

NBER WORKING PAPER SERIES

SHOULD ROBOTS BE TAXED?

Joao Guerreiro
Sergio Rebelo
Pedro Teles

Working Paper 23806
<http://www.nber.org/papers/w23806>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2017, Revised January 2019

We thank Gadi Barlevy, V.V. Chari, Bas Jacobs, and Nir Jaimovich for their comments. Teles thanks the support of 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 N° 649396. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Joao Guerreiro, Sergio Rebelo, and Pedro Teles. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Should Robots be Taxed?

Joao Guerreiro, Sergio Rebelo, and Pedro Teles

NBER Working Paper No. 23806

September 2017, Revised January 2019

JEL No. H21,O33

ABSTRACT

We use a model of automation to show that with the current U.S. tax system, a fall in automation costs could lead to a massive rise in income inequality. This inequality can be reduced by making the current income-tax system more progressive and by taxing robots. But this solution involves a substantial efficiency loss. A Mirrleesian optimal income tax can reduce inequality at a smaller efficiency cost. An alternative approach is to amend the current tax system to include a lump-sum rebate. With the rebate in place, it is optimal to tax robots as long as there is partial automation.

Joao Guerreiro
Department of Economics
Northwestern University
Evanston, IL 60208
Evanston, IL 60201
United States
jguerreiro@u.northwestern.edu

Pedro Teles
Banco de Portugal
R. Francisco Ribeiro 2
1150 Lisboa
Portugal
and Univ Catolica Portuguesa and CEPR
pteles@ucp.pt

Sergio Rebelo
Northwestern University
Kellogg School of Management
Department of Finance
Leverone Hall
Evanston, IL 60208-2001
and CEPR
and also NBER
s-rebelo@northwestern.edu

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 him to write a novel called “Player Piano.” It describes a world in which school children take a test at an early age that determines their fate. Those who pass, become engineers and design robots used in production. Those who fail, have no jobs and live from government transfers. 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 wide range of economic activities, Bill Gates re-ignited this debate by proposing that robots should be taxed. Policies that address the impact of automation on the labor market have been discussed in the European Parliament and have been implemented in South Korea.

In this paper, we use a simple model of automation to compare the equilibrium that emerges under the current U.S. tax system (which we call the status quo), the first-best solution to a planner’s problem without information constraints, and the second-best solutions associated with different configurations of the tax system.

Our model has two types of workers which we call routine and non-routine. Routine workers perform tasks that can be automated by using intermediate inputs that we refer to as robots.¹ We find that robot taxes are optimal as long as there is partial automation. These taxes increase the wages of routine workers, and decrease those of non-routine workers, giving the government an additional instrument to reduce income inequality. With full automation, it is not optimal to tax robots. Routine workers no

¹See Acemoglu and Autor (2011) and Cortes, Jaimovich and Siu (2017) for a discussion of the impact of automation on the labor market for routine workers.

longer work, so taxing robots distorts production decisions without reducing income inequality.²

We model the current U.S. tax system using the after-tax income function proposed by Feldstein (1969), Persson (1983), and Benabou (2000) and estimated by Heathcote, Storesletten and Violante (2017). With this tax system, as the cost of automation falls, the wages of non-routine workers rise while the wages of routine workers fall to make them competitive with robot use. The result is a large rise in income inequality and a substantial decline in the welfare of routine workers.

The level of social welfare obtained in the status quo is much worse than that achieved in the first-best solution to an utilitarian social planner problem without information constraints. But this first-best solution cannot be implemented when the government does not observe the worker type. The reason is that the two types of workers receive the same level of consumption but non-routine workers supply more labor than routine workers. As a result, non-routine workers have an incentive to act as routine workers and receive their bundle of consumption and hours worked.

To circumvent this problem, 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 workers' labor input. The government can observe total income and consumption of the two types of workers, as well as the use of robots by firms. We assume that taxes on robots are linear for the reasons emphasized in Guesnerie (1995): non-linear taxes on intermediate inputs are difficult to implement in practice because they create arbitrage opportunities. A Mirrleesian optimal tax system can improve welfare relative to the status quo. In fact, it can yield a level of welfare that is close to that of the first-best allocation.

²These results show that the reason why it can be optimal to tax robots in our model differs from the rationale used by Bill Gates to motivate robot taxation. Gates argued that robots should be taxed to replace the tax revenue that the government collected from routine workers before their jobs were automated. In our model, when there is full automation the government collects no tax revenue from routine workers yet it is optimal not to tax robots.

We also study the optimal policy when the tax schedule is constrained to take a simple, exogenous form. We consider the income tax schedule proposed by Heathcote, Storesletten and Violante (2017) and linear robot taxes. We compute the parameters of the income tax function and the robot tax rate that maximize social welfare. We find that income inequality can be reduced by raising marginal tax rates and taxing robots. Tax rates on robot use can be as high as 30 percent. Routine workers supply a constant number of hours over time even though their wages fall. This solution yields poor outcomes in terms of efficiency and distribution.

We consider a modification of the Heathcote, Storesletten and Violante (2017) tax schedule that allows for lump-sum transfers that ensure that all workers receive a minimum income. We find that this modification improves both efficiency and distribution relative to a tax system without transfers. Hours worked by routine and non-routine workers diverge over time. Full automation occurs in finite time, so hours worked by routine workers fall to zero. Once full automation occurs, routine workers pay no income taxes and the tax system can be designed so that the labor-supply decisions of non-routine workers are not distorted. The economy with full automation resembles the world of “Player Piano.” Only non-routine workers have jobs. Routine workers live off government transfers and, despite losing their jobs, are better off than in the status quo. The fact that full automation occurs in finite time reflects the rudimentary nature of the tax system available to the government. When the government has access to a more flexible non-linear tax schedule, as in the Mirrleesian solution, full automation occurs only asymptotically.

One might expect optimal robot taxation to follow from well-known principles of optimal taxation in the public finance literature. We know from the intermediate-goods theorem of Diamond and Mirrlees (1971) that it is not optimal to distort production decisions by taxing intermediate goods. Since robots are in essence an intermediate good, taxing them should not be optimal.

The intermediate-good theorem relies on the assumption that “net trades” of differ-

ent goods can be taxed at different rates. In our context, this assumption means that the government can use different tax schedules for routine and non-routine workers. We study two environments where there are limits to the government’s ability to tax different workers at different rates, Mirrlees (1971)-type information constraints and a simple exogenous tax system common to both types of workers. We find that it is optimal to tax robots in both environments.

This finding seems to contradict the key result in Atkinson and Stiglitz (1976). These authors show that when the income tax system is non-linear it is not optimal to distort production decisions by taxing intermediate goods. But, as stressed by Naito (1999) and Jacobs (2015), Atkinson and Stiglitz (1976)’s result depends critically on the assumption that workers with different productivities are perfect substitutes in production. This assumption does not hold in our model. Taxing robots can be optimal because it affects relative productivities, loosening the incentive constraint of non-routine workers.

We extend our model to allow workers to switch their occupations by paying a cost. In the first-best solution, workers who have a low cost of becoming non-routine workers do so. Those with a high cost become routine workers. In the Mirrlees solution to the model with occupational choice, it is optimal to use robot taxes to loosen the incentive constraint of non-routine workers. The planner can use the income tax schedule to redistribute income or to induce more agents to become non-routine workers. When the cost of becoming non-routine are high (low), the planner resorts more (less) to using the income tax schedule to redistribute income.

We generalize our static model to a dynamic setting in which robots are an investment good. The properties of the Mirleesian solution of the dynamic model are similar to those of the static model. It is optimal to tax robots to loosen the incentive constraint of non-routine workers. The levels of taxation are similar to those of the static model. The tax rate on robots converges to zero as the degree of automation converges to one.

The paper is organized as follows. In Section 2, we describe our static model of

automation. Subsection 2.1 describes the status-quo equilibrium, i.e. the equilibrium under the current U.S. income tax system and no robot taxes. Subsection 2.2 describes the first-best solution to the problem of an utilitarian planner. In subsection 2.3, we analyze a Mirrleesian second-best solution to the planner’s problem. In subsection 2.4, we study numerically the optimal tax system that emerges when income taxes are constrained to take the functional form proposed by Heathcote, Storesletten and Violante (2017) both with and without lump-sum rebates. In subsection 2.5, we compare the implications of different policies for social welfare and for the utility of different agents. Subsection 2.6 discusses the model with endogenous occupation choice. In Section 3, we analyze a dynamic model of automation. Section 4 relates our findings to classical results on production efficiency and capital taxation in the public finance literature. Section 5 concludes. To streamline the main text, we relegate the more technical proofs to the appendix.

2 A simple model

We first discuss a simple model of automation that allows us to address the optimal tax policy questions posed in the introduction. The model has two types of households who draw utility from consumption of private and public goods and disutility from labor. One household type supplies routine labor and the other non-routine labor. The consumption good is produced with non-routine labor, routine labor, and robots. Robots and routine labor are used in a continuum of tasks. They are both perfect substitutes in performing these tasks.³

Households There is a continuum of unit measure of households. A mass π_n of households is composed of non-routine workers while π_r households are composed of

³See Autor, Levy and Murmane (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.

routine workers. The index $j = n, r$, denotes the non-routine and routine labor type, respectively.

An household of type j derives utility from consumption, c_j , and from the provision of a public good, G . The household also derives disutility from the hours of labor it supplies, l_j . Households have a unit of time per period, so $l_j \leq 1$. The household's utility function is given by

$$U_j = u(c_j, l_j) + v(G). \quad (1)$$

Denote by $u_x = \partial u(c, l) / \partial x$ where $x = c, l$ and $u_{xy} = \partial^2 u(c, l) / \partial x \partial y$. We assume that $u_c > 0$, $u_l < 0$, $u_{cc}, u_{ll} < 0$ and that consumption and leisure are normal goods: $u_{lc}/u_l - u_{cc}/u_c \geq 0$, and $u_{ll}/u_l - u_{cl}/u_c \geq 0$, where one of these conditions is a strict inequality. Furthermore, we assume that utility satisfies the single-crossing property, which is equivalent to assuming that $u_{ll}l/u_l + 1 - u_{cl}l/u_c > 0$. Finally, we assume that $v'(G) > 0$, $v''(G) < 0$ and that $u(c, l)$ satisfies standard Inada conditions.

Household j chooses c_j and l_j 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 the household type j and $T(\cdot)$ denotes the income tax schedule.

Robot producers Final good producers can use robots in tasks $i \in [0, 1]$. The cost of producing a robot is the same across tasks and is equal to ϕ units of output. Robots are produced by competitive firms. A representative firm producing robots chooses x_i to maximize profits $p_i x_i - \phi x_i$. It follows that in equilibrium $p_i = \phi$ and profits are zero.

Final good producers The representative producer of final goods hires non-routine labor (N_n), routine labor (n_i) for each task i , and buys intermediate goods (x_i) which we refer to as robots, also for each task i . There is a continuum of tasks that can be performed by either routine labor or robots. We denote by m the fraction of these tasks that are automated, i.e. performed by robots. For convenience, we assume that

automated tasks are those in the interval $[0, m]$.⁴

The production function is given by

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

The problem of the firm 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 ((2)). The variable τ_x is an ad-valorem tax rate on intermediate goods.

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

$$w_n = \frac{\alpha Y}{N_n}, \quad (3)$$

$$(1 + \tau_x) \phi = (1 - \alpha) Y \left(\int_0^m x_s^{\frac{\rho-1}{\rho}} ds + \int_m^1 n_s^{\frac{\rho-1}{\rho}} ds \right)^{-1} x_i^{-\frac{1}{\rho}}, \text{ for } i \in [0, m], \quad (4)$$

$$w_r = (1 - \alpha) Y \left(\int_0^m x_s^{\frac{\rho-1}{\rho}} ds + \int_m^1 n_s^{\frac{\rho-1}{\rho}} ds \right)^{-1} n_i^{-\frac{1}{\rho}}, \text{ for } i \in (m, 1]. \quad (5)$$

It follows that it is optimal to use the same level of routine labor, n_i , in the $1 - m$ tasks that have not been automated and that the optimal use of robots is also the same in the m automated tasks.

The optimal level of automation is $m = 0$ if $w_r < (1 + \tau_x)p_x$. The firm chooses to fully automate ($m = 1$) and to employ no routine workers ($n_i = 0$) 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 this case, equations (4) and (5) imply that the levels of routine labor and robots are the same across tasks,

$$m x_i = X, \text{ for } i \in [0, m], \text{ and } (1 - m) n_i = N_r, \text{ for } i \in (m, 1], \quad (6)$$

where N_r denotes total routine hours and X denotes total robots. Using the fact that $x_i = n_j$, with interior automation we obtain $m = X/(N_r + X)$, and we can write the

⁴Since tasks are symmetric, there is no loss of generality associated with this assumption.

production function as $Y = A(X + N_r)^{1-\alpha} N_n^\alpha$. Since the technology has constant returns to scale, profits are zero.

Government The government chooses taxes and the optimal level of government spending, subject to the budget constraint

$$G \leq \pi_r T(w_r l_r) + \pi_n T(w_n l_n) + \tau_x p_x \int_0^m x_i di. \quad (7)$$

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 $\{T(\cdot), \tau_x\}$ such that: (i) given prices and taxes, allocations solve the households' problem; (ii) given prices and taxes, allocations solve the firms' problem; (iii) the government budget constraint is satisfied; and (iv) markets clear.

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

$$\begin{aligned} (1 - m)n_i &= N_r = \pi_r l_r, \\ N_n &= \pi_n l_n. \end{aligned} \quad (8)$$

The market-clearing condition for robots is redundant by Walras' law. The market-clearing condition for the output market is

$$\pi_r c_r + \pi_n c_n + G \leq Y - \phi \int_0^m x_i di. \quad (9)$$

In an equilibrium with automation ($m \in (0, 1)$) in which $w_r = (1 + \tau_x)\phi$, we also have

$$x_i = \frac{X}{m} = \frac{N_r}{1 - m}, \text{ for } i \in [0, m].$$

With interior automation, given aggregate labor supplies $\pi_r l_r$ and $\pi_n l_n$, the equilibrium level of automation satisfies

$$m = 1 - \left[\frac{(1 + \tau_x)\phi}{(1 - \alpha)A} \right]^{1/\alpha} \frac{\pi_r l_r}{\pi_n l_n}, \quad (10)$$

and the wage rates of both non-routine and routine labor are independent of preferences,

$$w_n = \alpha \frac{A^{1/\alpha}(1-\alpha)^{\frac{1-\alpha}{\alpha}}}{[(1+\tau_x)\phi]^{\frac{1-\alpha}{\alpha}}}, \quad (11)$$

$$w_r = (1+\tau_x)\phi. \quad (12)$$

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, and so is the wage of the non-routine worker. An increase in τ_x raises the wage of routine workers and lowers the wage rate of non-routine agents. It is also useful to note that in any equilibrium the income shares of total production are given by

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

An increase in automation reduces the income share of routine workers and does not change the share of non-routine workers. In this sense, an increase in automation leads to an increase in pre-tax income inequality.

2.1 The status-quo equilibrium

In this section, we describe the status-quo equilibrium, i.e. the equilibrium under the current U.S. income tax system and no taxes on robot use ($\tau_x = 0$). We model the U.S. income tax system using the functional form for after-tax income proposed by Feldstein (1969), Persson (1983), and Benabou (2000) and estimated by Heathcote, Storesletten and Violante (2017). In this specification, the income tax paid by household j is given by⁵

$$T(w_j l_j) = w_j l_j - \lambda(w_j l_j)^{1-\gamma}, \quad (13)$$

where $\gamma < 1$. Using PSID data, Heathcote, Storesletten and Violante (2017) estimate that $\gamma = 0.181$, which means that income taxes are close to linear. They find that their

⁵Income in Heathcote, Storesletten and Violante (2017) includes other sources of income, other than labor earnings.

specification fits the data with an R^2 of 0.91. 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 (negative), the average tax rate rises (falls) with income, so the tax system is progressive (regressive).

We assume in all our numerical work that the utility function takes the form:

$$u(c_j, l_j) + v(G) = \log(c_j) - \zeta \frac{l_j^{1+\nu}}{1+\nu} + \chi \log(G). \quad (14)$$

These preferences, which have been used by Ales, Kurnaz and Sleet (2015) and Heathcote, Storesletten and Violante (2017), have two desirable properties. First, they are consistent with balanced growth. Second, they are consistent with the empirical evidence reviewed in Chetty (2006).

For these preferences and the status-quo tax specification, both households choose to work the same number of hours, $l_j = [(1 - \gamma) / \zeta]^{1/(1+\nu)}$, which only depend on the preference parameters, ζ and ν , and the progressivity parameter, γ .

Model calibration We set $\zeta = 10.63$, so that in the status-quo equilibrium agents choose to work 1/3 of their time endowment. We set $\nu = 4/3$, so that the Frisch elasticity is equal to 0.75, which is consistent with the estimates discussed by Chetty, Guren, Manoli, and Weber (2011).

Following Heathcote et al. (2017), we choose $\chi = 0.233$ so that the optimal ratio of government to output is 18.9 percent, the same weight observed in the U.S. economy.⁶ The tax on robots is zero in the status quo ($\tau_x = 0$). We assume that the level of progressivity of the tax system is $\gamma = 0.181$, the value estimated by Heathcote, et al. (2017). We adjust λ to satisfy the government budget constraint.⁷ On the production

⁶When utility takes the form (14), the optimal ratio of government spending to output is the same for all the tax systems we consider.

⁷An alternative approach would have been to keep the tax schedule constant and adjust the level of government spending to balance the government budget. However, this approach would make it more difficult to compare the solutions for the different tax systems.

side, we normalize A to one, choose $\alpha = 0.53$ and $\pi_r = 0.55$. These choices are consistent with the share of labor income received by non-routine workers and the fraction of workers that are routine estimated by Chen (2016).⁸ These values of A , α , π_r and parameter choices are used in all our numerical experiments.

In our quantitative analysis, we consider a sequence of static economies where the cost of a robot falls geometrically over time, $\phi_t = \phi_0 e^{-g_\phi t}$. We set g_ϕ , the rate of decline in the price of robots, equal to 0.0083. This value allows our model to match the decline in task content estimated by Acemoglu and Restrepo (2018b) for the period 2000-2008. We choose this period to abstract from the financial crises and focus on the period where automation takes off in the U.S. We set $\phi_0 = 0.4226$ which is the lowest value of ϕ consistent with no automation in the status-quo equilibrium (see equation (10)). We assume that time zero corresponds to year 2000 and label our figures accordingly.

Figure 1 describes the effect of changes in the cost of automation. As time goes by, ϕ falls causing the wage of routine workers to fall and that of non-routine workers to rise. Since the utility function is logarithmic and wages are the only income source, hours worked remain constant for both routine and non-routine workers. This property reflects the offsetting nature of income and substitution effects. Given that as ϕ falls, wages of routine workers fall and their hours worked remain constant, their income, consumption, and utility fall. In contrast, non-routine workers benefit from rising income, consumption and utility.

As ϕ falls, the parameter that controls the level of taxation, λ , rises, which implies a decline in the overall level of taxation. This decline reflects the increasing share of tax revenue paid by non-routine workers pay and the fact that, as ϕ falls, their income rises faster than output.

In sum, our analysis suggests that the current U.S. tax system will lead to massive

⁸The equilibrium is independent of the value of ρ , the parameter that controls the elasticity of substitution between different tasks. The reason for this result is that all the factors (non-routine workers and/or robots) used in equilibrium to perform these tasks have the same marginal cost.

income and welfare inequality in response to a fall in the costs of automation.

2.2 The first-best allocation

The first-best allocation for this economy maximizes the weighted average of utilities, subject only to resource feasibility, without taking into account implementability constraints arising from information or other constraints.

The weighted average of utilities assigns positive weights ω_r and ω_n to routine and non-routine agents, respectively. These weights are normalized so that $\pi_r\omega_r + \pi_n\omega_n = 1$. The planning problem is to choose $\{c_r, l_r, c_n, l_n, G, m, \{x_i, n_i\}\}$ to maximize social welfare,⁹

$$\pi_r\omega_r [u(c_r, l_r) + v(G)] + \pi_n\omega_n [u(c_n, l_n) + v(G)]. \quad (15)$$

The first-best allocation features production efficiency, because the robots' marginal productivity equals their marginal cost, ϕ . We focus on the case in which the planner values redistribution to routine workers (the relatively poorer agent). If $\omega_r \geq 1$, then the planner gives more consumption to the routine worker, $c_r \geq c_n$, and requires the more productive non-routine worker to work more. This result implies that in the first-best the utility of routine workers is higher than that of non-routine workers. Clearly, the first-best solution cannot be implemented if the planner cannot discriminate between household types. In this solution, non-routine households would have an incentive to act as routine to benefit from a more generous consumption bundle.

In all our quantitative exercises, we assume that $\omega_r = 1$. Figure 2 illustrates the properties of the first-best solution. In panel A, we see that full automation occurs only asymptotically. However, 21 percent of the tasks are automated in the first period, 50 percent by 2017, and 75 percent by 2042. The real wage rate for both types of workers

⁹One interpretation of the social welfare function is as follows. Workers are identical ex-ante because they do not know whether their skills can be automated or not, i.e. whether they will be routine or non-routine workers. The planner maximizes the worker's ex-ante expected utility.

are the same as in the status-quo equilibrium.¹⁰ The consumption and utility of both types of worker rise as ϕ falls. Figure 2 also shows that implementing the first-best solution requires large transfers from non-routine to routine workers.

2.3 Mirrleesian optimal taxation

In this section, we characterize the optimal non-linear income tax when the planner observes a worker's total income but does not observe the worker's type or labor supply, as in the canonical Mirrlees (1971) problem. For the reasons emphasized in Guesnerie (1995), we assume that robot taxes are linear.

In the Mirrlees (1971) model, the productivities of different agents are exogenous. In our model, these productivities are endogenous and depend on τ_x . This property is central to the question we are interested in studying: is it optimal to distort production decisions by taxing the use of robots to redistribute income from non-routine to routine workers to increase social welfare?

In the analytical description of the optimal policy, we focus attention on plans with interior automation, $m > 0$.¹¹ 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 ($w_n \geq w_r$) in an equilibrium with automation (see equations (11) and (12)).

The Mirrleesian planning problem is to choose the allocations $\{c_j, l_j\}_{j=r,n}$, G , and the robot tax τ_x to maximize social welfare, defined in equation (15), subject to the resource constraint

$$\pi_r c_r + \pi_n c_n + G \leq \pi_n w_n l_n \frac{\tau_x + \alpha}{\alpha(1 + \tau_x)} + \frac{\pi_r w_r l_r}{1 + \tau_x}. \quad (16)$$

¹⁰The reason for this property is as follows. Equations (11) and (12) imply that wages depend on technological parameters (α and A), the cost of automation, and the value of τ_x . Since $\tau_x = 0$ in the status quo and there is production efficiency in the first-best allocation, the wages are the same in both allocations.

¹¹We do not discuss the case where $m = 0$ because in this case the results in Stiglitz (1982) apply to our model.

and two incentive constraints (IC)

$$u(c_n, l_n) \geq u\left(c_r, \frac{w_r}{w_n} l_r\right), \quad (17)$$

$$u(c_r, l_r) \geq u\left(c_n, \frac{w_n}{w_r} l_n\right), \quad (18)$$

The wages of the two types of workers are given by equations (11) and (12).

Any competitive equilibrium satisfies equations (16), (17), and (18). In addition, any allocation that satisfies these three equations can be decentralized as a competitive equilibrium.

Household 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 IC constraints, (17), and (18), so these conditions are necessary.

We show in the Appendix that equation (16) is necessary by combining the first-order conditions to the firms' problems with the resource constraint, (9). In addition, we show that conditions (16), (17), and (18), are also sufficient. To see that equations (17) and (18) summarize the household 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 $w_j l_j$, $j = r, n$. These results are summarized in the following proposition.

Lemma 1. *Equations (16), (18) and (17) characterize the set of implementable allocations. These conditions are necessary and sufficient for a competitive equilibrium.*

Since the government can choose an arbitrary tax function, it is only bound by the incentive constraints which 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.

The tax on intermediate goods provides the government with an additional instrument relative to the Mirrlees (1971) setting. The planner can use this instrument to affect the income of the two types of workers but its use distorts production.

To bring the analysis closer to a canonical Mirrleesian approach, we maximize the planner's objective in two steps. First, we set τ_x to a given level and solve for the optimal allocations. Second, we find the optimal level of τ_x . We define $W(\tau_x)$ as the maximum level of social welfare, (15), subject to the incentive constraints, (17) and (18), and the resource constraint, (16) for a given value of τ_x .

An optimal choice of τ_x requires that $W'(\tau_x) = 0$. We characterize optimal allocations in which the incentive constraint of the non-routine worker binds, and the incentive constraint of the routine worker is slack. This pattern holds in all our numerical exercises.

The expression for net output in the right-hand side of equation (16) can be written as

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

The term $(\tau_x + \alpha) / [\alpha(1 + \tau_x)^{1/\alpha}]$ 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 tax on robots. With automation is incomplete, the planner is willing to pay a resource cost, in terms of this production inefficiency, in order to loosen the incentive constraints that are also functions of the robot tax.

Proposition 1. *Suppose the optimal allocation is such that the non-routine workers' incentive constraint binds and the incentive constraint for routine workers does not bind.*

Then, if automation is incomplete ($m < 1$ and $l_r > 0$), robot taxes are strictly positive ($\tau_x > 0$). The optimal tax on robots satisfies

$$\frac{\tau_x}{1 + \tau_x} = \frac{\alpha}{1 - \alpha} \frac{\pi_r \phi l_r}{\pi_n w_n l_n} \left[1 + \frac{\omega_r u_l(c_r, l_r)}{\mu \phi} \right], \quad (19)$$

where μ denotes the multiplier on the resource constraint (16).

This proposition is proved in the appendix. To see the intuition for this result, suppose that $\tau_x < 0$. A marginal increase in τ_x has two benefits. First, it strictly increases output and hence the amount of goods available for consumption. Second, it reduces the relative wage w_n/w_r and makes the non-routine worker less inclined to mimic the routine workers. This property can be easily seen from the incentive constraint of the non-routine worker: $u(c_n, l_n) \geq u\left(c_r, \frac{w_n}{w_r} l_n\right)$.

Consider instead $\tau_x = 0$. Since a zero tax on robots maximizes output, for fixed labor supplies, 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 informational restriction. 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 ($m = 1, l_r = 0$) there is no informational gains from taxing robots. Since the robot tax distorts production and does not help loosen the incentive constraint of the non-routine agent, the optimal value of τ_x is zero. We prove this result in the appendix.

We now turn to the study of the optimal wedges. The optimality conditions imply the following marginal rates of substitution:

$$\begin{aligned} \frac{u_l(c_n, l_n)}{u_c(c_n, l_n)} &= w_n \frac{\tau_x + \alpha}{\alpha(1 + \tau_x)}, \\ \frac{u_l(c_r, l_r) l_r}{u_c(c_r, l_r)} &= \frac{\omega_r - \eta_n \frac{u_c(c_r, w_r l_r / w_n)}{u_c(c_r, l_r)}}{\omega_r - \eta_n \frac{u_l(c_r, w_r l_r / w_n) \frac{1}{w_n}}{u_l(c_r, l_r) \frac{1}{w_r}}} \frac{w_r l_r}{1 + \tau_x}, \end{aligned}$$

where $\eta_n \pi_r$ denotes the Lagrange multiplier of the incentive constraint of the non-routine worker.

One property of the original Mirrlees (1971) model is that the labor-supply decision of the high-ability agent should not be distorted. In our model, non-routine workers are subsidized at the margin when automation is incomplete. This subsidy corrects for the difference between the productivity as perceived by the firm (w_n) and the marginal increase in the resources available to the planner from a marginal increase in $\pi_n l_n$, which is equal to $w_n (\tau_x + \alpha) / \alpha (1 + \tau_x)$, where $(\tau_x + \alpha) / \alpha (1 + \tau_x) > 1$.¹²

Routine workers are taxed at the margin when automation is incomplete for two reasons. First, this tax corrects the distortion created by robot taxes, which make the wages of routine workers higher than the marginal increase in the resources available to the planner from a marginal increase in $\pi_r l_r$. Second, taxing routine workers makes it less appealing for non-routine workers to mimic routine workers and loosens the IC of non-routine workers.

Figure 3 illustrates the properties of the equilibrium associated with Mirrleesian optimal taxation. The process of automation begins later in the Mirrleesian solution than in the first best. This property reflects the presence of robot taxes in the Mirrleesian solution. These taxes increase the wages of routine workers and decrease the wages of non-routine workers. This wage compression loosens the incentive constraint of non-routine workers, which allows the government to redistribute more income from non-routine to routine workers.

The path for the tax rate on robots has a hump shape. The economy starts with inequality in wages that makes redistribution desirable. Since initially the cost of distorting automation is relatively small, the planner chooses a level of robot taxes which halts the process of automation. As the costs of automation fall, robot taxes increase to prevent automation from occurring. After this initial period, robot taxes

¹²The general-equilibrium effects emphasized by Stiglitz (1982) are reduced in our model to the impact of the robot tax on pre-tax wages.

fall. As robots become cheaper, it is inefficient to use routine workers so their labor supply falls. This decline in routine hours makes robot taxes less useful as a tool for income redistribution. In the limit, routine hours converge to zero and so do robot taxes.

Consumption of non-routine workers is higher than that of routine workers. Since non-routine workers work harder than routine workers, the former need to receive higher consumption to satisfy their incentive constraint. Both types of workers see their consumption rise as ϕ approaches zero. This outcome is achieved through large transfers to the routine workers.

In the limit, routine households stop working and live off government transfers. Those transfers are generous enough that the utilities of the two worker types are equalized. The reason for this equalization is that, once routine workers supply zero hours, there is no difference between the non-routine worker pretending to be routine and that of the routine worker.

2.4 Optimal policy with simple income taxes

In this subsection, we compare the Mirrleesian allocation with the solution to a Ramsey (1927)-style optimal taxation problem in which the tax schedule is assumed to take the simple form described in equation (13). This function has two parameters, γ and λ . When γ is zero, the tax system is linear with a rate λ . The tax system is progressive when γ is greater than zero and regressive when γ is lower than zero. This function has been widely used to study the U.S. income tax system and has recently been estimated by Heathcote, Storesletten and Violante (2017). Our goal is to assess how close the Ramsey solution with these simple taxes is to the Mirrleesian allocation. We also consider a version of the problem where we allow for lump-sum transfers. These transfers can be interpreted as a form of universal basic income.

We characterize the competitive equilibrium for this economy in the Appendix. Using these equations, we can write the ratio of the consumption of routine and non-

routine workers as:

$$\frac{c_r}{c_n} = \left[\frac{(1-\alpha)(1-m)}{\alpha} \frac{\pi_n}{\pi_r} \right]^{1-\gamma}. \quad (20)$$

Equation (20) shows that there are two ways to make the ratio c_r/c_n closer to one. One way is to raise τ_x which leads to a fall in the level of automation, m . The other way is to make γ closer to one, i.e. make the tax system more progressive. Both approaches have drawbacks. Taxing robot distorts production. Increasing progressivity reduces incentives to work.

Since we are interested in studying optimal taxation in an economy with automation, we focus on equilibria where $m > 0$. In this case, equation (20) can be written as:

$$c_r = c_n \left[\frac{(1-\alpha)}{\alpha} \frac{\pi_n}{\pi_r} \left(\frac{\phi(1+\tau_x)}{(1-\alpha)A} \right)^{1/\alpha} \frac{\pi_r l_r}{\pi_n l_n} \right]^{1-\gamma}. \quad (21)$$

This condition results from the fact that the government must set the same income tax schedule for both routine and non-routine agents.

The planner chooses allocations $\{c_r, l_r, c_n, l_n, G\}$ and the tax parameters $\{\tau_x, \gamma\}$ to maximize welfare, (15), subject to equation (21) and the following conditions

$$u_c(c_r, l_r)c_r + \frac{u_l(c_r, l_r)l_r}{1-\gamma} = 0, \quad (22)$$

$$u_c(c_n, l_n)c_n + \frac{u_l(c_n, l_n)l_n}{1-\gamma} = 0, \quad (23a)$$

and the resource constraint with interior automation, (16).

The constraints (22) and (23a) are the usual Ramsey implementability conditions. We show in the Appendix that equations (21)-(16) are necessary and sufficient conditions for a competitive equilibrium. We define η/c_r as the multiplier on the constraint (21) and by μ the multiplier on the resource constraint. In the Appendix we show that the optimal tax rate on robots satisfies:

$$\frac{\tau_x}{1+\tau_x} = \frac{\alpha}{1-\alpha} \frac{\eta(1-\gamma)}{\mu w_n \pi_n l_n}. \quad (24)$$

Since the marginal utility of public expenditures is always positive, the marginal value of resources to the planner, given by the multiplier μ , is strictly positive. The multiplier η captures the marginal value of redistributing income to routine households, which is limited by the assumptions on the income-tax function. If $\eta > 0$, the marginal value of additional redistribution of income towards routine workers is positive and robot taxes are strictly positive. The intuition for this result is that since the government has to use the same income tax function for both types of workers, taxing robots helps redistribute income by increasing the pre-income tax wage of routine workers and lowering that of non-routine workers.

Figure 4 shows that the form of the tax function constrains heavily the outcomes that can be achieved. As discussed above, the planner can redistribute income by taxing robots or by increasing progressivity. Taxing robots distorts production and higher progressivity reduces incentives to work. Since initially the cost of distorting automation is relatively small, the planner chooses a level of robot taxes consistent with no automation. Robot taxes are heavily used, reaching values as high as $\tau_x = 0.33$. As the costs of automation decline, the progressivity of the income tax rises. But there is still a large divergence in wage rates, consumption and utility across the two types of workers.

Optimal policy with lump-sum transfers The simple tax system studied in the previous section can require very high production distortions. This is because the alternative is to use progressive income taxes which heavily distort the labor supply. In this section, we study a simple modification of the tax system: we allow the government to give out a positive lump-sum transfer, Ω . In this specification, the after-tax income of household j is $\lambda(w_j N_j)^{1-\gamma} + \Omega$.

Combining the budget constraints of the two worker types, we obtain

$$\frac{c_r - \Omega}{c_n - \Omega} = \left[\frac{(1 - \alpha)(1 - m)}{\alpha} \frac{\pi_n}{\pi_r} \right]^{1-\gamma}. \quad (25)$$

Comparing this equation with (20), it is clear that higher transfers bring c_r and c_n closer together, holding everything else constant.

The planning problem is to choose allocations $\{c_n, c_r, l_r, l_n, G\}$ and the tax parameters $\{\tau_x, \gamma, \Omega\}$, subject to the following conditions

$$u_c(c_j, l_j)(c_j - \Omega) + \frac{u_l(c_j, l_j)l_j}{1 - \gamma} = 0, \quad j = r, n, \quad (26)$$

$$c_r - \Omega = (c_n - \Omega) \left[\frac{(1 - \alpha) \pi_n}{\alpha \pi_r} \left(\frac{\phi(1 + \tau_x)}{(1 - \alpha)A} \right)^{1/\alpha} \frac{\pi_r l_r}{\pi_n l_n} \right]^{1-\gamma}, \quad (27)$$

$$u(c_j, l_j) \geq u(\Omega, 0) \text{ if } \Omega \geq 0, \quad j = r, n, \quad (28)$$

and the resource constraint with interior automation, (16).

Conditions (26) are obtained from the budget constraints for each household type, combining the first-order conditions to replace prices and taxes. The second condition (27) imposes that the tax system is the same for both household types. With positive lump-sum transfers and regressivity, $\gamma < 0$, the solution to the households problem may not be interior, meaning that the household may choose to work zero hours and set consumption equal to the transfer. The conditions (28) impose that the household's allocation does not yield lower utility than that corner solution. We show in the appendix that these conditions are necessary and sufficient for a competitive equilibrium in the quantities $\{c_n, c_r, l_r, l_n, G\}$ and the tax parameters $\{\tau_x, \gamma, \Omega\}$.

The optimal tax on robots satisfies

$$\frac{\tau_x}{1 + \tau_x} = \frac{\alpha}{1 - \alpha} \frac{\eta(1 - \gamma)}{\mu w_n \pi_n l_n} \left(\frac{c_r - \Omega}{c_r} \right), \quad (29)$$

where η/c_r is the multiplier for the no-discrimination constraint (27) and μ is the multiplier for the resource constraint.¹³ If the lump-sum transfer is chosen to be zero, then the expression (29) is the same as in case without lump-sum transfers. The optimal plan may feature $c_r = \Omega$, in which case the routine worker supplies zero labor hours.

¹³See the appendix for the derivation.

In this limiting case, there is no role for redistributing by affecting the relative wages, and robot taxes should be set to zero.

Figure 5 illustrates the properties of this allocation. Workers have two sources of income: wages and transfers. For this reason, income and substitution effects of changes in wages are no longer offsetting. As a consequence, the two types of workers supply a different number of hours and their hours vary with ϕ .

The optimal solution with lump-sum transfers features full automation even when robots are still relatively expensive. In contrast, full automation occurs only asymptotically with Mirrleesian taxes and optimal simple taxes. Once full automation occurs, routine households have no labor income. Since only non-routine workers pay income taxes, the planner designs the tax system to avoid distorting their marginal labor-supply decisions. This result is achieved by increasing the regressivity of the tax system (i.e. lowering $\gamma < 0$) so that, given the level of taxation implied by λ , the marginal income tax rate for non-routine workers is zero. In this way, the taxation of non-routine households is effectively equivalent to lump-sum taxes.¹⁴ The level of transfers is chosen so that the non-routine worker is indifferent between the interior solution, with positive labor, and the corner solution with zero labor and consumption equal to transfers. For a higher level of transfers, both agents would supply zero labor.¹⁵

The availability of lump-sum transfers is essential for full automation to occur in finite time. Without lump-sum transfers, a routine worker who drops out of the labor force has zero consumption. For this reason, routine workers never drop out of the labor force.

One surprising result is that full automation occurs with this restricted tax system and not in the Mirrleesian solution which also allows for lump-sum transfers. The intu-

¹⁴The same logic implies that in a representative-agent economy it is possible to use the Heathcote et al. (2018) tax function to obtain the same allocation as with lump-sum taxes. This allocation is achieved by choosing a regressive tax system such that the marginal tax rate is zero. Government expenditures are financed with the revenue raised by the infra-marginal tax rates.

¹⁵In the appendix, we show a numerical example of the individual agent's problem with one such solution.

ition for this result is as follows. Because preferences are separable and have constant Frisch elasticity, the marginal disutility of labor converges to zero as labor hours approach zero. For this reason, a Mirrleesian planner finds it optimal to have routine workers supply positive labor, even if productivity is very low. With a general tax function it is possible to implement very different marginal distortions for both types of worker. As we have seen before, the Mirrleesian solution features a negative marginal distortion for non-routine workers, and a positive marginal distortion for routine workers. This solution cannot be obtained with the simple tax function. The restrictions on the marginal distortions associated with this function are such that the planner prefers a corner solution for the labor supply of routine workers. By excluding routine workers from the labor force, the planner can design the simple tax system to target only non-routine workers, reducing their marginal tax rate to zero so that in effect they are taxed in a lump-sum fashion.

In such an equilibrium, income is redistributed through a large lump-sum transfer. This transfer can be interpreted as a minimum income that is guaranteed to all agents in the economy. When automation is incomplete, robot taxes are used as an additional source of redistribution and τ_x can be as high as 37 percent. Complete automation occurs shortly after 2050, once the cost of robots drops below $\phi = 0.27$.

2.5 Comparing different policies

In this section, we compare the first-best allocation with the allocations associated with different policies in terms of social welfare and the utility of routine and non-routine workers. In the figures discussed below, we use the labels FB, SQ, OT, ST and STL to refer to the first-best, status-quo, Mirrleesian optimal taxes, simple taxes, and simple taxes with lump-sum transfers, respectively.

Figure 6 shows the welfare of the utilitarian social planner under the different policies, between the years 2000 and 2150. Recall that the level of ϕ in the year 2000 was chosen as the lowest value for which there is no automation in the status quo. Social

welfare rises as the costs of automation fall both for the first best and for all the policies we consider. We see that the Mirrlees allocation is relatively close in terms of welfare to the first-best allocation. The solution with simple taxation and lump-sum transfers ranks next in terms of welfare, followed by the solution with simple taxes without rebates. The status quo is by far the worst allocation.

A fall in the cost of automation can have very different consequences for routine and non-routine workers. To illustrate this property, we measure the utility of the two types of workers relative to the status-quo equilibrium in year 2000. We call this allocation the no-automation benchmark. Panel A (B) of Figure 7 shows how much routine (non-routine) workers would have to be compensated in the no-automation benchmark to be as well off as in the policy under consideration, for the different years. The measure is computed as a percentage of consumption.

Panel A of Figure 7 shows that the utility of routine workers in the first-best allocation improves as ϕ falls over time. In contrast, in the status quo, routine workers become increasingly worse as automation becomes more pronounced. With Mirrleesian optimal taxation, routine workers are always made better off. With simple income taxes, routine workers are not made better off until after 2150. We can see that including a universal form of income is a simple way of recovering gains for routine workers. Indeed, shortly after 2050 the routine worker is almost as well off in this solution as in the solution with Mirrleesian taxes.

Panel B of Figure 7 shows that non-routine workers prefer the no-automation benchmark to the first best while automation costs are relatively high (almost until 2100). This preference reflects the large transfers that non-routine workers make to routine workers in the first best. Once automation costs fall by 62 percent relative to their 2000 value, which happens in 2116, non-routine workers prefer the first best to the no-automation benchmark. The reason is that the wage of non-routine workers is high enough to compensate for the transfers they make to routine workers. Starting in 2015, non-routine workers prefer the status quo to all other allocations. This preference

results from a combination of high wages and relatively low taxes in the status quo.

Routine workers always rank the first-best allocation first, Mirrleesian optimal taxation second, simple taxes with lump-sum transfers third, simple taxes without lump-sum transfers fourth, and the status quo last. The utility for the allocation with simple taxes and lump-sum transfers approaches that under Mirrleesian taxes as the cost of automation falls. In contrast, non-routine workers rank the status quo first and the first best last. Mirrleesian optimal taxation and simple taxes with and without transfers rank in between the two extremes.

2.6 The simple model with endogenous occupation choice

In this section, we study the optimal tax policy in a version of our model that allows agents to choose whether to become routine or non-routine workers. In this model, taxing robots affects the relative wages of routine and non-routine workers thereby affecting occupation choices.

Our analysis is related to Saez (2004), Scheuer (2014), Rothschild and Scheuer (2013), and Gomes, Lozachmeur and Pavan (2017). These authors characterize Mirrlees-style optimal tax plans in models with endogenous occupation choice. Saez (2004) considers a setting in which agents choose their occupation but hours worked are fixed. Income is proportional to the wage rate so the government can infer a worker's occupation from the workers' income. This property allows the government to design the income tax schedule to effectively tax different occupations at different rates. As a result, it is not optimal to distort production. Gomes, Lozachmeur and Pavan (2017) consider a setting in which agents choose both their occupation and hours worked. They find that the optimal tax plan does not feature production efficiency.

In our model, workers have different preferences for the two occupations. They choose both their occupation and the number of hours worked. In our numerical examples, production efficiency is generally not optimal as long as the costs of changing occupations are significantly high.

Households There is a continuum of measure one of households indexed by the occupation preference parameter $\theta \in \Theta \subseteq \mathbb{R}$. This parameter is drawn from a distribution F with continuous density f . Household preferences are given by

$$u(c_\theta, l_\theta) + v(G) - \mathcal{O}_\theta \theta, \quad (30)$$

where c_θ and l_θ denote the consumption and income of household θ , respectively. The indicator $\mathcal{O}_\theta$ denotes the household's occupation choice. It takes the value 1 when the household chooses a non-routine occupation and 0 otherwise. The wage rate earned by the household depends on the individual occupation choice. It is equal to w_r if the household chooses a routine occupation and equal to w_n otherwise.

The utility representation above has the following interpretation: households have heterogeneous preferences with respect to different occupations. A household with a positive θ prefers, all else equal, a routine occupation. A household with a negative θ prefers, all else equal, a non-routine occupation.

The households maximize their utility subject to the budget constraint

$$c_\theta \leq w_\theta l_\theta - T(w_\theta l_\theta). \quad (31)$$

It is useful to define the set of households that choose to become non-routine workers, $\Theta_n \equiv \{\theta \in \Theta : \mathcal{O}_\theta = 1\}$ and the set of households that choose to become routine workers $\Theta_r \equiv \Theta \setminus \Theta_n$.

Mirrleesian optimal taxation The production side is the same as in previous sections. We maintain the assumption that the only instrument the government has to directly affect production is a proportional tax on robots. With these assumptions, the firms' decisions, when automation is interior, can be summarized by the constraint

$$\int_{\Theta} c_\theta f(\theta) d\theta + G \leq \frac{\tau_x + \alpha}{\alpha(1 + \tau_x)} \int_{\Theta_n} w_\theta l_\theta f(\theta) d\theta + \frac{\int_{\Theta_r} w_\theta l_\theta f(\theta) d\theta}{1 + \tau_x}. \quad (32)$$

As in Rothschild and Scheuer (2013), we characterize a direct implementation where households declare their type θ and get assigned an allocation $(c_\theta, l_\theta, \mathcal{O}_\theta)$. Income

and consumption are observable, but the government cannot observe labor, wage or sectoral choice. Given this informational asymmetry, the constraints that guarantee truth-telling are as follows. The first condition is the same incentive constraint on the choice of hours worked that we have used before:

$$u(c_\theta, l_\theta) \geq u\left(c_{\theta'}, \frac{w_{\theta'}}{w_\theta} l_{\theta'}\right), \quad \forall \theta, \theta' \in \Theta. \quad (33)$$

This *labor-supply incentive constraint* guarantees that households choose the assigned allocation, given the occupation choice.

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

$$u(c_\theta, l_\theta) - \mathcal{O}_\theta \theta \geq u(c_{\theta'}, l_{\theta'}) - \mathcal{O}_{\theta'} \theta, \quad \forall \theta, \theta' \in \Theta. \quad (34)$$

This *occupation-choice incentive constraint* ensures that households choose their assigned occupation. The other conditions are that the occupation choice, $\mathcal{O}_\theta$, defines the sets Θ_n and Θ_r and that (32) holds.¹⁶

To characterize the necessary and sufficient conditions for incentive-compatible occupation choice, it is useful to define $U_\theta \equiv u(c_\theta, l_\theta)$.

Lemma 2. *An allocation is incentive compatible for occupation choice if and only if for $\theta, \theta' \in \Theta_i$ then $U_\theta = U_{\theta'} \equiv U_i$ for $i = r, n$, and there exists a threshold $\theta^* = U_n - U_r$ such that (i) if $\theta \leq \theta^*$ then $\theta \in \Theta_n$; (ii) if $\theta > \theta^*$ then $\theta \in \Theta_r$.*

This result follows directly from the incentive constraints. The first part of the lemma states that all agents who share the same occupation choice must have the same

¹⁶These constraints do not explicitly take into account the possibility that agent θ might choose an allocation $(C_{\theta'}, Y_{\theta'})$ at a different occupational choice than $\mathcal{O}_{\theta'}$. However, those additional constraints are redundant. To see this result, suppose that agent θ deviates to an allocation $(C_{\theta'}, Y_{\theta'})$ and occupational choice $\mathcal{O}_{\hat{\theta}}$ which is different from that of $\mathcal{O}_{\theta'}$. From the intensive margin incentive constraint for agent $\hat{\theta}$ it follows that $u(C_{\theta'}) - v(Y_{\theta'}/w_{\hat{\theta}}) - \mathcal{O}_{\hat{\theta}} \theta \leq u(C_{\hat{\theta}}) - v(Y_{\hat{\theta}}/w_{\hat{\theta}}) - \mathcal{O}_{\hat{\theta}} \theta \leq u(C_\theta) - v(Y_\theta/w_\theta) - \mathcal{O}_\theta \theta$, where the last inequality follows from the extensive margin incentive constraint for θ . This condition also guarantees that when choosing his own assignment, $\mathcal{O}_\theta$ is optimal.

utility level. This property results from the fact that they can mimic the choices of another household in the same group at zero cost. So, if an allocation was better than all others for the same occupation, all agents would choose that allocation.

The second part of the proposition establishes that the incentive constraints for the extensive margin choice can be summarized by a single threshold rule. Agents with θ lower than θ^* choose a non-routine occupation and the remaining agents become routine workers.

The planner maximizes a weighted average of utilities, where the weight on agents of type θ is $\omega(\theta) \geq 0$ with $\int \omega(\theta) f(\theta) d\theta = 1$.

Proposition 2. *In the optimal plan, if $\theta, \theta' \in \Theta_i$ then $w_\theta l_\theta = w_{\theta'} l_{\theta'} \equiv w_i l_i$ and $c_\theta = c_{\theta'} \equiv c_i$, for $i = r, n$.*

Agents who choose the same occupation have the same preferences for consumption and leisure and the same productivity. They differ only in their value of θ which enters separably in their utility function. Since the planner has an utilitarian welfare function, the optimal plan sets the same consumption and hours worked for all agents with the same occupation.

Using these results, we can see that, for a fixed τ_x , the optimal plan solves the following optimization problem:

$$W(\tau_x) = u(c_n, l_n) \int_{-\infty}^{\theta^*} \omega(\theta) f(\theta) d\theta + u(c_r, l_r) \int_{\theta^*}^{\infty} \omega(\theta) f(\theta) d\theta + v(G) - \int_{-\infty}^{\theta^*} \omega(\theta) \theta f(\theta) d\theta,$$

subject to

$$\theta^* = u(c_n, l_n) - u(c_r, l_r),$$

$$F(\theta^*)c_n + [1 - F(\theta^*)]c_r + G \leq F(\theta^*)w_n l_n \frac{\tau_x + \alpha}{\alpha(1 + \tau_x)} + \frac{[1 - F(\theta^*)]w_r l_r}{1 + \tau_x}.$$

and two incentive constraints which are the same as those of the model with fixed occupations (equations (17) and (18)). Optimizing with respect to τ_x requires $W'(\tau_x) = 0$.

The solution to this problem can be decentralized with a mechanism in which the government sets τ_x to its optimal level and sets income taxes such that

$$T(w_n l_n) = w_n l_n - c_n, \quad \text{and} \quad T(w_r l_r) = w_r l_r - c_r,$$

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

Numerical analysis We now explore the properties of the occupation choice model with some numerical examples. We use the same preference and production parameters as in previous sections.¹⁷ We set $\Theta = \mathbb{R}$ and assume that $F(\theta)$ is a normal distribution with mean zero and standard deviation σ . This choice ensures that half the population has a preference for routine work and the other half for non-routine work. We solve the model for two values of σ , 1 and 7.5.

Figure 8 shows the first-best solution for both values of σ . We can see that lower dispersion of θ is associated with a higher share of non-routine workers. This property makes intuitive sense, since having more agents with θ close to zero implies that it is easier to switch them to more productive non-routine occupations. This higher fraction of non-routine workers also results in higher consumption and lower working hours for all households. Because it is easier to switch agents to non-routine occupations, automation advances faster when dispersion is lower.

Figure 9 shows the Mirrleesian solution for the same levels of σ . When dispersion is high, robot taxes are positive and are similar to the benchmark case without occupation choice. However, in the case of $\sigma = 1$ dispersion is low and the tax on robots is always zero.

Taxation has a direct distribution effect on after-tax income and an indirect distribution effect on the choice of occupation. The optimal plan balances the costs and benefits of these two forms of redistribution. For $\sigma = 1$, the costs of indirect redistribution are low, so the planner induces a higher fraction of the population to become

¹⁷In the numerical analysis we compare the Mirrleesian solution to the first-best solution in this environment. We solve for the first best in the appendix.

non-routine. In fact, for low automation costs the share of non-routine workers becomes almost as high as in the first-best allocation. Since taxes on robots are desirable only insofar as they improve the direct redistribution mechanism, the reduced importance of this form of redistribution also implies that taxes on robots should be lower.

For $\sigma = 7.5$, the costs of indirect redistribution are high, so the planner resorts to using more direct redistribution. This approach results in higher consumption and lower hours worked for routine workers. Because hours worked decline faster when σ is higher, automation also advances more rapidly as ϕ falls over time. We can also see that the share of non-routine workers is much lower than in the first best.

3 A dynamic model

In this section, we generalize the previous results to a dynamic model in which robots are an investment good.

Households and preferences As in our static environment, the model features two types of households, routine (r) and non-routine (n). We normalize the size of the population to one, and assume that there are π_r routine workers and π_n non-routine workers. These households differ permanently in the type of labor they supply.

The utility function of household j is given by

$$U_j = \sum_{t=0}^{\infty} \beta^t [u(c_{j,t}) - h(l_{j,t}) + v(G_t)],$$

where $u'(c) > 0$, $h'(l) > 0$, $v'(G) > 0$, and $u''(c) < 0$, $h''(l) > 0$, $v''(G) < 0$ denote respectively the first and second derivatives of each function.

As is well known, without weak separability between consumption and leisure, the uniform taxation result in Atkinson and Stiglitz (1976) fails. In a dynamic setting, this failure would mean that the optimal plan features intertemporal distortions, i.e. it is optimal to tax consumption in different periods at different rates. This reason

to tax capital is orthogonal to the one we focus on and, for that reason, we assume separability.¹⁸

Technology Production of the final good combines robots with routine and non-routine workers, according to the production function, $Y_t = F(N_{n,t}, N_{r,t}, X_t)$ where Y_t denotes total production at time t , and $N_{r,t}$ and $N_{n,t}$ denote the total number of routine and non-routine labor hours supplied by the households, respectively.

We assume that the production function has constant returns to scale. An household's effective contribution to production $y_{j,t}$ is observable and equal to the household's marginal productivity multiplied by the total number of hours supplied $y_{j,t} = F_j(t)l_{j,t}$, for $j = r, n$, where $F_j(t) = \partial F(N_{r,t}, N_{n,t}, X_t) / \partial N_{j,t}$. We define the marginal productivity of robots in period t as $F_x(t) = \partial F(N_{r,t}, N_{n,t}, X_t) / \partial X_t$. As in the previous section, we assume that robots share a higher degree of complementarity with non-routine work than with routine work. For that end, we assume that $\varepsilon_{F_r/F_n, x}(t) = d \log (F_r(t)/F_n(t)) / d \log X_t < 0$.

A unit of robot capital is produced using ϕ_t units of output. The stock of robots evolves according to:

$$X_{t+1} = (1 - \delta_X)X_t + i_t/\phi_t,$$

where δ_X denotes the depreciation rate of robots. The resource constraint in period t can be written as

$$\pi_r c_{r,t} + \pi_n c_{n,t} + G_t + \phi_t [X_{t+1} - (1 - \delta_X) X_t] \leq F(\pi_n l_{n,t}, \pi_r l_{r,t}, X_t). \quad (35)$$

First-best allocation Given any Pareto weights for routine and non-routine households, $\omega_r, \omega_n \geq 0$, normalized such that $\omega_r \pi_r + \omega_n \pi_n = 1$, the planner's objective

¹⁸Since we assume that utility is time separable, weak separability between the vector of consumptions and the vector of hours worked requires strict separability between consumption and leisure in each period.

function is

$$\sum_{t=0}^{\infty} \beta^t \left\{ \sum_{j=n,r} \omega_j \pi_j [u(c_{j,t}) - h(l_{j,t})] + v(G_t) \right\}. \quad (36)$$

The first-best allocation maximizes this welfare function subject to the resource constraints (35). The solution to this problem implies the following efficiency conditions, which equate marginal rates of substitution to marginal rates of transformation

$$\frac{h'(l_{j,t})}{u'(c_{j,t})} = F_j(t),$$

and

$$\frac{u'(c_{j,t})}{\beta u'(c_{j,t+1})} = \frac{F_x(t+1) + \phi_{t+1}(1 - \delta_x)}{\phi_t},$$

for $j = r, n$. In addition to these conditions, the first-best allocation also requires that

$$\omega_r u'(c_{r,t}) = \omega_n u'(c_{n,t}).$$

These conditions together with the resource constraint with equality characterize the Pareto frontier, as we let the weights vary.

Planning problem We solve the problem of a planner who maximizes a weighted-sum of household utilities subject to incentive and resource constraints. The solution of this problem cannot yield lower welfare than the solution of a non-linear taxation problem, and, as we show in the appendix, it is possible to find taxes that implement the optimal solution.

The incentive constraints for this problem are

$$\sum_{t=0}^{\infty} \beta^t [u(c_{r,t}) - h(l_{r,t})] \geq \sum_{t=0}^{\infty} \beta^t \left[u(c_{n,t}) - h\left(\frac{F_n(t)}{F_r(t)} l_{n,t}\right) \right], \quad (37)$$

$$\sum_{t=0}^{\infty} \beta^t [u(c_{n,t}) - h(l_{n,t})] \geq \sum_{t=0}^{\infty} \beta^t \left[u(c_{r,t}) - h\left(\frac{F_r(t)}{F_n(t)} l_{r,t}\right) \right]. \quad (38)$$

The planning problem is to maximize (36) subject to (37) and (38) and the resource constraints (35). We characterize the optimal allocations in which the incentive constraint of the non-routine worker (38) binds and the incentive constraint of the routine

worker (37) does not bind. This pattern generally holds whenever the planner values redistribution to the routine worker. We check that this property holds in our numerical exercises.

It is useful to define the following agent-specific intratemporal and intertemporal wedges

$$\tau_{j,t}^n \equiv 1 - \frac{h'(l_{j,t})}{u'(c_{j,t})} \frac{1}{F_j(t)}, \quad (39)$$

and

$$\tau_{j,t+1}^k \equiv 1 - \frac{u'(c_{j,t})}{\beta u'(c_{j,t+1})} \frac{\phi_t}{F_x(t+1) + \phi_{t+1}(1 - \delta_x)}. \quad (40)$$

The next proposition states results analogous to those we obtained for the static model: as long as automation is incomplete, there are positive wedges in the accumulation of robots.

Proposition 3. *In the optimal plan, the intertemporal wedge is the same for the two worker types, $\tau_{r,t+1}^k = \tau_{n,t+1}^k$. If routine labor hours are strictly positive ($l_{r,t+1} > 0$), then the intertemporal wedge is strictly positive, $\tau_{r,t+1}^k = \tau_{n,t+1}^k > 0$.*

Proof: See Appendix.

As in the static model, it is optimal to tax robots when the hours worked by routine households are positive. Since robots are a form of capital, their use is taxed by creating a positive intertemporal wedge which distorts the accumulation of capital goods. Notice also that in this economy the intertemporal marginal rates of substitution of the two agents are equated.

To compare the dynamic model with the static model, it is useful to interpret the intertemporal wedge as a tax on the rental cost of capital. This rental cost, inclusive of tax, is given by

$$\left[\phi_t \frac{u'(c_{j,t})}{\beta u'(c_{j,t+1})} - \phi_{t+1}(1 - \delta_x) \right] (1 + \tau_{t+1}^x), \quad (41)$$

where τ_{t+1}^x denotes the tax rate. This rate is not indexed by j because the optimal intertemporal wedges are the same for all worker types.

The intuition for the expression for the rental cost is as follows. Renting out a unit of robot capital at time $t + 1$ requires investing ϕ_t in robots at time t . This investment has an opportunity cost of $\phi_t u'(c_{j,t}) / \beta u'(c_{j,t+1})$. At the end of the period, the rentor receives $1 - \delta_x$ robot units each of which is worth ϕ_{t+1} . Since the rental cost inclusive of taxes is equated to the marginal productivity of robots ($F_x(t + 1)$), the optimal tax rate, τ_{t+1}^x , is given by:

$$\left[\phi_t \frac{u'(c_{j,t})}{\beta u'(c_{j,t+1})} - \phi_{t+1}(1 - \delta_x) \right] (1 + \tau_{t+1}^x) = F_x(t + 1).$$

Quantitative analysis In this section, we solve the planning problem and quantify the effects of advances in automation for the optimal taxation of robots and income. We assume that the cost of robots declines geometrically over time, $\phi_t = \phi e^{-g_\phi t}$, which means that the dynamic problem features investment-specific technical change as in Greenwood, Hercowitz and Krusell (1997). The baseline scenario assumes perfect substitutability between robots and routine workers, as in related work on the effects of automation (see Acemoglu and Restrepo (2018a) and Cortes, Jaimovich, and Siu (2017)). The production function is given by

$$F(\pi_n l_n, \pi_r l_r, X) = A (X + \pi_r l_r)^{1-\alpha} (\pi_n l_n)^\alpha,$$

where A denotes total factor productivity and α is the share of non-routine workers in production. The assumption that the elasticity between routine and non-routine labor is unitary is important because it ensures the existence of a balanced growth path which is reached asymptotically, as the hours supplied by routine workers converge to zero.

We use the same preferences as in our static model (equation (14)):

$$U_j = \sum_{t=0}^{\infty} \beta^t \left[\log(c_{r,t}) - \zeta \frac{l_{r,t}^{1+\nu}}{1+\nu} + \chi \log(G_t) \right].$$

Recall that these preferences are compatible with balanced growth and consistent with the empirical evidence reviewed in Chetty (2006).

The optimal plan maximizes $\omega_r \pi_r U_r + \omega_n \pi_n U_n$ subject to the incentive constraint (38) and the flow resource constraint (35). We check numerically that the incentive constraint (37) does not bind.

The model is calibrated with the same values of ω_r , ω_n , π_r , π_n , ζ , ν , χ , α , A , and g_ϕ used in the static model. We choose $\beta = 1/1.04$, $\delta_x = 0.1$. We set the initial stock of robots in 2000 equal be zero so, as in the static model, there is no automation in the year 2000. We choose the cost of robots in 1999 to be the lowest value consistent with no automation in 2000 in the status-quo equilibrium. We solve our dynamic model using a modified version of the algorithm proposed by Slavík and Yazici (2014) which we discuss in the appendix.

Figure 10 displays the optimal policy solution. The properties of this solution are similar to those of the static model. The maximum value of the optimal robot tax is 14 percent, which is higher than the maximum value attained in the static model (9 percent). The tax rate on robots converges to zero as the degree of automation, defined as $m_t = X_t / (X_t + \pi_r l_{r,t})$, converges to one.

Technical progress induces a fall over time in the relative productivity of routine workers. For this reason, it is optimal for the number of routine hours of work to decline over time. As routine hours fall, there is less incentive for the planner to tax robots to distort the ratio of wages and loosen the incentive constraint of non-routine workers. As routine hours converge to zero, the optimal robot tax converges to zero. This mechanism is also present in our static model.

4 Relation to the public finance literature

In this section we discuss how our results relate to classical results on production efficiency and taxation of capital in the public finance literature.

Relating our results to Diamond and Mirrlees (1971) Our results stand in sharp contrast to the celebrated Diamond and Mirrlees (1971) result that an optimal tax system should ensure efficiency in production and therefore leave intermediate goods untaxed. In our framework, this property would imply that the tax on robots should be zero.

At the heart of the failure of the Diamond and Mirrlees (1971) intermediate-good theorem in our model is the fact that the government cannot discriminate between the two types of workers. If tax functions could be worker specific, production efficiency would be recovered in our model. To see this result, consider type-specific tax functions of the form used by Heathcote et al. (2018) with different tax levels, λ_r and λ_n , but with the same progressivity parameter

$$T_i(w_i l_i) = w_i l_i - \lambda_i (w_i l_i)^{1-\gamma}.$$

In this case, household optimality requires

$$\frac{-u_l(c_j, l_j)l_j}{u_c(c_j, l_j)} = \lambda_j (1 - \gamma) (w_j l_j)^{1-\gamma}, \quad \text{and} \quad c_j = \lambda_j (w_j l_j)^{1-\gamma}.$$

Given that the planner can choose λ_r and λ_n to target each marginal rate of substitution independently, the only constraints faced by the planner are the resource constraint (16) which can be written as

$$\pi_r c_r + \pi_n c_n + G \leq \frac{\tau_x + \alpha}{\alpha(1 + \tau_x)^{1/\alpha}} \frac{\alpha A^{1/\alpha} (1 - \alpha)^{\frac{1-\alpha}{\alpha}}}{\phi^{\frac{1-\alpha}{\alpha}}} \pi_n l_n + \phi \pi_r l_r,$$

and the implementability conditions

$$u_c(c_j, l_j)c_j + \frac{u_l(c_j, l_j)l_j}{1 - \gamma} = 0, \quad \text{for } j = r, n.$$

These three conditions are necessary and sufficient for an equilibrium. Recall that the term $(\tau_x + \alpha) / \alpha(1 + \tau_x)^{1/\alpha} \leq 1$, and is strictly less than one if $\tau_x \neq 0$.

The robot tax only affects directly the resource constraint and not the implementability conditions. Since the robot tax does not interfere with incentives, it is

chosen to maximize output for given levels of hours worked. This objective is achieved by not distorting production, setting $\tau_x = 0$.

When the tax system requires that all workers face the same income-tax function ($\lambda_r = \lambda_n$), the planner must satisfy the following additional implementability constraint

$$\frac{c_r}{c_n} = \left(\frac{w_r l_r}{w_n l_n} \right)^{1-\gamma}. \quad (42)$$

The value of τ_x no longer appears only in the resource constraint; it also appears in equation (42) because the wage ratio is a function of τ_x . To relax restriction (42), it might be optimal to choose values of τ_x that are different from zero. This result depends crucially on the fact that different labor types interact differently with the intermediate good, which means that distorting the use of intermediate goods affects in different ways the wage rates of routine and non-routine workers. If the production function was weakly separable in labor types and intermediate inputs, the wage ratio would be independent of the usage of intermediate inputs and production efficiency would be optimal. In our model, robots are substitutes of routine workers and complements of non-routine workers. A tax on robots decreases the wage rate of non-routine workers and increases the wage rate of routine workers. This property implies that it can be optimal to use robot taxes.

Relating our results to Atkinson and Stiglitz (1976) Our result that in the Mirrleesian optimal taxation problem production efficiency is not optimal stands in contrast with the well-known result in Atkinson and Stiglitz (1976) that, for preferences that are separable in commodities and leisure, uniform commodity taxation is optimal. Since uniform taxation can be interpreted as production efficiency, their result seems to contradict ours.

The difference between our results and those of Atkinson and Stiglitz (1976) stem from the determinants of worker productivity. In Atkinson and Stiglitz (1976), workers' productivities are exogenous. In our setup workers' productivity are endogenous so,

it may be optimal to deviate from production efficiency to induce changes in those productivities. In particular, by taxing robots the Mirrleesian planner is able to change pre-tax wages through general-equilibrium effects, relaxing the incentive constraint, and improving welfare.¹⁹

Naito (1999) shows that uniform taxation may not be optimal in an economy in which the intensity of high- and low-skilled workers in production varies across goods. This form of production non-separability implies that commodities interact differently with different agent types and, as a result, it might be optimal to deviate from uniform commodity taxation.²⁰ Similarly, in our model the assumption that production is not separable in the use of robots and the two labor types is key to generate deviations from production efficiency.

The intuition for the importance of general-equilibrium effects of taxation on wages and prices is the same we emphasized in our discussion of proposition 1. Because the planner does not know the type of the agent and only observes income, it is restricted to use incentive-compatible tax systems. Since different types interact differently with the intermediate good, distorting production decisions may help in the screening process. To see this property, it is useful to write the incentive constraint as: $u(c_i, l_i) \geq u(c_j, w_j l_j / w_i)$. Crucially, this incentive constraint involves the wage ratio. Whenever the taxation of intermediate goods affects this ratio, production efficiency may no longer be optimal. When intermediate goods are not separable in production from the two labor types, taxing intermediate goods affects the wage ratio and it might be optimal to distort production.

The importance of general-equilibrium effects of taxes on wages in shaping the optimal tax policy was originally emphasized by Stiglitz (1982) and Stern (1982) in

¹⁹An important assumption is that when workers imitate others, they retain their productivity. See Scheuer and Werning (2016) for a discussion.

²⁰Jacobs (2015) shows that production efficiency is generally not optimal in a model where commodity prices are exogenous but wages are not. In his model, goods are produced with commodities and labor according to production functions that are worker specific. Taxation of commodities has a differential impact on the marginal productivities and wages of the different workers.

a Mirrlees (1971) environment. Mirrlees assumes that production is linear in labor, so taxation does not affect wages through general-equilibrium effects. Stiglitz (1982) and Stern (1982) show that when production is not linear in labor, the optimal tax schedule is more regressive than in the Mirrlees model and the top marginal income tax is negative instead of zero.²¹ The reason for this result is that it is optimal to encourage high-skilled workers to exert more effort so as to reduce their relative wages, making their incentive constraint easier to satisfy.

Relating our results to the literature on the optimal taxation of capital The public-finance literature discusses several reasons why it might be optimal to tax capital, introducing intertemporal distortions.²² First, it might be optimal to use intertemporal distortions to confiscate the initial stock of capital. Second, intertemporal distortions can be optimal when the elasticities of the marginal utility of consumption and labor are time varying. Third, intertemporal distortions can be used to provide insurance in models with idiosyncratic risk. All three reasons are absent in our model.

We consider Mirleesian taxes which allow for lump-sum taxation so there is no reason to confiscate the initial stock of capital. In addition, we assume that utility is separable in consumption and labor and the disutility of labor is isoelastic. Werning (2007) shows that under these conditions and with perfect substitutability of labor types, the optimal tax on capital is zero. Because our dynamic model abstracts from idiosyncratic risk, the reasons for capital taxation discussed in Golosov, Kocherlakota and Tsyvinski (2003) do not apply.

Our results are related to work by Slavík and Yazici (2014). These authors study optimal taxation in a model with two types of capital, structures and equipment, and

²¹Rothschild and Scheuer (2013) generalize the results of Stern (1982) and Stiglitz (1982) to an environment in which occupational choice is endogenous and there is a continuous distribution of agent types.

²²For a recent overview of the literature on optimal capital taxation in a dynamic Ramsey setting see Chari, Nicolini and Teles (2018).

no technical progress. They assume that equipment raises the marginal product of skilled workers relative to that of unskilled workers. In their set up, the optimal tax on equipment rises over time. In contrast, the optimal tax on robots in our model converges to zero. As discussed in Section 3, this property reflects the presence of technical progress in our model.

In sum, the classical results on production efficiency in the public finance literature depend on one of two key assumptions: (i) the government can tax differently every consumption good and labor type; or (ii) the environment is such that production distortions do not help in shaping incentives. Both assumptions fail in our model. On the one hand, the government cannot design the income tax system to independently target each type of worker. On the other hand, robots are substitutes for routine workers and complements to non-routine workers, so a tax on robots affects the ratio of the wages of these two types of workers.

5 Conclusions

Our analysis suggests that without changes to the current U.S. tax system, a sizable fall in the costs of automation would lead to a massive rise in income inequality. Even though routine workers keep their jobs, their wages fall to make them competitive with the possibility of automating production.

Income inequality can be reduced by raising the marginal tax rates paid by high-income individuals and by taxing robots to raise the wages of routine workers. But this solution involves a substantial efficiency loss. A Mirrleesian optimal income tax can reduce inequality at a smaller efficiency cost than the variants of the U.S. tax system discussed above, coming close to the levels of social welfare obtained in the first-best allocation.

An alternative, less ambitious, approach is to amend the tax system to include a transfer that is independent of income. The desirability of this type of universal basic

income system has been debated since Thomas More proposed it in his 1516 book, *Utopia*. With this transfer in place, it is optimal in our model to tax robots for values of the automation cost that lead to partial automation. For values of the automation cost that lead to full automation, it is not optimal to tax robots. Routine workers lose their jobs and live off government transfers, just like in Kurt Vonnegut’s “Player Piano.”

6 References

Acemoglu, Daron and David Autor. “Skills, tasks and technologies: implications for employment and earnings.” in *Handbook of Labor Economics*, vol. 4B, edited by Orley Ashenfelter and David Card, chap. 12. Elsevier, (2011): 1043-1171.

Acemoglu, Daron and Pascual Restrepo. “Robots and Jobs: evidence from US labor markets.” mimeo, Massachusetts Institute of Technology (2018a).

Acemoglu, Daron and Pascual Restrepo. “Automation and New Tasks: The Implications of the Task Content of Technology for Labor Demand.” mimeo, Massachusetts Institute of Technology (2018b).

Ales, Lawrence, Musab Kurnaz, and Christopher Sleet. “Technical change, wage inequality, and taxes.” *American Economic Review*, 105(10) (2015): 3061-3101.

Atkinson, Anthony Barnes, and Joseph E. Stiglitz. “The design of tax structure: direct versus indirect taxation.” *Journal of Public Economics* 6.1-2 (1976): 55-75.

Autor, David H., David Dorn, and Gordon H. Hanson. “The China Syndrome: Local Labor Market Effects of Import Competition in the United States” *The American Economic Review*, Vol. 103, No. 6 (2013), 2121-2168.

Autor, David H., Frank Levy, and Richard J. Murnane. “The skill content of recent technological change: An Empirical Exploration,” *The Quarterly Journal of Economics*, Vol. 118, No. 4 (2003), pp. 1279-1333.

Chari, V. V., Juan Pablo Nicolini, and Pedro Teles. “Optimal capital taxation

revisited,” manuscript, Federal Reserve Bank of Minneapolis, (2018).

Chen, Yiqun G. “The Rise of the Machines and the U.S. Labor Market,” manuscript, (2016)

Cortes, Guido Matias, Nir Jaimovich, and Henry E. Siu. “Disappearing routine jobs: who, how, and why? ” *Journal of Monetary Economics*, Volume 91, (2017): 69-87.

Diamond, Peter A., and James A. Mirrlees. “Optimal taxation and public production I: production efficiency.” *The American Economic Review* 61.1 (1971): 8-27.

Golosov, Mikhail, Narayana Kocherlakota, and Aleh Tsyvinski. “Optimal Indirect and Capital Taxation.” *Review of Economic Studies*, Vol. 70 (2003), 569-587.

Gomes, Renato, Jean-Marie Lozachmeur, and Alessandro Pavan. “Differential taxation and occupational choice.” *The Review of Economic Studies* (2017): rdx022.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell. “Long-run implications of investment-specific technological change” *American Economic Review* 87.3 (1997): 342-362.

Guesnerie, Roger. *A contribution to the pure theory of taxation*. Cambridge-UK: Cambridge University Press, 1995.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. “Optimal tax progressivity: An analytical framework.” *The Quarterly Journal of Economics* 132.4 (2017): 1693-1754.

Jacobs, Bas “Optimal inefficient production,” manuscript, Erasmus University, Rotterdam (2015).

Mirrlees, James A. “An exploration in the theory of optimum income taxation.” *The Review of Economic Studies* 38.2 (1971): 175-208.

Naito, Hisahiro. “Re-examination of uniform commodity taxes under a non-linear income tax system and its implications for production efficiency.” *Journal of Public Economics* 71 (1999): 165-188.

Ramsey, Frank P. “A contribution to the theory of taxation.” *The Economic Journal* 37.145 (1927): 47-61.

Rothschild, Casey and Florian Scheuer. “Redistributive taxation in the Roy model,” *The Quarterly Journal of Economics*, Volume 128, Issue 2, (2013): 623–668.

Saez, Emmanuel. “Direct or indirect tax instruments for redistribution: short-run versus long-run.” *Journal of Public Economics* 88 (2004): 503-518.

Scheuer, Florian. “Entrepreneurial taxation with endogenous entry.” *American Economic Journal: Economic Policy* 6:2 (2014): 126-163.

Scheuer, Florian and Iván Werning “Mirrlees meets Diamond-Mirrlees,” Working paper No. w22076 (2016), National Bureau of Economic Research.

Slavík, Ctirad, and Hakki Yazici. “Machines, buildings, and optimal dynamic taxes.” *Journal of Monetary Economics* 66 (2014): 47-61.

Stern, Nicholas. “Optimum taxation with errors in administration.” *Journal of Public Economics*, 17:2 (1982): 181 - 211.

Stiglitz, Joseph. “Self-selection and Pareto efficient taxation.” *Journal of Public Economics*, 17:2 (1982): 213-240.

Werning, Iván. “Optimal Fiscal Policy with Redistribution.” *The Quarterly Journal of Economics*, Vol. 122, No.3 (2007), pp. 925-967.

A Appendix

A.1 The first-best allocation

We define the first-best allocation in this economy as the solution to an utilitarian welfare function, absent informational constraints. This absence implies that the planner can perfectly discriminate among agents and enforce any allocation. The optimal plan solves the following problem

$$\begin{aligned}
W &= \max \omega_r \pi_r [u(c_r, l_r) + v(G)] + \omega_n \pi_n [u(c_n, l_n) + v(G)]. \\
\pi_r c_r + \pi_n c_n + G &\leq A \left[\int_0^m x_i^\rho di + \int_m^1 n_i^\rho di \right]^{\frac{1-\alpha}{\rho}} (\pi_n l_n)^\alpha - \phi \int_0^m x_i di, \quad [\mu], \\
\int_m^1 n_i di &= \pi_r l_r, \quad [\eta].
\end{aligned}$$

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

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

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 where they are used, $x_i = x$ 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

$$\begin{aligned}
W &= \max \omega_r \pi_r [u(c_r, l_r) + v(G)] + \omega_n \pi_n [u(c_n, l_n) + v(G)]. \\
\pi_r c_r + \pi_n c_n + G &\leq A \left[mx^\rho + (1-m) \left(\frac{\pi_r l_r}{1-m} \right)^\rho \right]^{\frac{1-\alpha}{\rho}} (\pi_n l_n)^\alpha - \phi mx, \quad [\mu].
\end{aligned}$$

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

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

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

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

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

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

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,

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

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

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

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

$$\omega_r u_l(c_r, l_r) \geq \frac{\mu}{\pi_r l_r} (1 - \alpha)(1 - m)Y,$$

$$\omega_n u_l(c_n, l_n) = \mu \frac{\alpha Y}{\pi_n l_n},$$

$$g'(G) = \mu.$$

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

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

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

$$g'(G) = \omega_r u'(c_r). \quad (43)$$

A.2 Proof of Lemma 1

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

$$p_i = \phi. \quad (44)$$

Optimality for final goods producers implies that

$$x_i = \begin{cases} \frac{\pi_r l_r}{1-m}, & i \in [0, m], \\ 0, & \text{otherwise} \end{cases} \quad (45)$$

$$n_i = \begin{cases} \frac{\pi_r l_r}{1-m}, & i \in (m, 1], \\ 0, & \text{otherwise} \end{cases} \quad (46)$$

$$m = \max \left\{ 1 - \left[\frac{(1 + \tau_x)\phi}{(1 - \alpha)A} \right]^{1/\alpha} \frac{\pi_r l_r}{\pi_n l_n}, 0 \right\}, \quad (47)$$

$$Y = A \left[\int_0^m x_i^\rho di + \int_m^1 n_i^\rho di \right]^{\frac{1-\alpha}{\rho}} (\pi_n l_n)^\alpha, \quad (48)$$

$$w_r = (1 - \alpha)(1 - m) \frac{Y}{\pi_r l_r}, \quad (49)$$

$$w_n = \alpha \frac{Y}{\pi_n l_n}. \quad (50)$$

The resource constraint is

$$\pi_r c_r + \pi_n c_n + G \leq Y - \int_0^m \phi x_i, \quad (51)$$

We can let equation (44) define the price of robots, equation (45) define x_i , equations (46), (47) and (48) determine n_i , m , and Y , respectively. Assuming that m is interior, the wage equations (49) and (50) can be written as (11) and (12). 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

$$\pi_r c_r + \pi_n c_n + G \leq \pi_n w_n l_n \frac{\tau_x + \alpha}{\alpha(1 + \tau_x)} + \frac{\pi_r w_r l_r}{1 + \tau_x}. \quad (52)$$

This derivation makes it clear that the resource constraint (52) summarize the equilibrium conditions of the production side of the economy.

Household optimality requires that

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

The following incentive constraint are necessary constraints

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

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 is done by setting

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

A.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

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

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

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

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 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 by marginally increasing τ_x . Furthermore, since optimality implies that $W'(\tau_x) = 0$ then the optimal tax on robots verifies that

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

The first order condition with respect to l_r implies that

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

where $\tilde{\omega}_r = \omega_r/\mu$. Replacing this in the optimal condition for τ_x we obtain

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

A.3.1 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

$$W'(\tau_x) = -\frac{\mu}{\alpha(1 + \tau_x)} \pi_n w_n l_n \frac{\tau_x}{1 + \tau_x} \frac{1 - \alpha}{\alpha} = 0 \Leftrightarrow \tau_x = 0. \quad (53)$$

A.4 Model with simple taxes

The competitive equilibrium for this economy is characterized by the following set of equations

$$c_n = \lambda (w_n l_n)^{1-\gamma} = \lambda \left(\frac{\alpha Y}{\pi_r} \right)^{1-\gamma}, \quad (54)$$

$$\frac{-u_l(c_n, l_n) l_n}{u_c(c_n, l_n)} = \lambda (1 - \gamma) (w_n l_n)^{1-\gamma}, \quad (55)$$

$$c_r = \lambda (w_r l_r)^{1-\gamma} = \lambda \left(\frac{(1-\alpha)(1-m)Y}{\pi_r} \right)^{1-\gamma}, \quad (56)$$

$$\frac{-u_l(c_r, l_r)l_r}{u_c(c_r, l_r)} = \lambda (1-\gamma) (w_r l_r)^{1-\gamma}, \quad (57)$$

$$m = \max \left\{ 1 - \left(\frac{\phi(1+\tau_x)}{(1-\alpha)A} \right)^{1/\alpha} \frac{\pi_r l_r}{\pi_n l_n}, 0 \right\}, \quad (58)$$

$$Y = \begin{cases} A(\pi_r l_r)^{1-\alpha} (\pi_n l_n)^\alpha, & \text{if } m = 0 \\ \frac{w_n}{\alpha} \pi_n l_n, & \text{if } m > 0 \end{cases} \quad (59)$$

$$\pi_n c_n + \pi_r c_r + G \leq \begin{cases} A(\pi_r l_r)^{1-\alpha} (\pi_n l_n)^\alpha, & \text{if } m = 0 \\ w_n \pi_n l_n \frac{\tau_x + \alpha}{\alpha(1+\tau_x)} + \frac{w_r \pi_r l_r}{1+\tau_x}, & \text{if } m > 0 \end{cases}, \quad (60)$$

where w_r and w_n are given by (12) and (11), respectively.

Taking the ratio between equations (54) and (56), we can see that a necessary condition is

$$\frac{c_r}{c_n} = \left[\frac{(1-\alpha)(1-m)}{\alpha} \frac{\pi_n}{\pi_r} \right]^{1-\gamma} \Leftrightarrow c_r = c_n \left[\frac{(1-\alpha)}{\alpha} \frac{\pi_n}{\pi_r} \left(\frac{\phi(1+\tau_x)}{(1-\alpha)A} \right)^{1/\alpha} \frac{\pi_r l_r}{\pi_n l_n} \right]^{1-\gamma}. \quad (61)$$

The conditions (21), (22), (23a) and (16) are necessary and sufficient for an interior automation equilibrium in terms the allocations $\{c_r, l_r, c_n, l_n, G\}$ and the tax parameters $\{\tau_x, \gamma\}$. They are necessary because they follow from the equilibrium conditions. They are sufficient because, given a solution for $\{c_r, l_r, c_n, l_n, G\}$ and $\{\tau_x, \gamma\}$ which satisfies the constraints, the other remaining conditions can be satisfied by the choice of the remaining variables. In particular, equations (12) and (11) can be satisfied by the choice of w_n and w_r , respectively. We can set λ such that

$$\lambda = \frac{1}{(1-\gamma) (w_n l_n)^{1-\gamma}} \frac{-u_l(c_n, l_n)l_n}{u_c(c_n, l_n)},$$

which satisfies (55). This choice of λ combined with (23a) also satisfies (54). Choosing λ in this way and combined with (21) implies that (57) is satisfied. Satisfying (57) with this choice of λ also implies that (56) is satisfied. The conditions (58) and (59) are used to solve for m and Y . The condition (60) is the same as (16).

We now derive equation (24). The Ramsey planner solves the following problem

$$\max \omega_r \pi_r u(c_r, l_r) + \omega_n \pi_n u(c_n, l_n) + v(G),$$

subject to

$$\begin{aligned} \left[\frac{\eta}{c_r} \right] \quad c_r &= c_n \left[\frac{(1-\alpha) \left(\frac{\phi(1+\tau_x)}{(1-\alpha)A} \right)^{1/\alpha} \frac{l_r}{l_n}}{\alpha} \right]^{1-\gamma} \\ [\lambda_r] \quad u_c(c_r, l_r) c_r + \frac{u_l(c_r, l_r) l_r}{1-\gamma} &= 0, \\ [\lambda_n] \quad u_c(c_n, l_n) c_n + \frac{u_l(c_n, l_n) l_n}{1-\gamma} &= 0, \\ [\mu] \quad \pi_r c_r + \pi_n c_n + G &\leq w_n \pi_n l_n \frac{\tau_x + \alpha}{\alpha(1+\tau_x)} + \phi \pi_r l_r. \end{aligned}$$

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

$$\begin{aligned} 0 &= \frac{\eta}{c_r} c_n \left[\frac{(1-\alpha) \left(\frac{\phi(1+\tau_x)}{(1-\alpha)A} \right)^{1/\alpha} \frac{l_r}{l_n}}{\alpha} \right]^{1-\gamma} \frac{1-\gamma}{\alpha(1+\tau_x)} + \mu \left[\frac{dw_n}{d\tau_x} \pi_n l_n \frac{\tau_x + \alpha}{\alpha(1+\tau_x)} + w_n \pi_n l_n \frac{1-\alpha}{\alpha(1+\tau_x)^2} \right] \\ 0 &= \frac{\eta}{c_r} c_r \frac{1-\gamma}{\alpha(1+\tau_x)} + \mu \frac{w_n \pi_n l_n}{\alpha(1+\tau_x)^2} \left[\frac{d \log w_n}{d \log(1+\tau_x)} (\tau_x + \alpha) + 1 - \alpha \right] \\ 0 &= \frac{\eta}{c_r} c_r (1-\gamma) + \mu \frac{w_n \pi_n l_n (1-\alpha)}{(1+\tau_x)} \left[-\frac{\tau_x + \alpha}{\alpha} + 1 \right] \\ \eta(1-\gamma) &= \mu w_n \pi_n l_n \frac{1-\alpha}{\alpha} \frac{\tau_x}{1+\tau_x} \Leftrightarrow \frac{\tau_x}{1+\tau_x} = \frac{\alpha}{1-\alpha} \frac{\eta(1-\gamma)}{\mu w_n \pi_n l_n}. \end{aligned}$$

A.5 Simple taxes with a lump-sum transfer

The conditions are necessary as they follow from manipulations of the necessary conditions for an equilibrium. Sufficiency is established as follows. Let $\{c_r, l_r, c_n, l_n, G\}$ and $\{\tau_x, \gamma, \omega\}$ denote some allocation that satisfies the conditions

$$u_c(c_r, l_r) (c_r - \Omega) + \frac{u_l(c_r, l_r) l_r}{1-\gamma} = 0,$$

$$\begin{aligned}
u_c(c_n, l_n) (c_n - \Omega) + \frac{u_l(c_n, l_n) l_n}{1 - \gamma} &= 0, \\
c_r - \Omega &= (c_n - \Omega) \left[\frac{(1 - \alpha) \pi_n}{\alpha \pi_r} \left(\frac{\phi(1 + \tau_x)}{(1 - \alpha)A} \right)^{1/\alpha} \frac{\pi_r l_r}{\pi_n l_n} \right]^{1-\gamma}, \\
\pi_r c_r + \pi_n c_n + G &\leq w_n \pi_n l_n \frac{\tau_x + \alpha}{\alpha(1 + \tau_x)} + \frac{w_r \pi_r l_r}{1 + \tau_x}.
\end{aligned}$$

First, let us set w_n and w_r according to their definitions (11) and (12), respectively.

Now set Y , λ , Ω and m such that

$$\begin{aligned}
Y &= \frac{\pi_n w_n l_n}{\alpha}, \\
\lambda &= \frac{-u_l(c_n, l_n) l_n}{u_c(c_n, l_n) (1 - \gamma) (w_n l_n)^{1-\gamma}}, \\
m &= \max \left\{ 1 - \left[\frac{\phi(1 + \tau_x)}{A(1 - \alpha)} \right]^{1/\alpha} \frac{\pi_r l_r}{\pi_n l_n}, 0 \right\}.
\end{aligned}$$

and note that

$$Y = \frac{\pi_n w_n l_n}{\alpha} = \frac{\pi_r w_r l_r}{\alpha} \frac{\pi_n w_n l_n}{\pi_r w_r l_r} = \pi_r w_r l_r \left[\frac{1}{(1 - \alpha) \left[\frac{\phi(1 + \tau_x)}{A(1 - \alpha)} \right]^{1/\alpha} \frac{\pi_r l_r}{\pi_n l_n}} \right] = \frac{\pi_r w_r l_r}{(1 - \alpha) (1 - m)}.$$

To show that the conditions for optimality of non-routine households are satisfied, we note that

$$\begin{aligned}
u_c(c_n, l_n) (c_n - \Omega) + \frac{u_l(c_n, l_n) l_n}{1 - \gamma} &= 0 \\
c_n &= -\frac{u_l(c_n, l_n) l_n}{u_c(c_n, l_n) (1 - \gamma)} + \Omega
\end{aligned}$$

and by definition of λ we obtain

$$c_n = \lambda (w_n l_n)^{1-\gamma} + \Omega.$$

Now to show that the conditions for optimality of the routine household are satisfied, we note that the no-discrimination constraint implies that

$$c_r - \Omega = (c_n - \Omega) \frac{\left(\frac{(1-\alpha)}{\pi_r} \left(\frac{\phi(1+\tau_x)}{(1-\alpha)A} \right)^{1/\alpha} \frac{\pi_r l_r}{\pi_n l_n} \right)^{1-\gamma}}{\left(\frac{\alpha}{\pi_n} \right)^{1-\gamma}} = (c_n - \Omega) \frac{\left(\frac{(1-\alpha)(1-m)Y}{\pi_r} \right)^{1-\gamma}}{\left(\frac{\alpha Y}{\pi_n} \right)^{1-\gamma}}$$

$$c_r - \Omega = \lambda(w_n l_n)^{1-\gamma} \frac{\left(\frac{(1-\alpha)(1-m)Y}{\pi_r} \right)^{1-\gamma}}{\left(\frac{\alpha Y}{\pi_n} \right)^{1-\gamma}} = \lambda(w_r l_r)^{1-\gamma}$$

which shows that the budget constraint is satisfied. Furthermore, from the implementability constraint

$$u_c(c_r, l_r) (c_r - \Omega) + \frac{u_l(c_r, l_r) l_r}{1 - \gamma} = 0$$

$$\Leftrightarrow c_r - \Omega = -\frac{u_l(c_r, l_r) l_r}{u_c(c_r, l_r) (1 - \gamma)}$$

and using what we have found above

$$-\frac{u_l(c_r, l_r) l_r}{u_c(c_r, l_r) (1 - \gamma)} = \lambda(w_r l_r)^{1-\gamma}.$$

which shows that the budget marginal condition is also satisfied.

The problem of the government is

$$\max \omega_r \pi_r u(c_r, l_r) + \omega_n \pi_n u(c_n, l_n) + v(G),$$

subject to

$$[\lambda_j] \quad u_c(c_j, l_j) (c_r - \Omega) + \frac{u_l(c_j, l_j) l_j}{1 - \gamma} = 0,$$

$$\left[\frac{\eta}{c_r} \right] \quad c_r - \Omega = (c_n - \Omega) \frac{\left(\frac{(1-\alpha)}{\pi_r} \left(\frac{\phi(1+\tau_x)}{(1-\alpha)A} \right)^{1/\alpha} \frac{\pi_r l_r}{\pi_n l_n} \right)^{1-\gamma}}{\left(\frac{\alpha}{\pi_n} \right)^{1-\gamma}},$$

$$[\phi] \quad u(c_j, l_j) \geq u(\Omega, 0)$$

$$[\mu] \quad \pi_r c_r + \pi_n c_n + G \leq w_n \pi_n l_n \frac{\tau_x + \alpha}{\alpha(1 + \tau_x)} + \phi \pi_r l_r.$$

The first order condition with respect to τ_x is given by

$$\begin{aligned} 0 &= \frac{\eta}{c_r} (c_n - \Omega) \frac{\left(\frac{(1-\alpha)}{\pi_r} \left(\frac{\phi(1+\tau_x)}{(1-\alpha)A} \right)^{1/\alpha} \frac{\pi_r l_r}{\pi_n l_n} \right)^{1-\gamma}}{\left(\frac{\alpha}{\pi_n} \right)^{1-\gamma}} \frac{1-\gamma}{\alpha} \frac{1}{1+\tau_x} \\ &\quad + \mu \left[\frac{dw_n}{d\tau_x} \pi_n l_n \frac{\tau_x + \alpha}{\alpha(1 + \tau_x)} + w_n \pi_n l_n \frac{1-\alpha}{\alpha(1 + \tau_x)^2} \right] \\ \eta \frac{1-\gamma}{\alpha} \left(\frac{c_r - \Omega}{c_r} \right) &= \mu \frac{w_n \pi_n l_n (1-\alpha)}{(1 + \tau_x)} \tau_x \\ \frac{\tau_x}{1 + \tau_x} &= \frac{\alpha}{1-\alpha} \frac{\eta(1-\gamma)}{\mu w_n \pi_n l_n} \left(\frac{c_r - \Omega}{c_r} \right). \end{aligned}$$

A.6 Proof of Lemma 2

With the definition U_θ we can write the incentive constraint for the extensive margin choice as

$$U_\theta - \mathcal{O}_\theta \theta \geq U_{\theta'} - \mathcal{O}_{\theta'} \theta.$$

Now, we use the fact that if $\theta, \theta' \in \Theta_i$ then $\mathcal{O}_\theta = \mathcal{O}_{\theta'}$ and the two incentive constraints

$$\begin{cases} U_\theta - \mathcal{O}_\theta \theta \geq U_{\theta'} - \mathcal{O}_{\theta'} \theta \\ U_{\theta'} - \mathcal{O}_{\theta'} \theta' \geq U_\theta - \mathcal{O}_\theta \theta' \end{cases} \Leftrightarrow \begin{cases} U_\theta \geq U_{\theta'} \\ U_{\theta'} \geq U_\theta, \end{cases}$$

which necessarily implies that $U_\theta = U_{\theta'}$, and we define this as U_i for $i = r, n$.

For the next part, of the proposition we note that if we define

$$\theta^* = U_n - U_r$$

we have, by construction, that for all $\theta \leq \theta^*$

$$U_n - \theta \geq U_r,$$

and for all $\theta > \theta^*$

$$U_n - \theta < U_r.$$

A.7 Proof of Proposition 2

The optimal plan, for a fixed τ_x , solves the following optimization problem

$$\begin{aligned}
W(\tau_x) = \max & U_n \int_{-\infty}^{\theta^*} \omega(\theta) dF(\theta) + U_r \int_{\theta^*}^{\infty} \omega(\theta) dF(\theta) - \int_{-\infty}^{\theta^*} \omega(\theta) \theta dF(\theta) + v(G), \\
\text{subject to} & \\
[\eta_r(\theta)] \quad & U_n \geq u\left(c_\theta, \frac{y_\theta}{w_n}\right), \quad \forall \theta \in \Theta_r \\
[\eta_n(\theta)] \quad & U_r \geq u\left(c_{\theta'}, \frac{y_{\theta'}}{w_r}\right), \quad \forall \theta' \in \Theta_n \\
[\eta] \quad & \theta^* = U_n - U_r, \\
[f(\theta')\psi(\theta')] \quad & U_n = u\left(c_{\theta'}, \frac{y_{\theta'}}{w_n}\right), \quad \forall \theta' \in \Theta_n \\
[f(\theta)\psi(\theta)] \quad & U_r = u\left(c_\theta, \frac{y_\theta}{w_r}\right), \quad \forall \theta \in \Theta_r \\
[\mu] \quad & \int_{\Theta} c_\theta dF(\theta) + G(\theta) \leq \int_{\Theta_n} Y_\theta dF(\theta) \frac{\tau_x + \alpha}{\alpha(1 + \tau_x)} + \frac{\int_{\Theta_r} Y_\theta dF(\theta)}{1 + \tau_x}.
\end{aligned}$$

The variables inside squared parenthesis define the Lagrange multipliers. Suppose, towards a contradiction, that we have found a solution $\{\{c_\theta, y_\theta\}, U_r, U_n, \theta^*\}$ where either:

- $y(\theta) \neq y(\theta')$ for some $\theta, \theta' \in \Theta_r$,
- $y(\theta) \neq y(\theta')$ for some $\theta, \theta' \in \Theta_n$.

Take $w_i \geq w_j$, and let's analyze first the case where $y(\theta) \neq y(\theta')$ for some $\theta, \theta' \in \Theta_i$. Define

$$y_i = \inf_{\theta \in \Theta_i} y(\theta),$$

and note that by single-crossing for all $y(\theta) > y_i$ the incentive constraint condition does not bind. For that reason, the first-order conditions for such θ are given by

$$\begin{aligned}
u_c\left(c_\theta, \frac{y_\theta}{w_i}\right) \psi(\theta) &= \mu, \\
u_l\left(c_\theta, \frac{y_\theta}{w_i}\right) \frac{1}{w_i} \psi(\theta) &= \mu \Gamma_i.
\end{aligned}$$

The combination of this two conditions plus the constraint that all $\theta \in \Theta_i$ must share the same utility level yields a unique solution for c_θ and y_θ :

$$\begin{aligned}\frac{u_l\left(c_\theta, \frac{y_\theta}{w_i}\right) \frac{1}{w_i}}{u_c\left(c_\theta, \frac{y_\theta}{w_i}\right)} &= \Gamma_i, \\ u\left(c_\theta, \frac{y_\theta}{w_i}\right) &= U_n.\end{aligned}$$

If there is no $\hat{\theta}$ such that $y(\hat{\theta}) = y_i$ we have reached a contradiction. If there exists $\hat{\theta}$ such that this is true then the first-order conditions are

$$\begin{aligned}u_c\left(c_{\hat{\theta}}, \frac{y_{\hat{\theta}}}{w_i}\right) \left[\psi(\hat{\theta}) - \eta_i \frac{u_c\left(c_{\hat{\theta}}, \frac{y_{\hat{\theta}}}{w_j}\right)}{u_c\left(c_{\theta_i}, \frac{y_\theta}{w_i}\right)} \right] &= \mu \\ -u_l\left(c_{\hat{\theta}}, \frac{y_{\hat{\theta}}}{w_i}\right) \frac{1}{w_i} \left[\psi(\hat{\theta}) - \eta_i \frac{u_l\left(c_{\hat{\theta}}, \frac{y_{\hat{\theta}}}{w_j}\right) \frac{1}{w_j}}{u_l\left(c_{\hat{\theta}}, \frac{y_{\hat{\theta}}}{w_i}\right) \frac{1}{w_i}} \right] &= \mu \Gamma_i\end{aligned}$$

Note that this implies that

$$\frac{-u_l\left(c_{\hat{\theta}}, \frac{y_{\hat{\theta}}}{w_i}\right) \frac{1}{w_i}}{u_c\left(c_{\hat{\theta}}, \frac{y_{\hat{\theta}}}{w_i}\right)} = \frac{\psi(\hat{\theta}) - \eta_i \frac{u_c\left(c_{\hat{\theta}}, \frac{y_{\hat{\theta}}}{w_j}\right)}{u_c\left(c_{\theta_i}, \frac{y_\theta}{w_i}\right)}}{\psi(\hat{\theta}) - \eta_i \frac{u_l\left(c_{\hat{\theta}}, \frac{y_{\hat{\theta}}}{w_j}\right) \frac{1}{w_j}}{u_l\left(c_{\hat{\theta}}, \frac{y_{\hat{\theta}}}{w_i}\right) \frac{1}{w_i}}} \Gamma_i.$$

The single-crossing property implies that

$$\frac{-u_l\left(c_{\hat{\theta}}, \frac{y_{\hat{\theta}}}{w_i}\right) \frac{1}{w_i}}{u_c\left(c_{\hat{\theta}}, \frac{y_{\hat{\theta}}}{w_i}\right)} \geq \Gamma_i = \frac{-u_l\left(c_\theta, \frac{y_\theta}{w_i}\right) \frac{1}{w_i}}{u_c\left(c_\theta, \frac{y_\theta}{w_i}\right)},$$

Note that given our assumptions on the utility function

$$\frac{\partial \left(\frac{-u_l\left(c, \frac{y}{w_i}\right) \frac{1}{w_i}}{u_c\left(c, \frac{y}{w_i}\right)} \right)}{\partial y} \geq 0, \text{ and } \frac{\partial \left(\frac{-u_l\left(c, \frac{y}{w_i}\right) \frac{1}{w_i}}{u_c\left(c, \frac{y}{w_i}\right)} \right)}{\partial c} \geq 0.$$

Therefore, if $y_{\hat{\theta}} < y_{\theta}$ then to have a higher marginal rate of substitution it requires that $c_{\hat{\theta}} < c_{\theta}$. This is impossible since utilities must be equal, thus reaching a contradiction.

The argument for agents of occupation j such that $w_j \leq w_i$ follows the same lines, with the change of taking

$$y_j = \sup_{\theta \in \Theta_j} y(\theta).$$

A.8 Simple Taxes with Lump-Sum Transfers - The household's problem with regressivity

In this section of the appendix we discuss the problem of the household when the income tax function is regressive. Under the proposed tax function with a lump-sum transfer, a household which has a wage rate w solves the following problem

$$\max u(c, l) \text{ subject to } c \leq \lambda (wl)^{1-\gamma} + \Omega.$$

For simplicity assume that preferences are given by

$$u(c, l) = \log c - \zeta \frac{l^{1+\nu}}{1+\nu},$$

for $\zeta, \nu > 0$. The solution to this problem satisfies the following conditions:


$$c\zeta l^{\nu} = (1 - \gamma) \lambda w^{1-\gamma} l^{-\gamma}, \tag{63}$$

$$c = \lambda (wl)^{1-\gamma} + \Omega. \tag{64}$$

Note that if the tax system is regressive, $\gamma < 0$, and lump-sum transfers are positive, $\Omega > 0$, then as $l \rightarrow 0$ both the right- and left-hand sides of (63) converge to zero. As a result, a corner solution may be the optimal choice.

This case actually happens in our solutions to the simple income taxes with lump-sum transfers problem. Indeed, when the routine worker drops out of the labor force, it is optimal to set the lump-sum transfer up to a level in which the non-routine worker is exactly indifferent between the corner solution, with $c = \Omega$ and $l = 0$, and the interior

solution, with $c > \Omega$ and $l > 0$. This is easiest seen in the following figure. In this figure we plot both the budget constraint for this case, and the indifference curve for the non-routine worker with the highest associated level of utility.



A.9 Endogenous occupational choice - The first-best allocation

The first-best planner in this environment solves

$$\begin{aligned} & \max_{\Theta} \int_{\Theta} [u(c_{\theta}, l_{\theta}) - \mathcal{O}_{\theta}\theta] dF(\theta) + v(G), \quad \text{subject to} \\ [\lambda] \quad & \int_{\Theta} c_{\theta} f(\theta) d\theta + G = A \left[\int_0^m x_i^{\rho} di + \int_m^1 n_i^{\rho} di \right]^{\frac{1-\alpha}{\rho}} N_n^{\alpha} - \phi \int_0^m x_i di \\ [\lambda_n] \quad & N_n = \int_{\Theta_n} l_{\theta} f(\theta) d\theta, \quad \text{and} \quad [\lambda_r] \quad \int_m^1 n_i di = \int_{\Theta_r} l_{\theta} f(\theta) d\theta \equiv N_r. \end{aligned}$$

The first-best allocation in this economy features production efficiency. This means that if m is interior

$$\begin{aligned} x_i &= \begin{cases} \frac{N_r}{1-m}, & i \in [0, m] \\ 0, & \text{otherwise} \end{cases} \\ n_i &= \begin{cases} 0, & i \in [0, m] \\ \frac{N_r}{1-m}, & \text{otherwise} \end{cases} \\ m &= 1 - \left[\frac{\phi}{(1-\alpha)A} \right]^{1/\alpha} \frac{N_r}{N_n}. \end{aligned}$$

If two agents choose the same occupation, they have the same productivity. So, the first best chooses the same allocation in terms of hours of work and consumption. Then, define C_i and N_i the bundle given to the agents that are in occupation $i = r, n$.

Using these optimality conditions we can rewrite the optimization problem as follows

$$\begin{aligned} & \max \int_{\Theta_n} [u(c_n, l_n) - \theta] dF(\theta) + \int_{\Theta_r} u(c_r, l_r) dF(\theta) + v(G), \quad \text{subject to} \\ [\lambda] \quad & \int_{\Theta_n} c_n dF(\theta) + \int_{\Theta_r} c_r dF(\theta) + G = w_n \int_{\Theta_n} l_n dF(\theta) + w_r \int_{\Theta_r} l_r dF(\theta) \end{aligned}$$

How does the decision of whether an agent becomes routine or non routine looks like? If the planner allocates household θ to a non-routine occupation, the contribution to social welfare is

$$u(c_n, l_n) - \theta + \lambda(w_n l_n - c_n).$$

If the planner allocates the household to a routine occupation the contribution is instead

$$u(c_r, l_r) + \lambda(w_r l_r - c_r).$$

Clearly, the planner should allocate household θ to a non-routine occupation if the first is greater than the second,

$$\theta \leq \lambda[(w_n l_n - c_n) - (w_r l_r - c_r)] + u(c_n, l_n) - u(c_r, l_r) \equiv \theta^*.$$

This equation defines a threshold rule, θ^* . All households with $\theta \leq \theta^*$ should become non-routine workers and those with $\theta > \theta^*$ should become routine workers. This threshold rule balances the private costs of choosing a non-routine occupation for household θ ,

$$u(c_r, l_r) - u(c_n, l_n) - \theta,$$

with the social benefit of generating more consumption goods

$$\lambda[(w_n N_n - C_n) - (w_r N_r - C_r)].$$

We can rewrite the optimization problem for the social planner as follows

$$\max_{\{C_n, N_n, C_r, N_r, \theta^*\}} F(\theta^*)u(c_n, l_n) + [1 - F(\theta^*)]u(c_r, l_r) + g(G) - \int_{-\infty}^{\theta^*} \theta f(\theta) d\theta, \quad \text{subject to.}$$

$$[\lambda] \quad F(\theta^*)c_n + [1 - F(\theta^*)]c_r + G = w_n F(\theta^*)l_n + w_r [1 - F(\theta^*)]l_r.$$

The solution to this optimization problem satisfies the following optimality conditions

$$\begin{aligned} \frac{-u_l(c_n, l_n)}{u_c(c_n, l_n)} &= w_n, \\ \frac{-u_l(c_r, l_r)}{u_c(c_r, l_r)} &\geq w_r, \\ u_c(c_n, l_n) &= u_c(c_r, l_r), \end{aligned}$$

$$\theta^* = \lambda [(w_n l_n - c_n) - (w_r l_r - c_r)] + u(c_n, l_n) - u(c_r, l_r).$$

Using the fact that at the optimum $c_n = c_r$ and that $\lambda w_i = -u_l(c_i, l_i)$, we can solve for the threshold rule as a function of (c_n, l_n) and (c_r, l_r)

$$\theta^* = [u(c_n, l_n) - u_l(c_n, l_n) l_n] - [u(c_r, l_r) - u_l(c_r, l_r) l_r].$$

A.10 Proof of proposition 3

Let us define $\beta^t \mu_t$ the multiplier for period t resource constraint, and η_n the multiplier for the incentive constraint of non-routine households. The first-order conditions with respect to $c_{n,t}$ and $c_{r,t}$ are given by

$$\begin{aligned} \beta^t \mu_t \omega_n u'(c_{n,t}) \left(1 + \frac{\eta_n}{\pi_n}\right) &= \mu_t, \\ \beta^t \mu_t \omega_r u'(c_{r,t}) \left(1 - \frac{\eta_n}{\pi_r}\right) &= \mu_t. \end{aligned}$$

These conditions imply that

$$\frac{u'(c_{n,t})}{\beta u'(c_{n,t+1})} = \frac{u'(c_{r,t})}{\beta u'(c_{r,t+1})} = \frac{\mu_t}{\beta \mu_{t+1}}.$$

Since the first order condition with respect to X_{t+1} is

$$\begin{aligned}\mu_t \phi_t &= \beta \mu_{t+1} [F_x(t+1) + \phi_{t+1}(1 - \delta_x)] \\ &\quad + \beta \frac{\eta_n}{X_{t+1}} h' \left(\frac{F_r(t+1)}{F_n(t+1)} l_{r,t+1} \right) \frac{F_r(t+1)}{F_n(t+1)} l_{r,t+1} \varepsilon_{F_r/F_n, x}(t+1),\end{aligned}$$

then

$$\begin{aligned}\frac{u'(c_{n,t})}{\beta u'(c_{n,t+1})} &= \frac{u'(c_{r,t})}{\beta u'(c_{r,t+1})} = \frac{F_x(t+1) + \phi_{t+1}(1 - \delta_x)}{\phi_t} \\ &\quad + \frac{\pi_n \eta_n}{\mu_{t+1} \phi_t X_{t+1}} h' \left(\frac{F_r(t+1)}{F_n(t+1)} l_{r,t+1} \right) \frac{F_r(t+1)}{F_n(t+1)} l_{r,t+1} \varepsilon_{F_r/F_n, x}(t+1).\end{aligned}$$

Note that $\mu_t > 0$, $\eta_n > 0$, and $\varepsilon_{w_r/w_n, x}(t+1) < 0$. Then, as long as $l_{r,t+1} > 0$,

$$\frac{u'(c_{r,t})}{\beta u'(c_{r,t+1})} = \frac{u'(c_{n,t})}{\beta u'(c_{n,t+1})} < \frac{F_x(t+1) + \phi_{t+1}(1 - \delta_x)}{\phi_t}.$$

If, instead, routine workers supply zero hours, $l_{r,t+1} = 0$, the optimal intertemporal wedge is zero, $\tau_{j,t+1}^k = 0$.

A.11 Solving the dynamic model

The dynamic Mirrlees plan in section 3 solves the following problem:

$$\begin{aligned}\max \quad & \sum_{t=0}^{\infty} \beta^t \left[\omega_r \pi_r \left[\log(c_{r,t}) - \zeta \frac{l_{r,t}^{1+\nu}}{1+\nu} \right] + \omega_n \pi_n \left[\log(c_{n,t}) - \zeta \frac{l_{n,t}^{1+\nu}}{1+\nu} \right] + \chi \log G_t \right], \text{ subject to.} \\ & \sum_{t=0}^{\infty} \beta^t \left[\log(c_{n,t}) - \zeta \frac{l_{n,t}^{1+\nu}}{1+\nu} \right] \geq \sum_{t=0}^{\infty} \beta^t \left[\log(c_{r,t}) - \zeta \frac{\left(\frac{F_r(t)}{F_n(t)} l_{r,t} \right)^{1+\nu}}{1+\nu} \right] \\ & \pi_r c_{r,t} + \pi_n c_{n,t} + G_t + \phi_t X_{t+1} \leq z [X_t + \pi_r l_{r,t}]^{1-\alpha} (\pi_n l_{n,t})^\alpha + \phi_t (1 - \delta_X) X_t, \\ & X_0 \geq 0 \text{ given}\end{aligned}$$

where $\frac{F_r(t)}{F_n(t)} = \frac{1-\alpha}{\alpha} \left(\frac{\pi_n l_{n,t}}{X_t + \pi_r l_{r,t}} \right)$. First, let us define the level of automation as

$$m_t \equiv \frac{X_t}{X_t + \pi_r l_{r,t}}.$$

It is useful to define the following scaled variables that are constant in the steady state: $\bar{c}_{i,t} \equiv c_{i,t}/e^{\frac{1-\alpha}{\alpha}g_\phi(t-1)}$, $\bar{G}_t \equiv G_t/e^{\frac{1-\alpha}{\alpha}g_\phi(t-1)}$, and $\bar{X}_{t+1} \equiv \phi_t X_{t+1}/e^{\frac{1-\alpha}{\alpha}g_\phi(t-1)}$. Using these variables, we can rewrite the Mirrleesian planning problem as:

$$\begin{aligned} \max \quad & \sum \beta^t \left[\omega_r \pi_r \left[\log(\bar{c}_{r,t}) - \zeta \frac{l_{r,t}^{1+\nu}}{1+\nu} \right] + \omega_n \pi_n \left[\log(\bar{c}_{n,t}) - \zeta \frac{l_{n,t}^{1+\nu}}{1+\nu} \right] + \chi \log \bar{G}_t \right], \text{ subject to.} \\ & \sum \beta^t \left[\log(\bar{c}_{n,t}) - \zeta \frac{l_{r,t}^{1+\nu}}{1+\nu} \right] \geq \sum \beta^t \left[\log(\bar{c}_{r,t}) - \zeta \frac{\left(\frac{1-\alpha}{\alpha} \bar{\phi}_t \left(\frac{\pi_n l_{n,t}}{\bar{X}_t + \phi(t) \pi_r l_{r,t}} \right) l_{r,t} \right)^{1+\nu}}{1+\nu} \right], \\ & \pi_r \bar{c}_{r,t} + \pi_n \bar{c}_{n,t} + \bar{G}_t + \bar{X}_{t+1} \leq z_\phi \left(\bar{X}_t + \bar{\phi}_t \pi_r l_{r,t} \right)^{1-\alpha} (\pi_n l_{n,t})^\alpha + (1 - \bar{\delta}_X) \bar{X}_t, \\ & \bar{X}_0 \geq 0 \text{ given.} \end{aligned}$$

where $z_\phi = \frac{ze^{-\frac{(1-\alpha)^2}{\alpha}g_\phi}}{\phi^{1-\alpha}}$, $(1 - \bar{\delta}_X) = (1 - \delta_X) e^{-\frac{g_\phi}{\alpha}}$, and $\bar{\phi}_t \equiv \phi e^{-\frac{g_\phi}{\alpha}(t-1-(1-\alpha))}$.

We compute the solution to this problem using a modified version of the algorithm proposed by Slavik and Yazici (2014). This algorithm involves truncating the problem at time T and assuming that a steady state is reached at this terminal time. We set

$T = 300$. The problem then becomes

$$\begin{aligned}
\max \quad & \sum_{t=0}^{T-1} \beta^t \left[\omega_r \pi_r \left(\log(\bar{c}_{r,t}) - \zeta \frac{l_{r,t}^{1+\nu}}{1+\nu} \right) + \omega_n \pi_n \left(\log(\bar{c}_{n,t}) - \zeta \frac{l_{n,t}^{1+\nu}}{1+\nu} \right) + \chi \log \bar{G}_t \right] \\
& + \frac{\beta^T}{1-\beta} \left[\omega_r \pi_r \left(\log(\bar{c}_{r,T}) - \zeta \frac{l_{r,T}^{1+\nu}}{1+\nu} \right) + \omega_n \pi_n \left(\log(\bar{c}_{n,T}) - \zeta \frac{l_{n,T}^{1+\nu}}{1+\nu} \right) + \chi \log \bar{G}_T \right], \text{ subject to.} \\
[\eta] \quad & \sum_{t=0}^{T-1} \beta^t \left(\log(\bar{c}_{n,t}) - \zeta \frac{l_{n,t}^{1+\nu}}{1+\nu} \right) + \frac{\beta^T}{1-\beta} \left(\log(\bar{c}_{n,T}) - \zeta \frac{l_{n,T}^{1+\nu}}{1+\nu} \right) \\
\geq \quad & \sum_{t=0}^{T-1} \beta^t \left(\log(\bar{c}_{r,t}) - \zeta \frac{\left[\frac{1-\alpha}{\alpha} \bar{\phi}_t \left(\frac{\pi_n l_{n,t}}{\bar{X}_t + \bar{\phi}_t \pi_r l_{r,t}} \right) l_{r,t} \right]^{1+\nu}}{1+\nu} \right) \\
& + \frac{\beta^T}{1-\beta} \left(\log(\bar{c}_{r,T}) - \zeta \frac{\left[\frac{1-\alpha}{\alpha} \bar{\phi}_T \left(\frac{\pi_n l_{n,T}}{\bar{X}_T + \bar{\phi}_T \pi_r l_{r,T}} \right) l_{r,T} \right]^{1+\nu}}{1+\nu} \right), \\
[\beta^t \mu_t] \quad & \pi_r \bar{c}_{r,t} + \pi_n \bar{c}_{n,t} + \bar{G}_t + \bar{X}_{t+1} \leq z_\phi (\bar{X}_t + \bar{\phi}_t \pi_r l_{r,t})^{1-\alpha} (\pi_n l_{n,t})^\alpha + (1 - \bar{\delta}_X) \bar{X}_t, \text{ for } t = 0, \\
\left(\frac{\beta^T}{1-\beta} \mu_T \right) \quad & \pi_r \bar{c}_{r,T} + \pi_n \bar{c}_{n,T} + \bar{G}_T \leq z_\phi (\bar{X}_T + \bar{\phi}_T \pi_r l_{r,T})^{1-\alpha} (\pi_n l_{n,T})^\alpha - \bar{\delta}_X \bar{X}_T, \\
& \bar{X}_0 \geq 0 \text{ given,}
\end{aligned}$$

where η , $\beta^t \mu_t$, and $\frac{\beta^T}{1-\beta} \mu_T$ denote the Lagrange multipliers of each constraint.

We initialize the algorithm by computing a stationary solution for a fixed value of ϕ which is equal to its terminal value (ϕ_{300}). We have eight unknowns: $\{\bar{c}_r, \bar{c}_n, l_r, l_n, \bar{G}, \bar{X}\}$, and the multipliers η and μ . We solve for these unknowns using the following eight equations: the incentive constraint of the non-routine workers, the resource constraint and the six first-order necessary conditions.

In the second step, we maintain the same fixed cost of robots $\phi_t^* = \phi_{300}$, but gradually adjust the initial condition $\bar{X}_0^*$. In the first iteration, we set $\bar{X}_0^*$ to the solution of the previous procedure, and then in each iteration we gradually reduce that level until we reach the target initial condition $\bar{X}_0$ (which in our case is zero). We look for the dynamic paths of $\{\bar{c}_{r,t}^*, \bar{c}_{n,t}^*, l_{r,t}^*, l_{n,t}^*, \bar{G}_t^*, \bar{X}_t^*, \mu_t^*\}_{t=0}^{300}$ and η^* which solve the first order-

conditions and the constraints that hold with equality, including the initial condition. The last solution of this process is used as the starting point for the third step.

Finally, we let the cost of robots vary over time. For each iteration $j \in \{0, \dots, T\}$ we set $\phi_t^* = \phi_{T-j}$ for $t = 0, \dots, T-j-1$, and $\{\phi_t^*\}_{t=T-j}^T = \{\phi_t\}_{t=T-j}^T$ at their actual values, and once again look for a solution that satisfies the first-order necessary conditions and the constraints, including the initial condition. The final iteration of this procedure produces the paths which solve the first-order condition of our original model. However, a final adjustment needs to be made to account for the possibility that the quantity of robots is at the corner with $\bar{X}_t = 0$. To implement this step, we take the final solution of this procedure and whenever $\bar{X}_1 < 0$, it sets $\bar{X}_1 = 0$ and recomputes the optimal path. In the second iteration, if there are negative values in the previous solution we set $\bar{X}_1, \bar{X}_2 = 0$ and recompute the solution. We proceed in this way until there are no negative values of $\bar{X}_t$ in the solution. In practice, corner solutions are generally found only in the earlier periods, when the costs of automation are still high.

B Figures

Figure 1: Status-Quo Equilibrium




Figure 2: First Best




Figure 3: Mirrleesian Optimal Taxation




Figure 4: Simple Taxes - Panel A




Figure 4: Simple Taxes - Panel B




Figure 5: Simple Taxes & Lump Sum Rebate - Panel A




Figure 5: Simple Taxes & Lump Sum Rebate - Panel B




Figure 6: Welfare




Figure 7: Consumption Equivalent




Figure 8: First Best with Occupational Choice




Figure 9: Mirlees Optimal Taxation with Occupational Choice (Panel A)




Figure 9: Mirrlees Second Best with Occupational Choice (Panel B)




Figure 10: Mirrleesian Optimal Taxation - Dynamic Model

