal plan is inherently *time inconsistent*. Since skills are chosen once and for all when workers are young, older workers cannot readjust their skills. As a result, a plan that promises low redistribution to favor the acquisition of non-routine skills is only optimal *ex ante*, that is, before skill decisions have been made. The same planning problem starting at a later date would deviate from the original plan and would use the direct redistribution mechanism to redistribute income toward older routine workers (the extensive margin would no longer be relevant for those workers). As a result, robots would be taxed initially.

6 Conclusions

Our analysis suggests that without changes to the current U.S. tax system, a sizable fall in the costs of automation will lead to a massive rise in income inequality.

We study the problem of a planner that implements a nonlinear income tax system and linear robot taxes. Our model has an overlapping-generations structure that incorporates the life-cycle aspects of labor supply. Before entering the labor force, workers choose whether to acquire routine or non-routine skills. The cost of becoming a non-routine worker is heterogeneous across the population.

Designing an optimal tax system involves balancing two objectives. First, the planner wants to give the young generations incentives to invest in skills and become non-routine workers. Second, the planner wants to redistribute income toward routine workers, since their wages fall as robots become cheaper. Taxing robots reduces the non-routine wage premium and helps redistribute income toward routine workers.

In our calibrated economy, we find that it is optimal to tax robots while the initial old generations of routine workers are in the labor force. Once they retire, optimal robot taxes are zero. In other words, it is optimal to tax robots in the short run but not in the long run.

The world economy has undergone many structural changes that destroyed some jobs while creating others. Isn't the advent of robotization just another one of these changes? Why should public policy intervene this time? What makes this time different is the speed with which automation can occur. Many of the prior structural changes occurred slowly. The older generations kept their jobs, and it was their children who had to adapt to the brave new world. Automation can destroy many of the jobs held by the older generations and lead to a dramatic rise in income inequality. Public policy can avoid turning modern economies into the bleak world described in Kurt Vonnegut's *Player Piano*.

References

Acemoglu D. and Autor D. (2011) Skills, tasks and technologies: Implications for employment and earnings, in *Handbook of Labor Economics*, vol. 4, pp. 1043–1171, Elsevier.

Acemoglu D., Manera A. and Restrepo P. (2020) Does the us tax code favor automation?, Working Paper 27052, National Bureau of Economic Research.

Acemoglu D. and Restrepo P. (2018) Artificial intelligence, automation and work, Tech. rep., National Bureau of Economic Research.

Acemoglu D. and Restrepo P. (2019) Automation and new tasks: How technology displaces and reinstates labor, *Journal of Economic Perspectives* **33**(2), 3–30.

Acemoglu D. and Restrepo P. (2020) Robots and jobs: Evidence from us labor markets, *Journal of Political Economy* **128**(6), 2188–2244.

Adão R., Beraja M. and Pandalai-Nayar N. (2018) Skill-biased technological transitions, Working paper.

Ales L., Kurnaz M. and Sleet C. (2015) Technical change, wage inequality, and taxes, *The American Economic Review* **105**(10), 3061–3101.

Atkinson A. B. and Stiglitz J. E. (1972) The structure of indirect taxation and economic efficiency, *Journal of Public economics* **1**(1), 97–119.

Atkinson A. B. and Stiglitz J. E. (1976) The design of tax structure: Direct versus indirect taxation, *Journal of Public Economics* **6**(1-2), 55–75.

Auclert A. and Rognlie M. (2018) Inequality and aggregate demand, Working Paper 24280, National Bureau of Economic Research.

Autor D. H., Katz L. F. and Krueger A. B. (1998) Computing inequality: Have computers changed the labor market?, *The Quarterly journal of economics* **113**(4), 1169–1213.

Autor D. H., Levy F. and Murnane R. J. (2003) The skill content of recent technological change: An empirical exploration, *The Quarterly Journal of Economics* **118**(4), 1279–1333.

Benabou R. (2000) Unequal societies: Income distribution and the social contract, *The American Economic Review* **90**(1), 96–129.

Bresnahan T. F., Brynjolfsson E. and Hitt L. M. (2002) Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence, *The Quarterly Journal of Economics* **117**(1), 339–376.

CASELLI F. (1999) Technological revolutions, *The American Economic Review* **89**(1), 78–102.

CHAMLEY C. (1986) Optimal taxation of capital income in general equilibrium with infinite lives, *Econometrica* pp. 607–622.

CHARI V. V. and HOPENHAYN H. (1991) Vintage human capital, growth, and the diffusion of new technology, *Journal of Political Economy* **99**(6), 1142–1165.

CHARI V. V. and KEHOE P. J. (1999) Optimal fiscal and monetary policy, *Handbook of Macroeconomics* **1**, 1671–1745.

CHARI V. V., NICOLINI J. P. and TELES P. (2019) Optimal capital taxation revisited, *Journal of Monetary Economics* .

CHEN Y. G. (2019) The rise of the machines and the us labor market, Working paper.

CHETTY R. (2006) A new method of estimating risk aversion, *The American Economic Review* **96**(5), 1821–1834.

CHETTY R., GUREN A., MANOLI D. and WEBER A. (2011) Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins, *The American Economic Review* **101**(3), 471–75.

CORTES G. M., JAIMOVICH N., NEKARDA C. J. and SIU H. E. (2014) The micro and macro of disappearing routine jobs: A flows approach, Working Paper 20307, National Bureau of Economic Research.

CORTES G. M., JAIMOVICH N. and SIU H. E. (2017) Disappearing routine jobs: Who, how, and why?, *Journal of Monetary Economics* **91**, 69–87.

COSTINOT A. and WERNING I. (2018) Robots, trade, and luddism: A sufficient statistic approach to optimal technology regulation, Working Paper 25103, National Bureau of Economic Research.

DA COSTA C. E. and WERNING I. (2002) Commodity taxation and social insurance, Working paper.

DIAMOND P. A. and MIRRLEES J. A. (1971) Optimal taxation and public production i: Production efficiency, *The American Economic Review* **61**(1), 8–27.

FELDSTEIN M. S. (1969) The effects of taxation on risk taking, *Journal of Political Economy* **77**(5), 755–764.

FERRIERE A. and NAVARRO G. (2014) The heterogeneous effects of government spending: It's all about taxes .

FLOOD S., KING M., RODGERS R., RUGGLES S. and WARREN J. R. (2018) Integrated public use microdata series, current population survey: Version 6.0.

GOLOSOV M., KOCHERLAKOTA N. and TSYVINSKI A. (2003) Optimal indirect and capital taxation, *The Review of Economic Studies* **70**(3), 569–587.

GOMES R., LOZACHMEUR J.-M. and PAVAN A. (2018) Differential taxation and occupational choice, *The Review of Economic Studies* **85**(1), 511–557.

GOOS M., MANNING A. and SALOMONS A. (2014) Explaining job polarization: Routine-biased technological change and offshoring, *American Economic Review* **104**(8), 2509–26.

GOURINCHAS P.-O. and PARKER J. A. (2002) Consumption over the life cycle, *Econometrica* **70**(1), 47–89.

GUVENEN F., KARAHAN F., OZKAN S. and SONG J. (2015) What do data on millions of u.s. workers reveal about life-cycle earnings risk?, Working Paper 20913, National Bureau of Economic Research.

HEATHCOTE J., STORESLETTEN K. and VIOLANTE G. L. (2017) Optimal tax progressivity: An analytical framework, *The Quarterly Journal of Economics* **132**(4), 1693–1754.

JACOBS B. (2015) Optimal inefficient production, Working paper, Mimeo, Erasmus University Rotterdam.

JOHNSON M. and KEANE M. P. (2013) A dynamic equilibrium model of the us wage structure, 1968–1996, *Journal of Labor Economics* **31**(1), 1–49.

JONES L. E., MANUELLI R. E. and ROSSI P. E. (1997) On the optimal taxation of capital income, *Journal of Economic Theory* **73**(1), 93–117.

JUDD K. L. (1985) Redistributive taxation in a simple perfect foresight model, *Journal of Public Economics* **28**(1), 59–83.

KRUSELL P., OHANIAN L. E., RÍOS-RULL J.-V. and VIOLANTE G. L. (2000) Capital-skill complementarity and inequality: A macroeconomic analysis, *Econometrica* **68**(5), 1029–1053.

MCFADDEN D. (1974) Conditional logit analysis of qualitative choice behavior, in *Frontiers in Econometrics*, pp. 105–142, New York: Academic Press.

MIRRLEES J. A. (1971) An exploration in the theory of optimum income taxation, *The Review of Economic Studies* **38**(2), 175–208.

NAITO H. (1999) Re-examination of uniform commodity taxes under a non-linear income tax system and its implication for production efficiency, *Journal of Public Economics* **71**(2), 165–188.

PERSSON M. (1983) The distribution of abilities and the progressive income tax, *Journal of Public Economics* **22**(1), 73–88.

ROTHSCHILD C. and SCHEUER F. (2013) Redistributive taxation in the roy model, *The Quarterly Journal of Economics* **128**(2), 623–668.

ROYS N. A. and TABER C. R. (2019) Skill prices, occupations, and changes in the wage structure for low skilled men, Working Paper 26453, National Bureau of Economic Research.

SACHS D., TSYVINSKI A. and WERQUIN N. (2016) Nonlinear tax incidence and optimal taxation in general equilibrium, Working Paper 22646, National Bureau of Economic Research.

SAEZ E. (2004) Direct or indirect tax instruments for redistribution: Short-run versus long-run, *Journal of Public Economics* **88**(3-4), 503–518.

SCHEUER F. (2014) Entrepreneurial taxation with endogenous entry, *American Economic Journal: Economic Policy* **6**(2), 126–63.

SCHEUER F. and WERNING I. (2016) Mirrlees meets diamond-mirrlees, Working Paper 22076, National Bureau of Economic Research.

SLAVÍK C. and YAZICI H. (2014) Machines, buildings, and optimal dynamic taxes, *Journal of Monetary Economics* **66**, 47–61.

STIGLITZ J. E. (1982) Self-selection and pareto efficient taxation, *Journal of Public Economics* **17**(2), 213–240.

THUEMMEL U. (2018) Optimal taxation of robots, Cesifo working paper.

TSYVINSKI A. and WERQUIN N. (2017) Generalized compensation principle, Working Paper 23509, National Bureau of Economic Research.

WERNING I. (2007) Optimal fiscal policy with redistribution, *The Quarterly Journal of Economics* **122**(3), 925–967.

A Appendix

A.1 Appendix to Section 3

A.1.1 The first-best allocation

We define the first-best allocation in this economy as the solution that maximizes welfare, defined in equation (27), absent informational constraints. This absence implies that the planner can perfectly discriminate among agents and enforce any allocation that satisfies the aggregate resource constraint. The optimal plan solves the following problem

$$W = \max \omega_r \pi_r \left[u(c_r, l_r) + v(G) \right] + \omega_n \pi_n \left[u(c_n, l_n) + v(G) \right].$$

$$\pi_r c_r + \pi_n c_n + G \leq A \left[\int_0^m x_i^\rho di + \int_m^1 n_i^\rho di \right]^{\frac{1-\alpha}{\rho}} (\pi_n l_n)^\alpha - \phi \int_0^m x_i di, \quad [\mu],$$

$$\int_m^1 n_i di = \pi_r l_r, \quad [\eta].$$

The first-order conditions with respect to n_i and x_i are

$$\mu(1-\alpha) A \left[\int_0^m x_i^\rho di + \int_m^1 n_i^\rho di \right]^{\frac{1-\alpha}{\rho} - 1} (\pi_n l_n)^\alpha n_i^{\rho-1} = \eta, \quad \forall i \in (m, 1]$$

$$(1-\alpha) A \left[\int_0^m x_i^\rho di + \int_m^1 n_i^\rho di \right]^{\frac{1-\alpha}{\rho} - 1} (\pi_n l_n)^\alpha x_i^{\rho-1} = \phi, \quad \forall i \in [0, m].$$

The first equation implies that the marginal productivity of routine labor should be constant across the activities that use routine labor. This property means that $(1-m)\, n_i = \pi_r l_r$ for $i \in (m, 1]$ and $n_i = 0$ otherwise. The same property applies to robots used in the activities that are automated, $x_i = x$ for $i \in [0, m]$ and $x_i = 0$ otherwise.

To characterize the optimal allocations, we replace n_i and x_i in the planner's problem, which can be rewritten as

$$W = \max \omega_r \pi_r \left[u(c_r, l_r) + v(G)\right] + \omega_n \pi_n \left[u(c_n, l_n) + v(G)\right].$$

$$\pi_r c_r + \pi_n c_n + G \leq A \left[m x^\rho + (1-m)\left(\frac{\pi_r l_r}{1-m}\right)^\rho\right]^{\frac{1-\alpha}{\rho}} (\pi_n l_n)^\alpha - \phi m x, \quad [\mu].$$

The first-order conditions with respect to x and m are, respectively,

$$(1-\alpha) A \left[m x^\rho + (1-m)\left(\frac{\pi_r l_r}{1-m}\right)^\rho\right]^{\frac{1-\alpha}{\rho}-1} N_n^\alpha x^{\rho-1} = \phi,$$

$$\frac{1-\alpha}{\rho} A \left[m x^\rho + (1-m)\left(\frac{\pi_r l_r}{1-m}\right)^\rho\right]^{\frac{1-\alpha}{\rho}-1} N_n^\alpha \left[x^\rho - (1-\rho)\left(\frac{\pi_r l_r}{1-m}\right)^\rho\right] = \phi x.$$

The ratio of these two equations implies that if automation is positive, $m \in (0,1)$, then $x = \pi_r l_r/(1-m)$. Using this condition, we obtain

$$W = \max \omega_r \pi_r \left[u(c_r, l_r) + v(G)\right] + \omega_n \pi_n \left[u(c_n, l_n) + v(G)\right].$$

$$\pi_r c_r + \pi_n c_n + G \leq A \left(\frac{\pi_r l_r}{1-m}\right)^{1-\alpha} (\pi_n l_n)^\alpha - \phi m \frac{\pi_r l_r}{1-m}, \quad [\mu].$$

The first-order condition with respect to the level of automation implies that

$$(1-\alpha) A \frac{1}{(1-m)^{2-\alpha}} (\pi_r l_r)^{1-\alpha} (\pi_n l_n)^\alpha - \phi \frac{\pi_r l_r}{(1-m)^2} = 0 \Leftrightarrow m = 1 - \left[\frac{\phi}{A(1-\alpha)}\right]^{1/\alpha} \frac{\pi_r l_r}{\pi_n l_n},$$

provided that m is interior. Then,

$$m = \max\left\{1 - \left[\frac{\phi}{A(1-\alpha)}\right]^{1/\alpha} \frac{N_r}{N_n}, 0\right\}.$$

The first-order conditions with respect to c_r, c_n, l_r, l_n, and G are

$$\omega_r u_c(c_r, l_r) = \mu,$$

$$\omega_n u_c(c_n, l_n) = \mu,$$

$$\omega_r u_l(c_r, l_r) \geq \frac{\mu}{\pi_r l_r}(1-\alpha)(1-m)Y,$$
$$\omega_n u_l(c_n, l_n) = \mu \frac{\alpha Y}{\pi_n l_n},$$
$$v'(G) = \mu.$$

The first-order condition with respect to N_r is presented with inequality because the constraint $N_r \geq 0$ may bind when automation costs are low. The combination of the first two equations implies that

$$\omega_r u_c(c_r, l_r) = \omega_n u_c(c_n, l_n).$$

The optimal marginal rates of substitution are given by the combination of the marginal utility of consumption and leisure for each individual

$$\frac{u_l(c_r, l_r)}{u_c(c_r, l_r)} \geq (1-\alpha)(1-m)\frac{Y}{\pi_r l_r},$$
$$\frac{u_l(c_n, l_n)}{u_c(c_n, l_n)} = \alpha \frac{Y}{\pi_n l_n}.$$

Finally, from the first-order conditions for G and c_r it follows that

$$v'(G) = \omega_r u'(c_r). \tag{52}$$

A.1.2 Necessity and sufficiency in the static model

Worker optimality implies that the utility associated with the bundle of consumption and income assigned to agent j, $\{c_j, l_j\}$, must be at least as high as the utility associated with any other bundle $\{c, l\}$ that satisfies the budget constraint $c \leq w_j l - T(w_j l)$, implying that $u(c_j, l_j) \geq u(c, l)$. In particular, routine workers must prefer their bundle, $\{c_r, l_r\}$, to the bundle that they would get if they pretended to be non-routine workers while keeping the routine wage, $\{c_n, w_n l_n / w_r\}$. Similarly, non-routine workers must prefer their bundle, $\{c_n, l_n\}$, to the bundle they would get if they pretended to be routine workers, $\{c_r, w_r l_r / w_n\}$. These requirements correspond to the two incentive constraints (IC), (29) and (30), so these conditions are necessary.

We show in the Appendix that equation (28) is necessary by combining the first-order conditions to the firms' problems with the resource constraint, (11). In addition, we show that conditions (28), (29), and (30) are also sufficient. To see that equations (29) and (30) summarize the worker problem, note that it is possible to choose a tax function such that agents prefer the bundle $\{c_j, l_j\}$ to any other bundle. For example, the government could choose a tax function that sets the agent's after-tax income to zero for any choice of $w_j l$ different from $w_j l_j$, $j = n, r$. These results are summarized in the following proposition.

Lemma 3. *Equations (28), (30), and (29) characterize the set of implementable allocations. These conditions are necessary and sufficient for a competitive equilibrium.*

In an equilibrium, robot producers set the price of robots equal to their marginal cost

$$p_i = \phi. \tag{53}$$

Optimality for final goods producers implies that

$$x_i = \begin{cases} \frac{\pi_r l_r}{1-m}, & i \in [0, m], \\ 0, & \text{otherwise} \end{cases} \tag{54}$$

$$n_i = \begin{cases} \frac{\pi_r l_r}{1-m}, & i \in (m, 1], \\ 0, & \text{otherwise} \end{cases} \tag{55}$$

$$m = \max\left\{1 - \left[\frac{(1+\tau^x)\phi}{(1-\alpha)A}\right]^{1/\alpha} \frac{\pi_r l_r}{\pi_n l_n}, 0\right\}, \tag{56}$$

$$Y = A\left[\int_0^m x_i^\rho di + \int_m^1 n_i^\rho di\right]^{\frac{1-\alpha}{\rho}} (\pi_n l_n)^\alpha, \tag{57}$$

$$w_r = (1-\alpha)(1-m)\frac{Y}{\pi_r l_r}, \tag{58}$$

$$w_n = \alpha \frac{Y}{\pi_n l_n}. \tag{59}$$

The resource constraint is

$$\pi_r c_r + \pi_n c_n + G \leq Y - \int_0^m \phi x_i. \tag{60}$$

We can let equation (53) define the price of robots, let equation (54) define x_i, and let equations (55), (56), and (57) determine n_i, m, and Y, respectively. Assuming that m is interior, the wage equations (58) and (59) can be written as (14) and (15). These equations can be used to solve for the equilibrium wage rates. Combining the results above, we can write the resource constraint as

$$\pi_r c_r + \pi_n c_n + G \leq \alpha \frac{A^{1/\alpha}(1-\alpha)^{\frac{1-\alpha}{\alpha}}}{[(1+\tau^x)\phi]^{\frac{1-\alpha}{\alpha}}} \frac{\tau^x + \alpha}{\alpha(1+\tau^x)} \pi_n l_n + \phi \pi_r l_r.$$

Replacing the wage rates, we can write

$$\pi_r c_r + \pi_n c_n + G \leq \pi_n w_n l_n \frac{\tau^x + \alpha}{\alpha(1+\tau^x)} + \frac{\pi_r w_r l_r}{1+\tau^x}. \tag{61}$$

This derivation makes it clear that the resource constraint (61) summarizes the equilibrium conditions of the production side of the economy.

Worker optimality requires that

$$u\left(c_j, l_j\right) \geq u\left(c, l\right), \quad \forall (c, l) : c \leq w_j l - T(w_j l).$$

The following incentive constraints are necessary conditions:

$$u\left(c_n, l_n\right) \geq u\left(c_r, \frac{w_r}{w_n} l_r\right)$$
$$u\left(c_r, l_r\right) \geq u\left(c_n, \frac{w_n}{w_r} l_n\right).$$

These are also sufficient conditions because the planner can set the tax schedule $T(\cdot)$ such that for all $Y \notin \{Y_n, Y_r\}$, the allocation is worse for both agents than their respective allocation. This goal can be accomplished by setting

$$T(y) = y - \max\left\{c | u(c_i, l_i) \geq u\left(c, \frac{y}{w_i}\right), \text{for } i = r, n\right\}.$$

Since the government can choose an arbitrary tax function, it is bound only by the incentive constraints that characterize the informational problem. This property means that the income tax function that is assumed here to implement the optimal allocation is without loss of generality. Any other implementation would at least have to satisfy the same two incentive constraints.

A.1.3 Proof of Proposition 1

The allocations solve the original optimization problem, or equivalently they solve

$$W(\tau^x) = \max \pi_r \omega_r u(c_r, l_r) + \pi_n \omega_n u(c_n, l_n) + v(G)$$

subject to

$$\begin{aligned}
&[\eta_r \pi_r] \;\; u(c_r, l_r) \geq u\left(c_n, \frac{w_n}{w_r} l_n\right), \\
&[\eta_n \pi_n] \;\; u(c_n, l_n) \geq u\left(c_r, \frac{w_r}{w_n} l_r\right), \\
&[\mu] \;\; \pi_r c_r + \pi_n c_n + G \leq \pi_n w_n l_n \frac{\tau^x + \alpha}{\alpha(1+\tau^x)} + \pi_r \frac{w_r l_r}{1+\tau^x}.
\end{aligned}$$

Assume that the routine IC does not bind, then $\eta_r = 0$. The envelope condition is

$$W'(\tau^x) = -\eta_n \pi_n u_l\left(c_r, \frac{w_r}{w_r} l_r\right) \frac{d\log\left(w_r/w_n\right)}{d\log(1+\tau^x)} \frac{1}{1+\tau^x} \frac{w_r l_r}{w_n} + \mu \left[\begin{array}{c} \pi_n w_n l_n \frac{\tau^x+\alpha}{\alpha(1+\tau^x)^2}\left[\frac{d\log w_n}{d\log(1+\tau^x)} + \frac{1-\alpha}{\tau^x+\alpha}\right] \\ +\pi_r \frac{w_r l_r}{(1+\tau^x)^2}\left[\frac{d\log w_r}{d\log(1+\tau^x)} - 1\right] \end{array}\right]$$

Using the wages, we have that

$$
\begin{aligned}
w_r &= \phi(1+\tau^x) \Rightarrow \frac{d\log w_r}{d\log(1+\tau^x)} = 1, \\
w_n &= \alpha\frac{A^{1/\alpha}(1-\alpha)^{\frac{1-\alpha}{\alpha}}}{[(1+\tau^x)\phi]^{\frac{1-\alpha}{\alpha}}} \Rightarrow \frac{d\log w_n}{d\log(1+\tau^x)} = -\frac{1-\alpha}{\alpha}, \\
\frac{w_r}{w_n} &= \frac{[(1+\tau^x)\phi]^{\frac{1}{\alpha}}}{\alpha A^{1/\alpha}(1-\alpha)^{\frac{1-\alpha}{\alpha}}} \Rightarrow \frac{d\log w_r/w_n}{d\log(1+\tau^x)} = \frac{1}{\alpha}.
\end{aligned}
$$

Plugging these conditions into the envelope condition, we obtain

$$
\begin{aligned}
W'(\tau^x) &= -\eta_n\pi_n u_l\left(c_r,\frac{w_r}{w_n}l_r\right)\frac{1}{\alpha(1+\tau^x)}\frac{w_r l_r}{w_n} + \mu\pi_n w_n l_n\frac{\tau^x+\alpha}{\alpha(1+\tau^x)^2}\left[-\frac{1-\alpha}{\alpha}+\frac{1-\alpha}{\tau^x+\alpha}\right] \\
&= \frac{1}{\alpha(1+\tau^x)}\left[-\eta_n u_l\left(c_r,\frac{w_r}{w_n}l_r\right)\frac{w_r l_r}{w_n} - \mu\pi_n w_n l_n\frac{\tau^x}{1+\tau^x}\frac{1-\alpha}{\alpha}\right].
\end{aligned}
$$

Because $\mu > 0$, then if $\tau^x \leq 0$ we obtain that

$$
W'(\tau^x) > 0,
$$

so that the planner always improves its objective by marginally increasing τ^x. Since optimality implies that $W'(\tau^x) = 0$, then the optimal tax on robots satisfies the following condition:

$$
\frac{\tau^x}{1+\tau^x} = \frac{\alpha}{1-\alpha}\frac{\eta_n\left(-u_l\left(c_r,\frac{w_r}{w_n}l_r\right)\frac{w_r l_r}{w_n}\right)}{\mu w_n l_n}.
$$

The first-order condition with respect to l_r implies that

$$
-\frac{\eta_n}{\mu}u_l\left(c_r,\frac{w_r}{w_r}l_r\right)\frac{w_r l_r}{w_n} = \frac{\widetilde{\omega}_r\pi_r u_l(c_r,l_r)\,l_r + \frac{\pi_r w_r l_r}{1+\tau^x}}{\pi_n} = \frac{\pi_r\phi l_r}{\pi_n}\left[1-\frac{\widetilde{\omega}_r\left(-u_l(c_r,l_r)\right)}{\phi}\right],
$$

where $\widetilde{\omega}_r = \omega_r/\mu$. Replacing this equation in the optimal condition for τ^x we obtain

$$
\frac{\tau^x}{1+\tau^x} = \frac{\alpha}{1-\alpha}\frac{\pi_r\phi l_r}{\pi_n w_n l_n}\left[1-\frac{\widetilde{\omega}_r\left(-u_l(c_r,l_r)\right)}{\phi}\right].
$$

A.1.4 The full automation case ($m = 1$, $l_r = 0$)

If the optimal plan features $l_r = 0$, then it must be that $l_n > 0$. This result implies that $\psi = 0$. From the envelope condition, we can see that

$$W'(\tau^x) = -\frac{\mu}{\alpha\,(1+\tau^x)}\pi_n w_n l_n \frac{\tau^x}{1+\tau^x}\frac{1-\alpha}{\alpha} = 0 \Leftrightarrow \tau^x = 0. \tag{62}$$

A.2 Appendix to Section 4

A.2.1 Proof of Lemma 1

The extensive-margin incentive compatibility constraints can equivalently be written as

$$U_{\theta,t} \geq U_{\theta',t} + \theta(s_{\theta,t} - s_{\theta',t}) \tag{63}$$

for all t and $\theta, \theta' \in \Theta$.

Suppose that conditions (63) are satisfied. Then, take $\theta, \theta' \in \Theta_{r,t}$, that is, such that $s_t(\theta) = s_t(\theta')$. As a result, those conditions imply

$$\begin{aligned} U_t(\theta) &\geq U_t(\theta') \\ U_t(\theta') &\geq U_t(\theta), \end{aligned}$$

which is equivalent to $U_t(\theta) = U_t(\theta')$. This condition must hold for all $\theta, \theta' \in \Theta_{j,t}$ for $j = n, r$. Then, define $U_{j,t} \equiv U_t(\theta)$ for $\theta \in \Theta_{j,t}$, which implies that $U_t(\theta) = U_{j,t}$ for all $\theta \in \Theta_{j,t}$. Then, define $\theta_t^* \equiv U_{n,t} - U_{r,t}$. For all $\theta < \theta_t^*$ we have

$$U_{n,t} - \theta > U_{r,t}, \tag{64}$$

which implies that $s_t(\theta) = 1$. For all $\theta > \theta^*$ we have

$$U_{n,t} - \theta < U_{r,t}, \tag{65}$$

which implies that $s_t(\theta) = 0$.

To show the reverse implication, suppose that the conditions in the lemma hold. Then, for all $\theta \in \Theta_{n,t}$ we have

$$\begin{aligned} U_{\theta,t} &= U_{n,t} = U_{\theta',t}, \quad \forall \theta' \in \Theta_{n,t} \\ \widetilde{U}_{\theta,t} &= U_{n,t} - \theta \geq U_{n,t} - \theta^* = U_{r,t} = \widetilde{U}_{\theta',t}, \quad \forall \theta' \in \Theta_{r,t}. \end{aligned}$$

Instead, if $\theta \in \Theta_{r,t}$, then

$$\begin{aligned} U_{\theta,t} &= U_{n,t} = U_{\theta',t}, \quad \forall \theta' \in \Theta_{r,t} \\ \widetilde{U}_{\theta,t} &= U_{r,t} = U_{n,t} - \theta^* \geq U_{n,t} - \theta = U_{\theta',t} - \theta, \quad \forall \theta' \in \Theta_{n,t}. \end{aligned}$$

As a result, the allocation is extensive-margin incentive compatible (i.e., it satisfies (63)).

A.2.2 Proof of Lemma 2

The proof strategy is as follows: take an allocation for which the properties in the lemma do not hold, and show that there exists a perturbation that strictly improves welfare. We start by showing that this allocation frees up resources, then show that it can deliver an increase in government spending that improves utility. Finally, we check that it still satisfies all constraints.

Define $\Omega_{j,t} \equiv \pi_{j,t}\omega_{j,t}$, for $t = 2-L,...,0$, $\Omega_{j,t} \equiv \int_{\Theta_{j,t}} \omega_{\theta,t}h(\theta)d\theta$, and $\Omega_t = \sum_{a=0}^{L-1}\sum_{j=n,r}\Omega_{j,t}$. We can write the optimal program as

$$\max \sum_{t=2-L}^{\infty}\sum_{j=n,r}\beta^{\max\{0,t-1\}}\Omega_{j,t}\hat{U}_{j,t} + \sum_{t=1}^{\infty}\beta^{t-1}\Omega_t v(G_t) - \sum_{t=1}^{\infty}\int_0^{\theta_t^*}\beta^{t-1}\omega_{\theta,t}h(\theta)\theta d\theta \tag{66}$$

$$\hat{U}_{n,1-\tilde{a}} \geq \hat{U}_{r,1-\tilde{a}} + \sum_{a=\tilde{a}}^{L_w-1}\beta^{a-\tilde{a}}\left[\psi\left(l^a_{r,1+a-\tilde{a}}\right) - \psi\left(\frac{F_{r,1+a-\tilde{a}}}{F_{n,1+a-\tilde{a}}}l^a_{r,1+a\tilde{a}}\right)\right], \tilde{a} = 1,...,L_w-1 \tag{67}$$

$$\hat{U}_{r,1-\tilde{a}} \geq \hat{U}_{n,1-\tilde{a}} + \sum_{a=\tilde{a}}^{L_w-1}\beta^{a-\tilde{a}}\left[\psi\left(l^a_{n,1+a-\tilde{a}}\right) - \psi\left(\frac{F_{n,1+a-\tilde{a}}}{F_{r,1+a-\tilde{a}}}l^a_{n,1+a\tilde{a}}\right)\right], \tilde{a} = 1,...,L_w-1 \tag{68}$$

$$\hat{U}_{n,t} \geq \hat{U}_{r,t} + \sum_{a=0}^{L_w-1}\beta^a\left[\psi\left(l^a_{\theta,t+a}\right) - \psi\left(\frac{F_{r,t+a}}{F_{n,t+a}}l^a_{\theta,t+a}\right)\right], \theta \in \Theta_{r,t}, t = 1,2,... \tag{69}$$

$$\hat{U}_{r,t} \geq \hat{U}_{n,t} + \sum_{a=0}^{L_w-1}\beta^a\left[\psi\left(l^a_{\theta,t+a}\right) - \psi\left(\frac{F_{n,t+a}}{F_{r,t+a}}l^a_{\theta,t+a}\right)\right], \theta \in \Theta_{n,t}, t = 1,2,... \tag{70}$$

$$\theta_t^* = \hat{U}_{n,t} - \hat{U}_{r,t}, t = 1,2,... \tag{71}$$

$$\hat{U}_{j,1-a} = \sum_{t=1}^{L-a}\beta^{t-1}u\left(c^{a+(t-1)}_{j,t}\right) - \sum_{t=1}^{L_w-a}\beta^{t-1}\psi\left(l^{a+(t-1)}_{j,t}\right), a = 1,..,L-1,\ j = n,r \tag{72}$$

$$\hat{U}_{j,t} = \sum_{a=0}^{L-1}\beta^a u\left(c^a_{\theta,t+a}\right) - \sum_{t=1}^{L_w-1}\beta^a\psi\left(l^a_{\theta,t+a}\right), \theta \in \Theta_{j,t}, t = 1,2,...,\ j = n,r, \tag{73}$$

plus the resource constraint (38).

Labor supply Take an allocation that satisfies all the constraints and for which there exists a triplet (t,a,j) such that for every $\theta \in \Theta_{j,t-a}$, there exists a subset

$\Theta^* \subset \Theta_{j,t-a}$, such that: (i) $l^a_{\theta,t} \neq l^a_{\theta',t}$, for all $\theta' \in \Theta^*$, and (ii) $\int_{\Theta_j} h(\theta) d\theta > 0$. The first property simply requires some dispersion in allocations and the second property requires that this set has a non-null measure.

Consider the following perturbation: for all $\theta \in \Theta_{j,t-a}$, define their new labor supply as $\underline{l}^a_{\theta,t} = \underline{l}^a_{j,t} \equiv \int_{\Theta_{j,t-a}} l^a_{j,t} \frac{h(\theta)}{\int_{\Theta_{j,t-a}} h(\theta) d\theta} d\theta$. Then, construct their new consumption, $\underline{c}^a_{\theta,t} = \underline{c}^a_{j,t}$, such that if

$$\underbrace{U_{\theta,t-a}}_{=U_{j,t}} + u(\underline{c}^a_{j,t}) - u(c^a_{\theta,t}) - \psi(\underline{l}^a_{j,t}) + \psi(l^a_{\theta,t}) = U_{j,t} \Leftrightarrow u(\underline{c}^a_{j,t}) - u(c^a_{\theta,t}) = \psi(\underline{l}^a_{j,t}) - \psi(l^a_{\theta,t}), \tag{74}$$

that is, such that their utility is unchanged. Integrating on both sides, we obtain

$$u(\underline{c}^a_{j,t}) - \int_{\Theta_{j,t-a}} u(c^a_{\theta,t}) \frac{h(\theta)}{\int_{\Theta_{j,t-a}} h(\theta) d\theta} d\theta = \psi(\underline{l}^a_{j,t}) - \int_{\Theta_{j,t-a}} \psi(l^a_{\theta,t}) \frac{h(\theta)}{\int_{\Theta_{j,t-a}} h(\theta) d\theta} d\theta. \tag{75}$$

Since u is concave and ψ is convex, we know that

$$\begin{aligned}
&\int_{\Theta_{j,t-a}} u(c^a_{\theta,t}) \frac{h(\theta)}{\int_{\Theta_{j,t-a}} h(\theta) d\theta} d\theta \leq u\left(\int_{\Theta_{j,t-a}} c^a_{\theta,t} \frac{h(\theta)}{\int_{\Theta_{j,t-a}} h(\theta) d\theta} d\theta \right) \\
&\int_{\Theta_{j,t-a}} \psi(l^a_{\theta,t}) \frac{h(\theta)}{\int_{\Theta_{j,t-a}} h(\theta) d\theta} d\theta > \psi\left(\int_{\Theta_{j,t-a}} l^a_{\theta,t} \frac{h(\theta)}{\int_{\Theta_{j,t-a}} h(\theta) d\theta} d\theta \right) = \psi(\underline{l}^a_{j,t}).
\end{aligned}$$

As a result,

$$\begin{aligned}
u(\underline{c}^a_{j,t}) &= \int_{\Theta_{j,t-a}} u(c^a_{\theta,t}) \frac{h(\theta)}{\int_{\Theta_{j,t-a}} h(\theta) d\theta} d\theta + \psi(\underline{l}^a_{j,t}) - \int_{\Theta_{j,t-a}} \psi(l^a_{\theta,t}) \frac{h(\theta)}{\int_{\Theta_{j,t-a}} h(\theta) d\theta} d\theta \\
&< u\left(\int_{\Theta_{j,t-a}} c^a_{\theta,t} \frac{h(\theta)}{\int_{\Theta_{j,t-a}} h(\theta) d\theta} d\theta \right)
\end{aligned}$$

which implies that

$$\underline{c}^a_{j,t} < \int_{\Theta_{j,t-a}} c^a_{\theta,t} \frac{h(\theta)}{\int_{\Theta_{j,t-a}} h(\theta) d\theta} d\theta, \tag{76}$$

because u is increasing.

Intuitively, these results tell us that this perturbation leads to the same $\hat{U}_{j,t-a}$ for all agents but relaxes resources in the economy. There is no change in aggregate labor supply by the workers, but aggregate consumption is strictly lower. As a result, the government can increase spending,

$$\underline{G}_t = G_t + \left[\int_{\Theta_{j,t-a}} c^a_{\theta,t} \frac{h(\theta)}{\int_{\Theta_{j,t-a}} h(\theta) d\theta} d\theta - \underline{c}^a_{j,t} \right] > G_t,$$

which leads to an increase in welfare while still satisfying the resource constraints.

It remains to be shown that this allocation satisfies all other implementability conditions. Because we hold fixed $\{\hat{U}_{n,t}, \hat{U}_{r,t}\}$, the equations in (71) are satisfied for all t. Since we did not change allocations for agents born prior to $t = 1$, then (67), (68), and (72) are still satisfied. The equations in (73) are still satisfied because we imposed them to construct $\underline{c}^a_{j,t}$.

We just need to show that the intensive-margin incentive constraints, (70) and (71), are still satisfied. Because allocations do not change for other workers, they are satisfied for all workers not in occupation j and for all workers in occupation j born in periods other than $t - a$. So, we only need to show that

$$\hat{U}_{-j,t-a} \geq \hat{U}_{j,t-a} + \sum_{a'=0}^{L_w - 1} \beta^a \left[\psi \left(\underline{l}^{a'}_{\theta,t-a+a'} \right) - \psi \left(\frac{F_{j,t-a+a'}}{F_{-j,t-a+a'}} \underline{l}^{a'}_{\theta,t-a+a'} \right) \right] \tag{77}$$

for all $\theta \in \Theta_{j,t-a}$, where $\underline{l}^a_{\theta,t} = \underline{l}^a_{j,t}$ and $\underline{l}^{a'}_{\theta,t-a+a'} = l^{a'}_{\theta,t-a+a'}$ for $a' \neq a$.

Define $\Psi_{j,t}(l) \equiv \psi(l) - \psi\left(\frac{F_{j,t}}{F_{-j,t}} l\right)$. By convexity of ψ, $\Psi_{j,t}$ is increasing in l if $F_{j,t}/F_{-j,t} < 1$ and decreasing if $F_{j,t}/F_{-j,t} \geq 1$. Then, if $F_{j,t}/F_{-j,t} \geq 1$ take $\theta' \in \Theta_{j,t-a}$ such that $l^a_{\theta',t} < \underline{l}^a_{\theta',t}$, or if $F_{j,t}/F_{-j,t} < 1$ take $\theta' \in \Theta_{j,t-a}$ such that $l^a_{\theta',t} > \underline{l}^a_{\theta',t}$. This result implies that $\Psi\left(l^a_{\theta',t}\right) > \Psi\left(\underline{l}^a_{j,t}\right)$.

Since for the original allocations

$$\hat{U}_{-j,t-a} \geq \hat{U}_{j,t-a} + \sum_{a'=0}^{L_w - 1} \beta^a \left[\psi \left(l^{a'}_{\theta,t-a+a'} \right) - \psi \left(\frac{F_{j,t-a+a'}}{F_{-j,t-a+a'}} l^{a'}_{\theta,t-a+a'} \right) \right]$$

holds for all $\theta \in \Theta_{j,t-a}$, then

$$\begin{aligned}
\hat{U}_{-j,t-a} &\geq \hat{U}_{j,t-a} + \sum_{a'=0}^{L_w-1} \beta^a \left[\psi\left(l^{a'}_{\theta',t-a+a'} \right) - \psi\left(\frac{F_{j,t-a+a'}}{F_{-j,t-a+a'}} l^{a'}_{\theta',t-a+a'} \right) \right] \\
&= \hat{U}_{j,t-a} + \sum_{a'=0}^{L_w-1} \beta^a \left[\psi\left(l^{a'}_{\theta',t-a+a'} \right) - \psi\left(\frac{F_{j,t-a+a'}}{F_{-j,t-a+a'}} l^{a'}_{\theta',t-a+a'} \right) \right] + \underbrace{\left[\Psi\left(l^{a}_{\theta',t} \right) - \Psi\left(l^{a}_{j,t} \right) \right]}_{\geq 0} \\
&\geq \hat{U}_{j,t-a} + \sum_{a'=0}^{L_w-1} \beta^a \left[\psi\left(l^{a'}_{\theta',t-a+a'} \right) - \psi\left(\frac{F_{j,t-a+a'}}{F_{-j,t-a+a'}} l^{a'}_{\theta',t-a+a'} \right) \right].
\end{aligned}$$

Consumption We showed that if $s_{\theta,t} = s_{\theta',t}$, then $l^a_{\theta,t+a} = l^a_{\theta',t+a}$, for $a = 0, 1, ..., L_w - 1$. We now want to show that the same property holds for consumption. This result can be most easily seen from the first-order conditions, which imply that

$$\frac{u'(c^0_{\theta,t})}{u'(c^a_{\theta,t+a})} = \frac{u'(c^0_{\theta',t})}{u'(c^a_{\theta',t+a})} \tag{78}$$

for all θ, θ'. These equations, combined with the fact that (73) implies that

$$\sum_{a=0}^{L-1} \beta^a u\left(c^a_{\theta,t+a}\right) = \sum_{a=0}^{L-1} \beta^a u\left(c^a_{\theta',t+a}\right) \tag{79}$$

for θ, θ' such that $s_{\theta,t} = s_{\theta',t}$, delivers the intended result.

A.2.3 Proof of Proposition 2

Denote by $\beta^{t-1}\mu_t$ the multiplier for period t's resource constraint and $\beta^{\min\{0,t-1\}}\eta_{j,t}$ the multiplier on period $t \in \{2-L, ..., 0, 1, 2, ...\}$ for workers in occupation j's intensive-margin incentive constraint.

The first-order condition with respect to X_t is given by

$$\mu_t \left[F_{X,t} - \phi_t \right] + \sum_{a=0}^{L_w-1} \eta_{n,t-a} \psi' \left(\frac{F_{r,t}}{F_{n,t}} l^a_{r,t} \right) \frac{d\frac{F_{r,t}}{F_{n,t}}}{dX_t} l^a_{r,t} + \sum_{a=0}^{L_w-1} \eta_{r,t-a} \psi' \left(\frac{F_{n,t}}{F_{r,t}} l^a_{r,t} \right) \frac{d\frac{F_{n,t}}{F_{r,t}}}{dX_t} l^a_{n,t} = 0. \tag{80}$$

If the incentive constraints of routine workers do not bind, then $\eta_{r,t-a} = 0$. If the incentive constraints of at least one routine worker binds, then at least one $\eta_{n,t-a} > 0$, which implies that

$$\sum_{a=0}^{L_w-1} \frac{\eta_{n,t-a}}{\mu_t} \psi' \left(\frac{F_{r,t}}{F_{n,t}} l_{r,t}^a \right) l_{r,t}^a > 0.$$

As a result,

$$F_{X,t} = \phi_t - \frac{d\frac{F_{r,t}}{F_{n,t}}}{dX_t} \sum_{a=0}^{L_w-1} \frac{\eta_{n,t-a}}{\mu_t} \psi' \left(\frac{F_{r,t}}{F_{n,t}} l_{r,t}^a \right) l_{r,t}^a > \phi_t, \tag{81}$$

because $d\frac{F_{r,t}}{F_{n,t}}/dX_t < 0$.

A.2.4 Normalizing the dynamic model

In Table 3, we define the normalized variables, which are constant in the steady state

Table 4: Detrended variables

Parameter/Variable	*Original Variable*	*Normalized Variable*
Consumption	$c_{j,t}^a = \phi_t^{-\frac{1-\alpha}{\alpha}} \overline{c}_{j,t}^a$	$\overline{c}_{j,t}^a = \phi_t^{\frac{1-\alpha}{\alpha}} c_{j,t}^a$
Government spending	$G_t = \phi_t^{-\frac{1-\alpha}{\alpha}} \overline{G}_t$	$\overline{G}_t = \phi_t^{\frac{1-\alpha}{\alpha}} G_t$
Robots	$X_t = \phi_t^{-\frac{1}{\alpha}} \overline{X}_t$	$\overline{X}_t = \phi_t^{\frac{1}{\alpha}} X_t$
Output	$Y_t = \phi_t^{-\frac{1-\alpha}{\alpha}} \overline{Y}_t$	$\overline{Y}_t = \phi_t^{\frac{1-\alpha}{\alpha}} Y_t$
Net output	$NY_t = \phi_t^{-\frac{1-\alpha}{\alpha}} \overline{NY}_t$	$\overline{NY}_t = \phi_t^{\frac{1-\alpha}{\alpha}} NY_t$
Non-routine wage	$F_{n,t} = \phi_t^{-\frac{1-\alpha}{\alpha}} \overline{F}_{n,t}$	$\overline{F}_{n,t} = \phi_t^{\frac{1-\alpha}{\alpha}} w_{n,t}$
Routine wage	$F_{r,t} = \phi_t^{-\frac{1-\alpha}{\alpha}} \overline{F}_{r,t}$	$\overline{F}_{r,t} = \phi_t^{\frac{1-\alpha}{\alpha}} F_{r,t}$

With this normalization, we can write output, net output, and wages, relative wage, marginal productivity of robots, and automation as follows:

$$
\begin{aligned}
&\overline{Y}_t = A\left[\overline{X}^{\varepsilon} + \phi_t^{\frac{\varepsilon}{\alpha}} N_{r,t}^{\varepsilon}\right]^{\frac{1-\alpha}{\varepsilon}} N_{n,t}^{\alpha}, \qquad \overline{NY}_t = \overline{Y}_t - \overline{X}_t,\\
&\overline{F}_{n,t} = \alpha A\left[\overline{X}^{\varepsilon} + \phi_t^{\frac{\varepsilon}{\alpha}} N_{r,t}^{\varepsilon}\right]^{\frac{1-\alpha}{\varepsilon}} N_{n,t}^{\alpha-1}, \quad \overline{F}_{r,t} = (1-\alpha) A\left[\overline{X}^{\varepsilon} + \phi_t^{\frac{\varepsilon}{\alpha}} N_{r,t}^{\varepsilon}\right]^{\frac{1-\alpha}{\varepsilon}} N_{n,t}^{\alpha}\phi_t^{\frac{\varepsilon}{\alpha}} N_{r,t}^{\varepsilon-1},\\
&\frac{\overline{F}_{r,t}}{\overline{F}_{n,t}} = \frac{1-\alpha}{\alpha}\frac{\phi_t^{\frac{\varepsilon}{\alpha}} N_{r,t}^{\varepsilon}}{\overline{X}^{\varepsilon} + \phi_t^{\frac{\varepsilon}{\alpha}} N_{r,t}^{\varepsilon}}\frac{N_{n,t}}{N_{r,t}}, \qquad \overline{F}_{X,t} = (1-\alpha) A\left[\overline{X}^{\varepsilon} + \phi_t^{\frac{\varepsilon}{\alpha}} N_{r,t}^{\varepsilon}\right]^{\frac{1-\alpha}{\varepsilon}} N_{n,t}^{\alpha}\overline{X}_t^{\varepsilon-1},\\
&m_t = \frac{\overline{X}_t^{\varepsilon}}{\overline{X}^{\varepsilon} + \phi_t^{\frac{\varepsilon}{\alpha}} N_{r,t}^{\varepsilon}}.
\end{aligned}
$$

The detrended optimization problem Define $\Omega_{j,t} \equiv \pi_{j,t}\omega_{j,t}$, for $t = 2-L,...,0$, $\Omega_{j,t} \equiv \int_{\Theta_{j,t}} \omega_{\theta,t}h(\theta)d\theta$, and $\Omega_t = \sum_{a=0}^{L-1}\sum_{j=n,r}\Omega_{j,t}$. We can define the optimal program as

$$
\begin{aligned}
\max\Bigg\{ &\sum_{\widetilde{a}=1}^{L-1}\sum_{j=n,r}\Omega_{j,t-\widetilde{a}}\left\{\sum_{a=\widetilde{a}}^{L-1}\beta^{a-\widetilde{a}}u\left(c_{j,1+a-\widetilde{a}}^{a}\right) - \sum_{a=\widetilde{a}}^{L_w-1}\beta^{a-\widetilde{a}}\psi\left(l_{j,1+a-\widetilde{a}}^{a}\right)\right\}\\
&+\sum_{t=1}^{\infty}\sum_{j=n,r}\beta^{t-1}\Omega_{j,t}\left\{\sum_{a=0}^{L-1}\beta^{a}u\left(c_{j,t+a}^{a}\right) - \sum_{t=1}^{L_w-1}\beta^{a}\psi\left(l_{j,t+a}^{a}\right)\right\}\\
&+\sum_{t=1}^{\infty}\beta^{t-1}\Omega_t v(G_t) - \sum_{t=1}^{\infty}\int_0^{\theta_t^*}\beta^{t-1}\omega_{\theta,t}h(\theta)\theta d\theta\Bigg\} \quad \text{s.to}
\end{aligned}
$$

$$
\begin{aligned}
[\eta_{j,1-a}] \quad &\sum_{a=\widetilde{a}}^{L-1}\beta^{a-\widetilde{a}}u\left(c_{j,1+a-\widetilde{a}}^{a}\right) - \sum_{a=\widetilde{a}}^{L_w-1}\beta^{a-\widetilde{a}}\psi\left(l_{j,1+a-\widetilde{a}}^{a}\right) \geq\\
&\sum_{a=\widetilde{a}}^{L-1}\beta^{a-\widetilde{a}}u\left(c_{-j,1+a-\widetilde{a}}^{a}\right) - \sum_{a=\widetilde{a}}^{L_w-1}\beta^{a-\widetilde{a}}\psi\left(\frac{F_{-j,1+a-\widetilde{a}}}{F_{j,1+a-\widetilde{a}}}l_{-j,1+a-\widetilde{a}}^{a}\right)\\
[\beta^{t-1}\eta_{j,t}] \quad &\sum_{a=0}^{L-1}\beta^{a}u\left(c_{j,t+a}^{a}\right) - \sum_{t=1}^{L_w-1}\beta^{a}\psi\left(l_{j,t+a}^{a}\right) \geq \sum_{a=0}^{L-1}\beta^{a}u\left(c_{-j,t+a}^{a}\right) - \sum_{a=0}^{L_w-1}\beta^{a}\psi\left(\frac{F_{-j,t+a}}{F_{j,t+a}}l_{-j,t+a}^{a}\right)\\
[\beta^{t-1}\xi_t] \quad &\theta_t^* = \hat{U}_{n,t} - \hat{U}_{r,t}\\
[\beta^{t-1}\mu_t] \quad &\sum_{a=0}^{L-1}\sum_{j=n,r}\pi_{j,t-a}c_{j,t}^{a} + G_t = A\left[\overline{X}^{\varepsilon} + \phi_t^{\frac{\varepsilon}{\alpha}} N_{r,t}^{\varepsilon}\right]^{\frac{1-\alpha}{\varepsilon}} N_{n,t}^{\alpha} - \overline{X}_t,
\end{aligned}
$$

We define the Lagrange multipliers in parentheses.

A.2.5 Optimal policy steady state

In what follows, we derive the steady state for all variables assuming that the aggregate allocations converge, that is, assume that the allocations for aggregate consumption, $\overline{C}_t$, aggregate labor supply $\overline{N}_{j,t}$ for $j = n, r$, robots X_t, and θ_t^*, and government spending converge to an interior steady state.

Since $\theta_t^* \to \theta^*$, then $\pi_{n,t} = H(\theta_t^*) \to H(\theta^*) \equiv \pi_n$, and $\pi_{r,t} \to \pi_t \equiv 1 - \pi_n$. Furthermore, since

$$\Omega_t v'(G_t) = \mu_t,$$

and both $\Omega_t \to \Omega$ and $G_t \to G$, then $\mu_t \to \mu$.

Along the balanced-growth path, $\phi_t \to 0$ and $\overline{F}_{r,t}/\overline{F}_{n,t} \to 0$. As a result, the incentive compatibility of routine workers can never bind. This property is shown in the following lemma.

Lemma 4. *Suppose that the allocations converge to a steady-state growth path with interior automation, then $\eta_{r,t-a} \to 0$, for all a.*

Proof. Since $\overline{\phi}_t \to 0$, the optimal labor supply by agents with routine skills is $l_{r,t} = 0$. This property implies that the utility of a worker with routine skills converges to

$$U_{r,t} \to \sum_{a=0}^{L-1} \beta^a \left[u(\overline{c}_r^a) + v(G)\right] - \sum_{a=0}^{L_w-1} \beta^a \psi(0) \equiv U_r,$$

while the utility from mimicking a non-routine worker converges to $-\infty$ since it must be that $l_n > 0$:

$$\lim_{t\to\infty} \sum_{a=0}^{L-1} \beta^a \left[u(\overline{c}_{n,t+a}^a) + v(G)\right] - \sum_{a=0}^{L_w-1} \beta^a \psi\left(\frac{\overline{F}_{n,t+a}}{\overline{F}_{r,t+a}} l_{r,t+a}^a\right) = -\infty,$$

as $\frac{\overline{F}_{n,t}}{\overline{F}_{r,t}} \to +\infty$. □

The first-order conditions with respect to consumption when young are given by

$$u'(\overline{c}_{n,t}^a)[\Omega_{n,t} + (\xi_{t-a} + \eta_{n,t-a})] = \mu_t \pi_{n,t}$$
$$u'(\overline{c}_{r,t}^a)[\Omega_{r,t} - (\xi_{t-a} + \eta_{n,t-a})] = \mu_t \pi_{r,t}$$

Using these expressions, we can make two important observations. First,

$$\frac{u'(c^0_{j,t})}{u'(c^a_{j,t+a})} = \frac{\mu_t}{\mu_{t+a}} \to 1,$$

which implies that in the steady state, $c^a_j = c^0_j$ for all $a = 1, ..., L-1$. Second, these expressions imply that $\psi_t + \eta_{n,t} \to \kappa$, for some $\kappa \in \mathbb{R}$.

Because the detrended marginal productivity of routine workers falls to zero, then $l^a_{r,t} \to 0$ for all a. As a result, in the steady state, the labor supply of non-routine workers of age τ is given by the following condition:

$$v'(l^a_n)(\Omega_n + \underbrace{\xi_t + \eta_{n,t}}_{=\kappa}) = \pi_n \mu \overline{F}_n e_a.$$

As a result, $l^a_{n,t} \to l^a_n$.

The marginal condition with respect to robots is simply

$$\overline{F}_X = 1 \Leftrightarrow \overline{X}_t = [(1-\alpha)A]^{1/\alpha} N_n,$$

and

$$\overline{F}_n = \alpha A^{1/\alpha}(1-\alpha)^{\frac{1-\alpha}{\alpha}}.$$

From the first-order condition, we then obtain that $\xi_t \to \xi$, which solves

$$\xi = h(\theta^*)\mu \left[\sum_{a=0}^{L-1} \beta^a \left(c^a_r - c^a_n\right) + \sum_{a=0}^{L_w-1} \beta^a \overline{F}_n e_a l^a_n \right],$$

which also implies that $\eta_{n,t} \to \eta_n$.

The necessary and sufficient conditions to solve for an interior steady state are the following.

1. Consumption and government spending:

$$u'(c^a_r)(\Omega_r - \xi - \eta_n) = \mu \pi_r, \qquad u'(c^a_n)(\Omega_n + \xi + \eta_n) = \mu \pi_n \tag{82}$$

and

$$C = \sum_{a=0} (c^a_r + c^a_n), \qquad \Omega v'(G) = \mu. \tag{83}$$

2. Labor supply:

$$l_r^a = 0, \qquad \psi'(l_n^a)(\Omega_n + \psi + \eta_n) = \mu \overline{F}_n \pi_n e_a. \tag{84}$$

and

$$N_r = 0, \qquad N_n = \sum_{a=0}^{L_w - 1} \pi_n e_a l_n^a \tag{85}$$

3. Robots:

$$\overline{X} = \left[(1-\alpha)A\right]^{1/\alpha} N_n. \tag{86}$$

4. Skill acquisition cutoff θ^*

$$\xi = h(\theta^*)\mu \left[\sum_{a=0}^{L-1} \beta^a \left(c_r^a - c_n^a\right) + \sum_{a=0}^{L_w-1} \beta^a \overline{F}_n e_a l_n^a \right],$$

5. Intensive-margin incentive compatibility, which need not necessarily bind:

$$\theta^* \geq 0, \qquad \eta_n \geq 0, \qquad \eta_n \theta^* = 0. \tag{87}$$

6. Extensive-margin incentive compatibility:

$$\theta^* = \sum_{a=0}^{L-1} \beta^a \left[u(c_n^a) - u(c_r^a)\right] - \sum_{a=0}^{L_w-1} \beta^a \left[\psi\left(l_n^a\right) - \psi(0)\right]. \tag{88}$$

7. Resource constraint:

$$C + G = \overline{F}_n N_n. \tag{89}$$

We have $9 + 2L + 2L_w$ equations in

$$\{\{c_n^a, c_r^a\}_{a=0,\ldots,L-1}, \{l_n^a, l_r^a\}_{a=0,\ldots,L_w-1}, C, G, N_r, N_n, \overline{X}, \theta^*, \mu, \psi, \eta_n\},$$

which are $9 + 2L + 2L_w$ unknowns.

A.2.6 Proof of Proposition 3

In the balanced-growth path, we have that $\eta_{r,-a} = 0$ and $l_r^a = 0$. As a result,

$$\mu\left[\overline{F}_X - 1\right] + \sum_{a=0}^{L_w-1} \eta_{n,-a}\psi'\left(\frac{\overline{F}_r}{\overline{F}_n} l_r^a\right) \frac{d\frac{\overline{F}_r}{\overline{F}_n}}{d\overline{X}_t} l_r^a + \sum_{a=0}^{L_w-1} \eta_{r,-a}\psi'\left(\frac{\overline{F}_n}{\overline{F}_r} l_n^a\right) \frac{d\frac{\overline{F}_n}{\overline{F}_r}}{d\overline{X}_t} l_n^a = 0$$

reduces to

$$\overline{F}_X = 1.$$

A.3 Appendix to Section 5

A.3.1 Status quo equilibrium equations

Below, we summarize the equilibrium equations for our model. We define the following variables:

$$q_{t,t+a} \equiv \prod_{s=0}^{a-1} R_{t+s}^{-1}, \tag{90}$$

for $a > 0$, and $q_{t,t} \equiv 1$.

Workers born at $t \geq 1$ The consumption policy function is

$$c_{j,t+a}^{a} = \frac{\beta^{a}}{q_{t,t+a}} \frac{1-\beta}{1-\beta^{T}} \mathcal{W}_{i,t}^{0} \tag{91}$$

for $t \geq 0$ and $a = 0, ..., L-1$. Here,

$$\mathcal{W}_{j,t}^{0} \equiv \sum_{a=0}^{L_w - 1} q_{t,t+a} \lambda_{t+a} \left(w_{j,t+a} e_a l_{j,t+a}^{a} \right)^{1-\gamma}. \tag{92}$$

The labor supply is given by

$$l_{j,t+a}^{a} = \left[\frac{q_{t,t+a}}{\beta^{a}} \frac{\lambda_{t+a} \left(w_{j,t+a} e_a \right)^{1-\gamma}}{\lambda_t \left(w_{j,t} e_0 \right)^{1-\gamma}} \right]^{\frac{1}{\nu+\gamma}} l_{j,t}^{0} \tag{93}$$

and

$$l_{j,t}^{0} = \left[\frac{1-\beta^{L}}{1-\beta} \frac{1-\gamma}{\zeta} \frac{1}{\sum_{s=0}^{L_w - 1} \beta^{-a\frac{1-\gamma}{\nu+\gamma}} \left[\frac{q_{t,t+s} \lambda_{t+s} \left(w_{j,t+s} e_s \right)^{1-\gamma}}{\lambda_t \left(w_{j,t} e_0 \right)^{1-\gamma}} \right]^{\frac{1+\nu}{\nu+\gamma}}} \right]^{\frac{1}{1+\nu}} \tag{94}$$

for $t \geq 0$ and $a = 0, ..., L_w - 1$.

Solving for $c_{j,t+a}^{a}$ and $l_{j,t+a}^{a}$ we can compute asset holdings recursively,

$$b_{j,t+a}^{a} = R_{t+a} \left[b_{j,t+a-1}^{a-1} + \lambda_{t+a} \left(w_{j,t+a} e_a l_{j,t+a}^{a} \right)^{1-\gamma} - c_{j,t+a}^{a} \right] \tag{95}$$

for $a = 0, 1, ..., L_w - 1$ using the fact that $b^{-1}_{j,t-1} = 0$, and

$$b^a_{j,t+a} = R_{t+a}\left[b^{a-1}_{j,t+a-1} - c^a_{j,t+a}\right] \tag{96}$$

for $a = L_w, ..., L - 1$.

Skill acquisition is determined by a threshold rule, which implies that the share of non-routine workers is given by

$$\pi_{n,t} = H(\theta^*_t), \qquad \pi_{r,t} = 1 - \pi_{n,t}, \tag{97}$$

where

$$\theta^*_t = \sum_{a=0}^{L-1} \beta^a \log c^a_{n,t+a} - \sum_{a=0}^{L_w-1} \beta^a v\left(l^a_{n,t+a}\right) - \sum_{a=0}^{L-1} \beta^a \log c^a_{r,t+a} + \sum_{a=0}^{L_w-1} \beta^a v\left(l^a_{r,t+a}\right). \tag{98}$$

Workers born at $t = 2 - L_w, ..., 0$ The consumption policy function is

$$c^{a+s-1}_{j,s} = \frac{\beta^{s-1}}{q_{1,s}} \frac{1-\beta}{1-\beta^{L-a}} \mathcal{W}^a_{j,1}, \tag{99}$$

for $s = 1, ..., L - a$ and $a = 1, ..., L_w - 1$, where

$$\mathcal{W}^a_{j,1} = \sum_{s=1}^{L_w-a} q_{1,s}\lambda_s \left(w_{j,s}e_{a+s-1}l^{a+s-1}_{j,s}\right)^{1-\gamma} + R_0 b^{a-1}_{j,0}. \tag{100}$$

Here, $b^{a-1}_{j,0}$ denotes the exogenous level of financial wealth with which these agents enter the economy.

Labor supply is given by

$$l^{a+s-1}_{j,s} = \left[\frac{q_{1,s}}{\beta^{s-1}} \frac{\lambda_s \left(w_{j,s}e_{a+s-1}\right)^{1-\gamma}}{\lambda_1 \left(w_{j,1}e_a\right)^{1-\gamma}}\right]^{\frac{1}{\gamma+\nu}} l^a_{j,1}, \tag{101}$$

for $s = 1, ..., L_w - a$ and $a = 1, ..., L_w - 1$, and

$$l^a_{j,1} = \left[\frac{1-\gamma}{\zeta}\frac{1-\beta^{L-a}}{1-\beta}\frac{1}{\sum_{s=1}^{L_w-a}\beta^{-(s-1)\frac{1-\gamma}{\gamma+\nu}}\left[q_{1,s}\frac{\lambda_s\left(w_{j,s}e_{a+s-1}\right)^{1-\gamma}}{\lambda_1\left(w_{j,1}e_a\right)^{1-\gamma}}\right]^{\frac{1+\nu}{\gamma+\nu}} + \frac{b^{a-1}_{j,0}}{\lambda_1\left(w_{j,1}e_a l^a_{j,1}\right)^{1-\gamma}}}\right]^{\frac{1}{1+\nu}}, \tag{102}$$

for $a = 1, ..., L_w - 1$.

Solving for $c^{a}_{j,t+a}$ and $l^{a}_{j,t+a}$ we can calculate asset holdings recursively,

$$b^{a+s-1}_{j,s} = R_s \left[b^{a+s-2}_{j,s-1} + \lambda_s \left(w_{j,s} e_{a+s-1} l^{a+s-1}_{j,s} \right)^{1-\gamma} - c^{a+s-1}_{j,s} \right] \tag{103}$$

for $s = 1, ..., L_w - a$ using the fact that $b^{a-1}_{j,0}$ is exogenous, and

$$b^{a+s-1}_{j,s} = R_s \left[b^{a+s-2}_{j,s-1} - c^{a+s-1}_{j,s} \right] \tag{104}$$

for $s = L_w - a + 1, ..., L - a$. This procedure is used for $a = 1, ..., L_w - 1$.

Workers born at $t = 2 - L, ..., 1 - L_w$ The consumption policy function is

$$c^{\widetilde{a}-1+s}_{j,s} = \frac{\beta^{s-1}}{q_{1,s}} \frac{1-\beta}{1-\beta^{L-\widetilde{a}}} b^{\widetilde{a}-1}_{j,0}, \tag{105}$$

for $s = 1, ..., L - \widetilde{a}$ and $\widetilde{a} = L_w, ...L - 1$.

We can solve for asset holdings recursively,

$$b^{\widetilde{a}+s-1}_{j,s} = R_s \left[b^{\widetilde{a}+s-2}_{j,s-1} - c^{\widetilde{a}+s-1}_{j,s} \right] \tag{106}$$

using the fact that $b^{\widetilde{a}-1}_{j,0}$ is exogenous, for $s = 1, ..., L - \widetilde{a}$. This procedure is used for $\widetilde{a} = 1, ..., L_w - 1$.

Firm's problem The first-order conditions with respect to routine labor, robots and non-routine labor are as follows:

$$w_{r,t} = (1-\alpha) A \left[X_t^{\frac{\varepsilon-1}{\varepsilon}} + N_{r,t}^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}(1-\alpha)-1} N^{\alpha}_{n,t} N_{r,t}^{-\frac{1}{\varepsilon}}, \tag{107}$$

$$\phi_t = (1-\alpha) A \left[X_t^{\frac{\varepsilon-1}{\varepsilon}} + N_{r,t}^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}(1-\alpha)-1} N^{\alpha}_{n,t} X_t^{-\frac{1}{\varepsilon}}, \tag{108}$$

$$w_{n,t} = \alpha A \left[X_t^{\frac{\varepsilon-1}{\varepsilon}} + N_{r,t}^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}(1-\alpha)} N^{\alpha-1}_{n,t}. \tag{109}$$

Government's budget constraint

$$G_t + B_{t-1} = \sum_{a=0}^{L_w-1} \sum_{j=n,r} \pi_{t-a} \left\{ w_{j,t} e_a l_{j,t}^a - \lambda_t \left(w_{j,t} e_a l_{j,t}^a \right)^{1-\gamma} \right\} + \tau^x \phi_t X_t + \frac{B_t}{R_t}. \quad (110)$$

Market clearing The market-clearing condition for aggregate labor supply is

$$N_{j,t} = \sum_{a=0}^{L_w-1} \pi_{j,t-a} e_a l_{j,t}^a. \quad (111)$$

Aggregate consumption, C_t, is given by

$$C_t = \sum_{a=0}^{L-1} \sum_{j=r,n} \pi_{j,t-a} c_{j,t}^a. \quad (112)$$

The goods market-clearing condition is

$$C_t + G_t = F(X_t, N_{r,t}, N_{n,t}) - \phi_t X_t. \quad (113)$$

The asset market-clearing condition is

$$\sum_{a=0}^{L-2} \sum_{j=n,r} \pi_{j,t-a} b_{j,t}^a = B_t. \quad (114)$$

Table 5 lists the normalized variables used in our analysis.

Table 5: Normalized variables

Parameter/Variable	*Original Variables*	*Normalized variables*
Tax level	$\lambda_t = \phi_t^{-\gamma\frac{1-\alpha}{\alpha}} \overline{\lambda}_t$	$\overline{\lambda}_t = \phi_t^{\gamma\frac{1-\alpha}{\alpha}} \lambda_t$
Government bonds	$B_t = \phi_{t+1}^{-\frac{1-\alpha}{\alpha}} \overline{B}_t$	$\overline{B}_t = \overline{\phi}_{t+1}^{\frac{1-\alpha}{\alpha}} B_t$
Government bonds 2	$\frac{B_t}{C_t} = b_t$	$\frac{\overline{B}_t}{\overline{C}_t} = b_t e^{-\frac{1-\alpha}{\alpha} g_\phi}$
Real interest rate	$R_t = e^{g_\phi \frac{1-\alpha}{\alpha}} \overline{R}_t$	$\overline{R}_t = e^{-g_\phi \frac{1-\alpha}{\alpha}} R_t$
Discount factor	$q_{t,t+a} = e^{-g_\phi \frac{1-\alpha}{\alpha}} \overline{q}_{t,t+a}$	$\overline{q}_{t,t+a} = e^{a g_\phi \frac{1-\alpha}{\alpha}} q_{t,t+a}$
Initial period assets	$b_{j,0}^{a-1} = \phi_1^{-\frac{1-\alpha}{\alpha}} \overline{b}_{j,0}^{a-1}$	$\overline{b}_{j,0}^{a-1} = \phi_1^{\frac{1-\alpha}{\alpha}} b_{j,0}^{a-1}$
Initial government bonds	$B_0 = \phi_1^{-\frac{1-\alpha}{\alpha}} \overline{B}_0$	$\overline{B}_0 = \phi_1^{\frac{1-\alpha}{\alpha}} B_0$
Present value wealth	$\mathcal{W}_{j,t}^a = \phi_t^{-\frac{1-\alpha}{\alpha}} \overline{\mathcal{W}}_{j,t}^a$	$\overline{\mathcal{W}}_{j,t}^a = \phi_t^{\frac{1-\alpha}{\alpha}} \mathcal{W}_{j,t}^a$

Pre-automation steady state We start by solving for the non-detrendred variables in a steady state with $X = 0$. For this initial steady state, we take occupations as being exogenous and set $\pi_n = 0.4356$ constant across generations to match the data on occupations. We also set $G/C = 0.2126$ and $B/C = 0.0427$.

We use a root finding algorithm to calibrate the discount factor β, the labor disutility parameter ζ, the share of non-routine workers in production α, and the steady-state interest rate R, so that

$$\beta R = 1,$$
$$\sum_{a=0}^{L_w-1} \sum_{j=n,r} \pi_j l_j^a = 1/3,$$
$$\frac{w_n}{w_r} = 1.1943.$$

The assumption that $\beta R = 1$ is consistent with Gourinchas and Parker (2002), the second condition implies that the average labor supply is equal to one-third, and the final condition implies that the non-routine wage premium matches that in the data.

Transition dynamics First, we solve for the final steady state equilibrium. In this steady state, $\pi_n = 1$ and l_r^τ for all τ. As a result,

$$\overline{X} = [(1-\alpha)\, A]^{1/\alpha}\, N_n,$$

and

$$\overline{w}_r = 0, \qquad \overline{w}_n = \alpha A^{1/\alpha}(1-\alpha)^{\frac{1-\alpha}{\alpha}}.$$

Computing this steady state requires iterating on the equilibrium real interest rate. In checking for a steady-state equilibrium, we always look for a solution that satisfies asset market clearing rather than goods market clearing. From Walras' law, if the asset market clears, then so does the goods market. Auclert and Rognlie (2018) note that searching for a solution that only satisfies the goods market clearing condition can be problematic. The reason is that, in a steady state, satisfying the goods market clearing and every budget constraint implies asset market clearing only if the interest rate is not one (i.e., $R \neq 1$).

We use the transition to calibrate the parameters $\{\widetilde{\phi}, g_\phi, \varepsilon, \mu, \sigma\}$. These parameters are chosen so that the competitive equilibrium matches the wage premium and occupation shares. The calibration procedure minimizes a sum of square deviations between the equilibrium and the data.

To solve for the equilibrium, we take the asset distribution and occupation shares in the initial steady state and compute a perfect foresight transition to this final steady state. We assume that convergence occurs after T periods (for our baseline exercise, we set $T = 50$). Given $\{w_{n,t}, w_{r,t}, R_t, \lambda_t\}_{t=0}^{T}$, we can solve every household problem to obtain consumption and labor at all periods and ages. As a result, we can also solve for θ_t^* for all t, and then $\pi_{n,t}$ and $\pi_{r,t}$. We can also back out $b_{j,t}^\tau$.

Aggregating these variables, we obtain $C_t = \sum_{\tau=0}^{L-1}\sum_{j=n,r}\pi_{j,t-\tau}c_{j,t}^\tau$ and $N_{j,t} = \sum_{\tau=0}^{L-1}\pi_{j,t-\tau}e_\tau l_{j,t}^\tau$. Because we fixed the ratio of spending to consumption and the ratio of debt to consumption, we can use C_t to also back out G_t and B_t for all t. We can also use the first-order condition with respect to robots to solve for X_t.

Given these solutions and the initial guesses $\{w_{n,t}, w_{r,t}, R_t, \lambda_t\}_{t=0}^{T}$, we check a set of four equations in every period:

$$
\begin{aligned}
\Delta_{1,t} &\equiv \overline{w}_{r,t} - \phi_t^{\frac{\varepsilon}{\alpha}} (1-\alpha) A \left[\overline{X}_t^{\varepsilon} + \phi_t^{\frac{\varepsilon}{\alpha}} N_{r,t}^{\varepsilon} \right]^{\frac{1-\alpha}{\varepsilon}-1} N_{n,t}^{\alpha} N_{r,t}^{\varepsilon-1}, \\
\Delta_{2,t} &\equiv \overline{w}_{n,t} - \alpha A \left[\overline{X}_t^{\varepsilon} + \phi_t^{\frac{\varepsilon}{\alpha}} N_{r,t}^{\varepsilon} \right]^{\frac{1-\alpha}{\varepsilon}} N_{n,t}^{\alpha-1}, \\
\Delta_{3,t} &\equiv \overline{G}_t + \tau^x \overline{X}_t + \overline{B}_{t-1} - \sum_{a=0}^{L_w-1} \sum_{j=n,r} \pi_{t-a} \left\{ \overline{w}_{j,t} e_a l_{j,t}^a - \overline{\lambda}_t \left(\overline{w}_{j,t} e_a l_{j,t}^a \right)^{1-\gamma} \right\} - \frac{\overline{B}_t}{\overline{R}_t}, \\
\Delta_{4,t} &\equiv \sum_{a=0}^{L-2} \sum_{j=n,r} \pi_{j,t-a} \overline{b}_{j,t}^a - \overline{B}_t.
\end{aligned}
$$

Formally, the model provides a mapping from $\mathbb{X} \equiv \{w_{n,t}, w_{r,t}, R_t, \lambda_t\}_{t=0}^{T-1}$ to $\Delta \equiv \{\Delta_{1,t}, \Delta_{2,t}, \Delta_{3,t}, \Delta_{4,t}\}_{t=0}^{T-1}$, and we denote this mapping by $\mathbb{M} : \mathbb{R}^{4T} \to \mathbb{R}^{4T}$. An equilibrium is $\mathbb{X}$ such that $\mathbb{M}(\mathbb{X}) = 0$.

A.3.2 Optimal Mirrleesian policy

Because our calibration procedure yields $\varepsilon = 1$, we specialize this presentation to this case. Assuming that automation is interior (which we verify ex post), we can change variables as in the static model; that is, instead of $\overline{X}_t$ we use the variable τ_t^x, which is such that

$$
\overline{X}_t = \left[\frac{(1-\alpha)A}{1+\tau_t^x} \right] N_{n,t} - \phi_t^{1/\alpha} N_{r,t}.
$$

This change of variables also implies that

$$
\overline{NY}_t = A^{\frac{1}{\alpha}} \frac{(1-\alpha)^{\frac{1-\alpha}{\alpha}}}{(1+\tau_X)^{\frac{1-\alpha}{\alpha}}} \frac{\tau_t^x + \alpha}{1+\tau_t^x} N_{n,t} + \phi_t^{\frac{1}{\alpha}} N_{r,t}
$$

and that relative wages are

$$\frac{\overline{F}_{r,t}}{\overline{F}_{n,t}} = \frac{1-\alpha}{\alpha}\phi_t^{\frac{1}{\alpha}}\left[\frac{1+\tau_t^x}{(1-\alpha)A}\right]^{\frac{1}{\alpha}}.$$

The optimal plan solves the following problem:

$$\max\sum_{t=1}^{\infty}\beta^t\sum_{i=r,n}\left\{\sum_{a=0}^{L-1}\beta^t\pi_{i,t-a}\log\left(\overline{c}_{i,t}^a\right)-\sum_{a=0}^{L_w-1}\beta^t\pi_{i,t-a}v\left(l_{i,t}^a\right)\right\}-\sum_{t=1}^{\infty}\beta^t\int_{-\infty}^{\theta_t^*}h(\theta)\,\theta d\theta+\sum_{t=1}^{\infty}\beta^t L\chi\log\left(\overline{G}_t\right)$$

$$\left[\beta^t\mu_t\right]\quad \sum_{i=r,n}\sum_{a=0}^{L-1}\pi_{i,t-a}c_{i,t}^a+G_t\leq A^{\frac{1}{\alpha}}\left[\frac{1-\alpha}{1+\tau_t^x}\right]^{\frac{1-\alpha}{\alpha}}\frac{\tau_t^x+\alpha}{1+\tau_t^x}N_{n,t}+\phi_t^{\frac{1}{\alpha}}N_{r,t}$$

$$\left[\beta^t\psi_t\right]\quad \theta_t^*=\left\{\sum_{a=0}^{L-1}\beta^t\log\left(\overline{c}_{n,t+a}^a\right)-\sum_{a=0}^{L_w-1}\beta^t v\left(l_{n,t+a}^a\right)\right\}-\left\{\sum_{a=0}^{L-1}\beta^t\log\left(\overline{c}_{r,t+a}^a\right)-\sum_{a=0}^{L_w-1}\beta^t v\left(l_{r,t+a}^a\right)\right\}$$

$$\left[\beta\eta_{n,1-\widetilde{a}}\right]\quad \sum_{a=\widetilde{a}}^{L-1}\beta^{a-\widetilde{a}}\log\left(\overline{c}_{n,1+a-\widetilde{a}}^a\right)-\sum_{a=\widetilde{a}}^{L_w-1}\beta^{a-\widetilde{a}}v\left(l_{n,1+a-\widetilde{a}}^a\right)\geq$$
$$\sum_{a=\widetilde{a}}^{L-1}\beta^{a-\widetilde{a}}\log\left(\overline{c}_{r,1+a-\widetilde{a}}^a\right)-\sum_{a=\widetilde{a}}^{L_w-1}\beta^{a-\widetilde{a}}v\left(\frac{F_{r,1+a-\widetilde{a}}}{F_{n,1+1-\widetilde{a}}}l_{r,1+a-\widetilde{a}}^a\right)$$

$$\left[\beta^t\eta_{n,t}\right]\quad \sum_{a=0}^{L-1}\beta^a\log\left(\overline{c}_{n,t+a}^a\right)-\sum_{a=0}^{L_w-1}\beta^a v\left(l_{n,t+a}^a\right)\geq\sum_{a=0}^{L-1}\beta^a\log\left(\overline{c}_{r,t+a}^a\right)-\sum_{a=0}^{L_w-1}\beta^a v\left(\frac{F_{r,t+a}}{F_{n,t+a}}l_{r,t+a}^a\right)$$

$$\left[\beta\eta_{r,1-\widetilde{a}}\right]\quad \sum_{a=\widetilde{a}}^{L-1}\beta^{a-\widetilde{a}}\log\left(\overline{c}_{r,1+a-\widetilde{a}}^a\right)-\sum_{a=\widetilde{a}}^{L_w-1}\beta^{a-\widetilde{a}}v\left(l_{r,1+a-\widetilde{a}}^a\right)\geq$$
$$\sum_{a=\widetilde{a}}^{L-1}\beta^{a-\widetilde{a}}\log\left(\overline{c}_{n,1+a-\widetilde{a}}^a\right)-\sum_{a=\widetilde{a}}^{L_w-1}\beta^{a-\widetilde{a}}v\left(\frac{F_{n,1+a-\widetilde{a}}}{F_{r,1+1-\widetilde{a}}}l_{n,1+a-\widetilde{a}}^a\right)$$

$$\left[\beta^t\eta_{n,t}\right]\quad \sum_{a=0}^{L-1}\beta^a\log\left(\overline{c}_{r,t+a}^{\tau}\right)-\sum_{a=0}^{L_w-1}\beta^a v\left(l_{r,t+a}^a\right)\geq\sum_{a=0}^{L-1}\beta^a\log\left(\overline{c}_{n,t+a}^a\right)-\sum_{a=0}^{L_w-1}\beta^a v\left(\frac{F_{n,t+a}}{F_{r,t+a}}l_{n,t+a}^a\right),$$

where we write the Lagrange multipliers in parentheses.

Optimality conditions We assume and later verify later that the intensive-margin incentive compatibility of routine workers never binds.

1. First-order conditions with respect to $\bar{c}^a_{j,t}$ for $\widetilde{a} \geq L_w$ (those who are retirees):

$$\bar{c}^a_{j,t} = \frac{1}{\mu_t}.$$

2. First-order conditions with respect to $\bar{c}^a_{j,1+a-\widetilde{a}}$ for $1 \leq \widetilde{a} \leq L_w - 1$, which are (these do not have an extensive-margin IC):

$$\frac{\pi_{r,1-\widetilde{a}} - \eta_{n,1-\widetilde{a}}}{c^a_{r,1+a-\widetilde{a}}} = \mu_{1+a-\widetilde{a}}\pi_{r,1-\widetilde{a}} \Leftrightarrow c_{r,1+a-\widetilde{a}} = \frac{1 - \frac{\eta_{n,1-\widetilde{a}}}{\pi_{r,1-\widetilde{a}}}}{\mu_{1+a-\widetilde{a}}}$$

$$\frac{\pi_{r,1-\widetilde{a}} + \eta_{n,1-\widetilde{a}}}{c^a_{n,1+a-\widetilde{a}}} = \mu_{1+a-\widetilde{a}}\pi_{r,1-\widetilde{a}} \Leftrightarrow c_{n,1+a-\widetilde{a}} = \frac{1 + \frac{\eta_{n,1-\widetilde{a}}}{\pi_{r,1-\widetilde{a}}}}{\mu_{1+a-\widetilde{a}}}.$$

3. First-order conditions with respect to $l^a_{i,1+a-\widetilde{a}}$ for $1 \leq \widetilde{a} \leq L_w - 1$. For routine workers:

$$v'\left(l^a_{r,1+a-\widetilde{a}}\right)\left\{\pi_{r,1-\widetilde{a}} - \eta_{n,1-\widetilde{a}}\left(\frac{\overline{F}_{r,1+a-\widetilde{a}}}{\overline{F}_{n,1+a-\widetilde{a}}}\right)^{1+\nu}\right\} = \pi_{r,1-\widetilde{a}}e_a\mu_{1+a-\widetilde{a}}E_{r,1+a-\widetilde{a}}.$$

For non-routine workers:

$$v'\left(l^a_{n,1+a-\widetilde{a}}\right)\left[\pi_{n,1-\widetilde{a}} + \eta_{n,1-\widetilde{a}}\right] = \pi_{n,1-\widetilde{a}}e_a\mu_{1+a-\widetilde{a}}E_{n,1+a-\widetilde{a}}.$$

4. Consumption $c^a_{i,t+a}$ (i.e., for those born in period $t \geq 1$). For routine workers:

$$\frac{\pi_{r,t} - \psi_t - \eta_{n,t}}{c^a_{r,t+a}} = \mu_{t+a}\pi_{r,t} \Leftrightarrow c^a_{r,t+a} = \frac{1 - \frac{\psi_t + \eta_{n,t}}{\pi_{r,t}}}{\mu_{t+a}}.$$

For non-routine workers:

$$\frac{\pi_{n,t} + \psi_t + \eta_{n,t}}{c^a_{n,t+a}} = \mu_{t+a}\pi_{n,t} \Leftrightarrow c^a_{n,t+a} = \frac{1 + \frac{\psi_t + \eta_{n,t}}{\pi_{n,t}}}{\mu_{t+a}}.$$

5. Government spending G_t:

$$\frac{L\chi}{G_t} = \mu_t \Leftrightarrow G_t = \frac{L\chi}{\mu_t}.$$

6. Labor supply $l^a_{i,t+a}$ (i.e., for those born in period $t \geq 1$). For routine workers:

$$v'\left(l^a_{r,t+a}\right)\left(\pi_{r,t}-\psi_t\right)-\eta_{n,t}v'\left(\frac{\overline{F}_{r,t+a}}{\overline{F}_{n,t+a}}l^a_{r,t+a}\right)\frac{F_{r,t+a}}{F_{n,t+a}}=e_a\pi_{r,t}\mu_{t+a}E_{r,t+a}$$

$$v'\left(l^a_{r,t+a}\right)\left(\pi_{r,t}-\psi_t-\eta_{n,t}\left(\frac{\overline{F}_{r,t+a}}{\overline{F}_{n,t+a}}\right)^{1+\nu}\right)=e_a\pi_{r,t}\mu_{t+a}E_{r,t+a}.$$

For non-routine workers:

$$v'\left(l^a_{n,t+a}\right)\left[\pi_{n,t}+\psi_t+\eta_{n,t}\right]=e_a\pi_{n,t}\mu_{t+a}E_n.$$

7. First-order condition with respect to the tax rate, $\tau_{x,t}$:

$$\sum_{a=0}^{L_w-1}\beta^{t-a}\eta_{n,t-a}\beta^a v'\left(\frac{F_{r,t}}{F_{n,t}}l^a_{r,t}\right)l^a_{r,t}\frac{F_{r,t}}{F_{n,t}}=\mu_t\left[A^{\frac{1}{\alpha}}\left(1-\alpha\right)^{\frac{1}{\alpha}}\frac{\tau_{X,t}}{\left(1+\tau_{X,t}\right)^{\frac{1}{\alpha}}}\right]N_{n,t}.$$

8. First-order condition with respect to θ^*_t:

$$\psi_t=h\left(\theta^*_t\right)\left[\sum_{a=0}^{L-1}\beta^a\mu_{t+a}\left(c^a_{r,t+\tau}-c^a_{n,t+a}\right)+\sum_{a=0}^{L_w-1}\beta^a\mu_{t+a}\left(E_ne_al^a_{n,t+a}-E_{r,t}e_al^a_{r,t+a}\right)\right].$$

9. We then need to add the intensive-margin IC for $a=1,...,L_w-1$ in $t=1$ (i.e., assume that $\eta_{n,1-a}>0$):

$$\sum_{s=0}^{L-1-a}\beta^s\log\left(\overline{c}^{a+s}_{n,1+s}\right)-\sum_{s=0}^{L_w-1-a}\beta^s v\left(l^s_{n,1+s}\right)=\sum_{s=0}^{L-1-a}\beta^s\log\left(\overline{c}^{a+s}_{r,1+s}\right)-\sum_{s=0}^{L_w-1-a}\beta^s v\left(\frac{\overline{F}_{r,1+s}}{\overline{F}_{n,1+s}}l^{a+s}_{r,1+s}\right).$$

This condition can be written as

$$\sum_{s=0}^{L-1-a}\beta^s\log\left\{\frac{\overline{c}^{a+s}_{n,1+s}}{\overline{c}^{a+s}_{r,1+s}}\right\}=\sum_{s=0}^{L_w-1-a}\beta^s\left\{v\left(l^s_{n,1+s}\right)-v\left(\frac{\overline{F}_{r,1+s}}{\overline{F}_{n,1+s}}l^{a+s}_{r,1+s}\right)\right\}.$$

10. This optimization requires $\eta_{n,t}\geq 0$ for $t\geq 1$, so we need to add the extra-conditions (here, we write these equations only for $t\geq 1$ as those are the ones for which the extensive-margin IC may mean that the intensive-margin IC does not bind)

$$\sum_{a=0}^{L-1}\beta^a\log\left(\overline{c}^a_{n,t+a}\right)-\sum_{a=0}^{L_w-1}\beta^a v\left(l^a_{n,t+a}\right)\geq\sum_{a=0}^{L-1}\beta^a\log\left(\overline{c}^a_{r,t+a}\right)-\sum_{a=0}^{L_w-1}\beta^a v\left(\frac{\overline{F}_{r,t+a}}{\overline{F}_{n,t+a}}l^a_{r,t+a}\right)$$

$$\eta_{n,t}\geq 0$$

$$\eta_{n,t}\left[\sum_{a=0}^{L-1}\beta^a\log\left(\overline{c}^a_{n,t+a}\right)-\sum_{a=0}^{L_w-1}\beta^a v\left(l^a_{n,t+a}\right)-\sum_{a=0}^{L-1}\beta^a\log\left(\overline{c}^a_{r,t+a}\right)+\sum_{a=0}^{L_w-1}\beta^a v\left(\frac{\overline{F}_{r,t+a}}{\overline{F}_{n,t+a}}l^a_{r,t+a}\right)\right]=0.$$

11. Intensive-margin IC:

$$\theta_t^* = \left\{ \sum_{a=0}^{L-1} \beta^t \log\left(\bar{c}_{n,t+a}^a\right) - \sum_{a=0}^{L_w-1} \beta^t v\left(l_{n,t+a}^a\right) \right\} - \left\{ \sum_{a=0}^{L-1} \beta^t \log\left(\bar{c}_{r,t+a}^a\right) - \sum_{a=0}^{L_w-1} \beta^t v\left(l_{r,t+a}^a\right) \right\}$$

$$\pi_{n,t} = H(\theta_t^*).$$

12. Resource constraint:

$$\underbrace{\sum_{i=r,n} \sum_{a=0}^{L-1} \pi_{i,t-a} c_{i,t}^a}_{\frac{L(1+\chi)}{\mu_t}} + G_t = \frac{\alpha A^{\frac{1}{\alpha}} (1-\alpha)^{\frac{1-\alpha}{\alpha}}}{(1+a_{X,t})^{\frac{1-\alpha}{\alpha}}} \frac{a_{X,t} + \alpha}{\alpha (1+a_{X,t})} N_{n,t} + \phi_t^{\frac{1}{\alpha}} N_{r,t}.$$

We start by solving these equations to obtain the steady state, as described in appendix A.2.5. For our calibration, the intensive-margin constraint is not binding in the steady state.

Next, we use a root finding algorithm to solve for the transition to this steady state. We assume that convergence to steady state occurs after $T = 50$ periods. To simplify the computational process, we proceed in steps. We first search for a solution that disregards all intensive-margin incentive constraints (i.e., such that $\eta_{n,t} = 0$ for all $t = 2-L, 3-L, ...$). We check that, in this solution, the intensive-margin constraints of the old generations at time 1 are violated. We then use this solution as an initial guess for a computational algorithm that includes these constraints for every old generation in period 1 (i.e., $\eta_{n,t} > 0$ for $t = 1-L, ..., 0$). Next, we sequentially check whether the intensive-margin constraints are violated for the generations at time $t \geq 1$ and add these constraints to the problem if they are. In our calibration, these constraints do not bind.

A.4 Robustness exercises

A.4.1 Mirrleesian optimal taxation from 1988

We solve for the optimal tax system starting in 1988. Figures 9 and 10 display our results. As in the optimal plan discussed in Section 5, we see that robots are taxed only while the initial old generations of routine workers are in the labor force. Optimal robot taxes in this plan are 9.1 percent between 1988 and 1997, 2.9 percent between 1998 and 2007, and 0.7 percent between 2008 and 2017.

Figure 9: Mirrleesian optimal taxation A

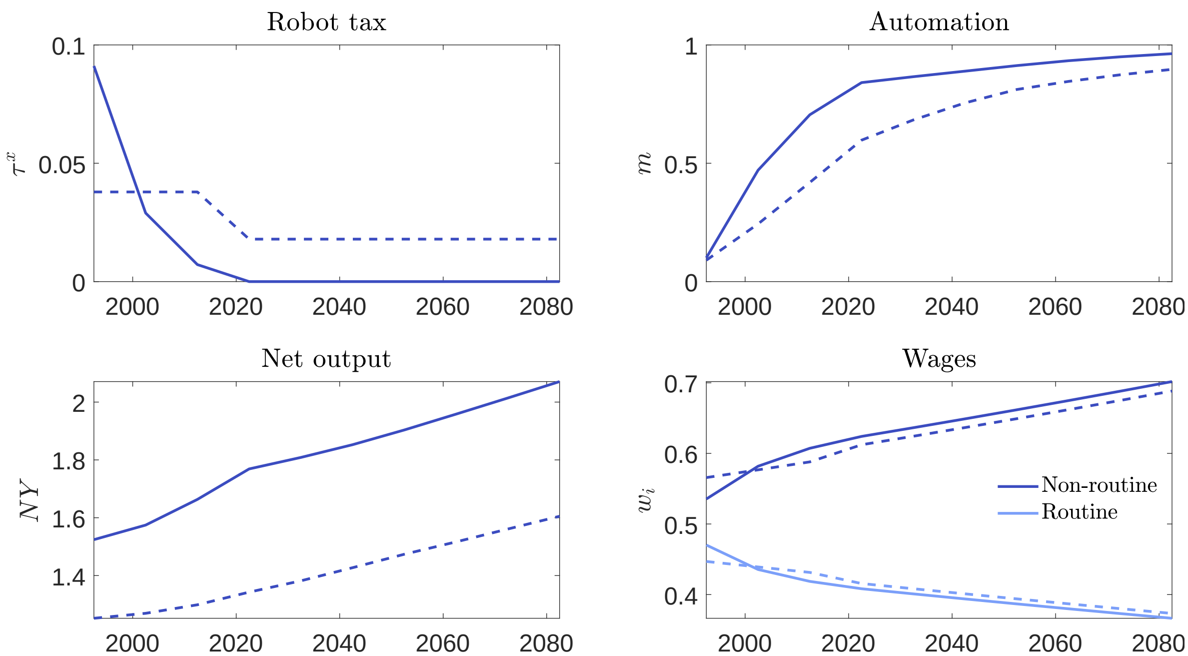

Notes: The solid lines show the allocations for a Mirrleesian optimal plan that starts to be implemented in 1988. The dashed lines correspond to the status quo equilibrium. The first and second panels show the evolution of robot taxes, τ_t^x, and automation, m_t, respectively. The third and fourth panels plot the equilibrium levels of net output and wages for the two worker types, respectively.

A.4.2 Robustness to parameters

Table 6 summarizes key properties of the optimal plan for different parameter choices.

Figure 10: Mirrleesian optimal taxation B

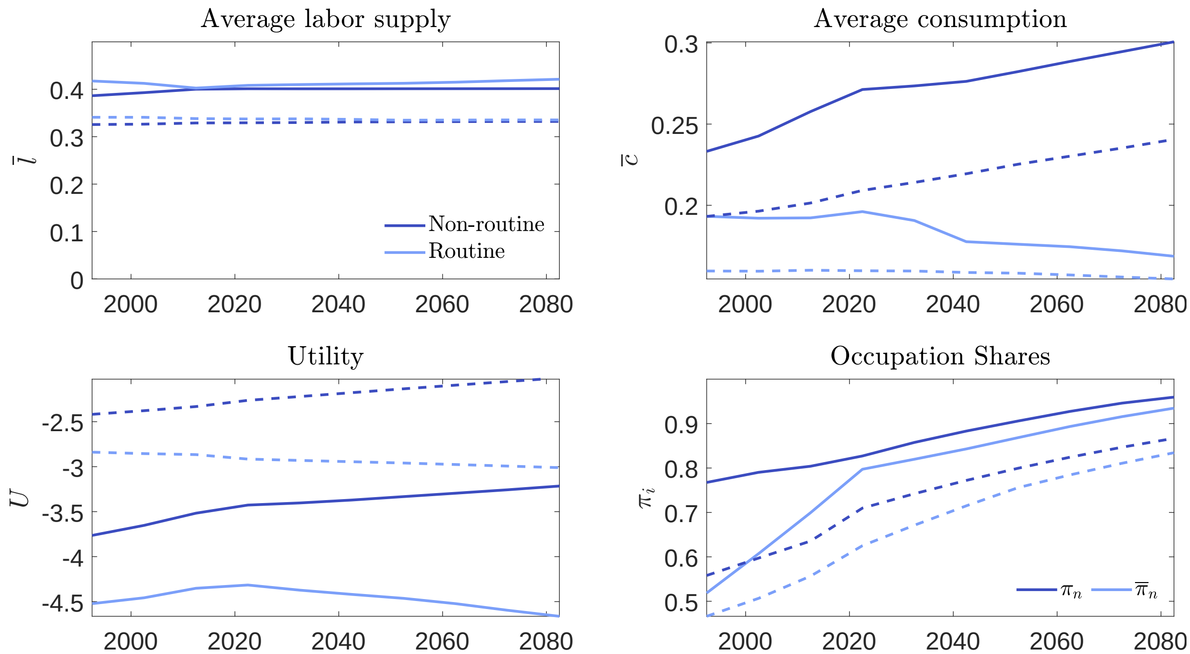

Notes: The solid lines show the allocations for a Mirrleesian optimal plan that starts to be implemented in 1988. The dashed lines correspond to the status quo equilibrium. The first and second panels show the equilibrium levels of the average across agents of the labor supply and consumption, respectively. The third panel plots the equilibrium levels of utility for both agents. Finally, the fourth panel shows the share of newborns who choose non-routine skills, $\pi_{n,t}$, and the share of non-routine workers in the workforce, $\overline{\pi}_{n,t}$.

Table 6: Robustness

			2018–2027	2028–2037	2038–2047
Baseline	Robot tax (%)	τ_t^x	5.1	2.2	0.6
	Newborn non-routine	$\pi_{n,t}$	0.88	0.88	0.88
	Wage premium	$w_{n,t}/w_{r,t}$	1.38	1.52	1.63
High cost level $\widetilde{\phi} = 0.44$	Robot tax (%)	τ_t^x	5.7	2.4	0.7
	Newborn non-routine	$\pi_{n,t}$	0.84	0.85	0.86
	Wage premium	$w_{n,t}/w_{r,t}$	1.31	1.45	1.56
Low cost level $\widetilde{\phi} = 0.42$	Robot tax (%)	τ_t^x	4.5	2.0	0.6
	Newborn non-routine	$\pi_{n,t}$	0.92	0.90	0.90
	Wage premium	$w_{n,t}/w_{r,t}$	1.45	1.58	1.69
Low skill elasticity $\sigma = 1.23$	Robot tax (%)	τ_t^x	5.9	2.9	0.9
	Newborn non-routine	$\pi_{n,t}$	0.54	0.62	0.63
	Wage premium	$w_{n,t}/w_{r,t}$	1.36	1.50	1.62
Faster tech. progress $g_\phi = 0.036$	Robot tax (%)	τ_t^x	5.2	1.8	0.4
	Newborn non-routine	$\pi_{n,t}$	0.91	0.93	0.94
	Wage premium	$w_{n,t}/w_{r,t}$	1.37	1.59	1.76
Slower tech. progress $g_\phi = 0.009$	Robot tax (%)	τ_t^x	5.1	2.4	0.8
	Newborn non-routine	$\pi_{n,t}$	0.86	0.85	0.84
	Wage premium	$w_{n,t}/w_{r,t}$	1.38	1.48	1.56

Note: This table summarizes key properties of the Mirrleesian optimal plan for different parameter configurations. The baseline parameters are reported in Table 1. For comparison purposes, we start by reporting the baseline results. For the first two robustness exercises, we vary the rate of technical progress g_ϕ fixing the level of ϕ_t in 2018. For the next two exercises, we consider changes in the level $\widetilde{\phi}$. In the final two exercises, we consider two economies with a higher variance of θ.