Should Robots Be Taxed?*

Joao Guerreiro[†], Sergio Rebelo[‡], and Pedro Teles[§]

April 2018

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 raising marginal income tax rates and taxing robots. But this solution involves a substantial efficiency loss for the reduced level of inequality. A Mirrleesian optimal income tax can reduce inequality at a smaller efficiency cost, but is difficult to implement. An alternative approach is to amend the current tax system to include a lump-sum rebate. In our model, with the rebate in place, it is optimal to tax robots only when there is partial automation.

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

*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.

[†]Northwestern University.

[‡]Northwestern University, NBER and CEPR.

[§]Católica-Lisbon School of Business & Economics, Banco de Portugal and CEPR

1 Introduction

The American writer Kurt Vonnegut began his career in the public relations division of General Electric. One day, he saw a new milling machine operated by a punch-card computer outperform the company's best machinists. This experience inspired 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.

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 only when there is partial automation. These taxes increase the wages of routine workers, and decrease the ones of the non-routine, giving the government an additional instrument to reduce income inequality. Once there is full automation, it is not optimal to tax robots. Routine workers do not work, so taxing robots distorts production decisions without reducing income inequality.[2]

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

[2] These results show that the reason why it can be optimal to tax robots in our model differs from

Under the current U.S. tax system, modeled using the after-tax income function estimated by Heathcote, Storesletten and Violante (2017), full automation never occurs. 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 agents receive the same 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. This assumption, which is standard in a Ramsey (1927) setting, restricts the outcomes that can be achieved when robot taxation is optimal.

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 solution. Unfortunately, Mirrleesian tax systems are known to be complex and potentially difficult to implement in practice.

For this reason, we study the optimal policy when the tax schedule is constrained

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.

to take a simple, exogenous form. Specifically, 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 and full automation never occurs, so routine workers keep their jobs. But 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 rebates 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 rebates.

In the three best systems in terms of welfare, the first-best, Mirrleesian optimal taxes and simple income taxes with lump-sum rebates there is full automation once the costs of automation are sufficiently low. These solutions resemble 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.

One might expect that optimal robot taxation would 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 different goods can be taxed at different rates. In our context, this assumption implies 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 when there is partial automation.

We know from the work of Atkinson and Stiglitz (1976) 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 loosens the incentive compatibility constraint of non-routine workers.

We extend our model to allow agents to switch their occupations by paying a cost. In the first-best solution, agents who have a low cost of becoming non-routine workers do so. Those with a high cost become routine workers. Once the costs of automation are sufficiently low, there is full automation; agents for whom it is too costly to become non-routine lose their jobs and live from government transfers.

In the Mirrlees solution to the model with occupational choice it is optimal to use robot taxes to loosen the incentive compatibility 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.

The paper is organized as follows. In Section 2 we describe our model of automation. Section 3 describes the status-quo equilibrium, i.e. the equilibrium under the current U.S. income tax system and no robot taxes. Section 4 describes the first-best solution to the problem of an utilitarian planner. In Section 5, we analyze a Mirrleesian second-best solution to the planner's problem. In Section 6, 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 Section 7, we compare the different policies we consider both in terms of social welfare and of the utility of different agents. Section 8 discusses the model with endogenous occupation choice. Section 9 relates our results to classical results on production efficiency in the public finance literature. Section 10 concludes.

To streamline the main text, we relegate the more technical proofs to the appendix.

2 Model

In this section, we discuss a simple model of automation that allows us to address the optimal tax policy questions posed in the introduction. The model has two representative households who draw utility from consumption of private and public goods and disutility from labor. One household 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 are two representative households, one composed of non-routine workers and the other of routine workers. The index $j = n, r$, denotes non-routine and routine labor, respectively.

Household j derives utility from consumption, C_j, and from the provision of a public good, G. Each household has one unit of time and derives disutility from labor (N_j). The household's utility function is given by

$$U_j = u(C_j) - v(N_j) + g(G). \tag{1}$$

The functions $u(C_j)$ and $g(G)$ are differentiable, strictly increasing and concave. The function $v(N_j)$ is differentiable, strictly increasing and convex. In order to guarantee that the optimal choices of consumption and leisure are interior solutions, we assume that the following conditions hold: $\lim_{C_j \to 0} u'(C_j) = \infty$ and $\lim_{N_j \to 1} v'(N_j) = \infty$.

Household j chooses C_j and N_j to maximize utility (1), subject to the budget constraint

$$C_j \leq w_j N_j - T(w_j N_j),$$

[3]Autor, Levy and Murmane (2003) study the importance of tasks performed by routine workers in different industries. They then discuss the impact of automating these tasks on the demand for routine labor.

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

We now describe the problem of the firms, starting with the firms that produce robots.

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 to automate task i maximizes profits

$$\pi_i = \max_{x_i} p_i x_i - \phi x_i.$$

It follows that

$$p_i = \phi,$$

and that $\pi_i = 0$.

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 when m tasks from the continuum $[0,1]$ are automated, the automated tasks are in the interval $[0,m]$.[4]

The production function is given by

$$Y = A\left[\int_0^m x_i^\rho di + \int_m^1 n_i^\rho di\right]^{\frac{1-\alpha}{\rho}} N_n^\alpha, \tag{2}$$

with $\alpha \in (0,1)$ and $\rho \in (-\infty, 1)$.

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

The problem of the firm is to maximize profits,

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

where Y is given by (2), and 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

$$w_n = \alpha A \left[\int_0^m x_i^\rho di + \int_m^1 n_i^\rho di \right]^{\frac{1-\alpha}{\rho}} N_n^{\alpha-1}, \tag{3}$$

$$(1+\tau_x)\phi = (1-\alpha) A \left[\int_0^m x_i^\rho di + \int_m^1 n_i^\rho di \right]^{\frac{1-\alpha}{\rho}-1} N_n^\alpha x_i^{\rho-1}, \text{ for } i \in [0,m] \tag{4}$$

$$w_r = (1-\alpha) A \left[\int_0^m x_i^\rho di + \int_m^1 n_i^\rho di \right]^{\frac{1-\alpha}{\rho}-1} N_n^\alpha n_i^{\rho-1}, \text{ for } i \in (m,1] \tag{5}$$

are satisfied. 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 $w_r < (1+\tau_x)\phi$. The firm chooses to fully automate ($m = 1$) and to employ no routine workers ($n_i = 0$) if $w_r > (1+\tau_x)\phi$. If $w_r = (1+\tau_x)\phi$, 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,

$$x_i = n_l, \text{ for } i \in [0,m] \text{ and } l \in (m,1].$$

Since the technology has constant returns to scale, profits are zero, $\pi = 0$.

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

$$G \leq T(w_r N_r) + T(w_n N_n) + \int_0^m \tau_x \phi x_i di. \tag{6}$$

Equilibrium An equilibrium is a set of allocations $\{C_r, N_r, C_n, N_n, G, n_i, x_i, m\}$, prices $\{w_r, w_n, p_i\}$, and a tax system $\{T(\cdot), \tau_x\}$ such that: (i) given prices and taxes, allocations solve each households' problem; (ii) given prices and taxes, allocations solve each firm's problem; (iii) the government budget constraint is satisfied; and (iv) markets clear.

The market clearing condition for routine labor is

$$\int_m^1 n_i di = N_r. \tag{7}$$

The market-clearing condition for the output market is

$$C_r + C_n + G \leq A \left[\int_0^m x_i^\rho di + \int_m^1 n_i^\rho di \right]^{\frac{1-\alpha}{\rho}} N_n^\alpha - \int_0^m \phi x_i di. \tag{8}$$

The market clearing condition (7), and firms first-order condition (5) imply

$$N_r = (1-m) n_i, \text{ for } i \in (m, 1].$$

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

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

In such an equilibrium, using equation (4) together with the two previous conditions, we obtain

$$m = 1 - \left[\frac{(1+\tau_x)\phi}{(1-\alpha)A} \right]^{1/\alpha} \frac{N_r}{N_n}. \tag{9}$$

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

$$w_n = \alpha \frac{A^{1/\alpha} (1-\alpha)^{\frac{1-\alpha}{\alpha}}}{[(1+\tau_x)\phi]^{\frac{1-\alpha}{\alpha}}}, \tag{10}$$

$$w_r = (1+\tau_x)\phi. \tag{11}$$

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.

It is also useful to note that in any equilibrium

$$\begin{aligned} \frac{w_r N_r}{Y} &= (1-\alpha)(1-m), \\ \frac{w_n N_n}{Y} &= \alpha. \end{aligned}$$

An increase in automation reduces the income share of routine workers while keeping constant the share of non-routine workers. In this sense, an increase in automation leads to an increase in pre-tax income inequality.

3 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 empirically-plausible functional form for after-tax income proposed by Heathcote, Storesletten and Violante (2017). Using their specification, the after-tax income of household j, $y(w_j N_j)$, is given by[5]

$$y(w_j N_j) = \lambda (w_j N_j)^{1-\gamma}, \tag{12}$$

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 specification yields a good fit to the data with an R^2 of 0.91.

This formulation implies that total and average taxes paid by agent j are given by

$$\begin{aligned} T(w_j N_j) &= w_j N_j - \lambda (w_j N_j)^{1-\gamma}, \\ \frac{T(w_j N_j)}{w_j N_j} &= 1 - \lambda (w_j N_j)^{-\gamma}. \end{aligned}$$

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

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

(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:

$$u(C_j) - v(N_j) + g(G) = \log(C_j) + \zeta \log(1 - N_j) + \chi \log(G). \tag{13}$$

For these preferences, and the status-quo tax specification, both households choose to work the same hours, which only depend on ζ and the progressivity parameter, γ,

$$N_j = \frac{1-\gamma}{1-\gamma+\zeta}.$$

We set $\zeta = 1.63$, so that in the status-quo equilibrium agents choose to work one third of their time endowment. We choose $\chi = 0.25$ and assume that the government sets government spending equal to 20 percent of output and adjusts λ to satisfy its budget constraint. Given our choice for χ, this policy is optimal in the sense that it maximizes the average utility of the two workers in the economy. On the production side, we normalize A to one and choose $\alpha = 0.5$, so that the wages of routine and non-routine workers are identical in the absence of automation. These parameter choices are used in all of our numerical experiments.[6]

We vary ϕ, the cost of producing robots on the interval $\Phi = (0; (1-\alpha)A]$. The upper bound of the interval, $(1-\alpha)A$, is the lowest value of ϕ consistent with no automation in the status quo (see equation (9)).

Figure 1 describes the effect of changes in the cost of automation. As ϕ falls, the wage of routine workers fall and the wage of non-routine workers rises. Since the utility function is logarithmic and wages are the only source of income, 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.

[6]The 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.

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 is due to the fact that non-routine workers pay an increasing share of the tax revenue and their income rises faster than output as ϕ falls.

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

4 The first-best allocation

A first-best allocation solves a problem where the social planner maximizes a social-welfare function subject to the economy's resource constraints. In this problem, there are no restrictions on the ability of the planner to transfer income across agents.

We assume that the planner has an utilitarian social welfare function which assigns equal weights to the utilities of the individual agents. The planner chooses $C_r, N_r, C_n, N_n, G, m, \{x_i, n_i\}$ to maximize average utility,

$$V = \frac{1}{2}\left[u(C_r) - v(N_r) + g(G)\right] + \frac{1}{2}\left[u(C_n) - v(N_n) + g(G)\right].$$

One 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.

The planner's resource constraints are

$$C_r + C_n + G \leq Y - \phi \int_0^m x_i di, \tag{14}$$

$$Y = A\left[\int_0^m x_i^\rho di + \int_m^1 n_i^\rho di\right]^{\frac{1-\alpha}{\rho}} N_n^\alpha, \tag{15}$$

$$\int_m^1 n_i di = N_r. \tag{16}$$

We derive in the appendix the solution to this problem. The first-best solution shares three properties with the competitive equilibrium. First, routine labor is evenly allocated to the tasks that have not been automated, $n_i = N_r/(1-m)$, for $i \in (m, 1]$ and is zero otherwise. Second, intermediate goods are evenly allocated to the activities that have been automated. Third, for $m \in (0, 1)$ the level of intermediate goods used in each automated activity is the same as the amount of labor used in non-automated activities, $x_i = N_r/(1-m)$ for $i \in [0, m]$.

The expression for the level of automation that occurs in the first best is the same as in the status-quo equilibrium,

$$m = \max\left\{1 - \left[\frac{\phi}{(1-\alpha)A}\right]^{1/\alpha} \frac{N_r}{N_n}, 0\right\}, \tag{17}$$

because in both cases robot use is not taxed.

Labor supply decisions in the first best are determined by the following conditions

$$\frac{v'(N_n)}{u'(C_n)} = \frac{\alpha Y}{N_n}, \tag{18}$$

$$\frac{v'(N_r)}{u'(C_r)} \geq \frac{(1-m)(1-\alpha)Y}{N_r}, \tag{19}$$

where $\alpha Y/N_n$ is the marginal product of non-routine workers and $(1-m)(1-\alpha)Y/N_r$ is the marginal product of routine workers. These conditions hold in a competitive equilibrium with zero marginal income taxes.

It is optimal to equate the marginal utility of consumption of the two types of workers. Since utility is separable in consumption, achieving that goal requires equalizing the consumption of the two households

$$C_r = C_n.$$

Finally, the optimal provision of public goods equates their marginal utility to their opportunity cost in terms of private consumption,

$$g'(G) = \frac{1}{2}u'(C_n).$$

Equations (18), (19), together with the fact that $C_r = C_n$, imply that the worker with the highest marginal product works the most. This property implies that non-routine workers have lower utility than routine workers. As a result, this equilibrium cannot be implemented when the government observes total income but does not observe agent types. The reason is that non-routine workers would choose to act as routine workers in order to receive higher utility. In the next section, we analyze the optimal tax system taking into account the information constraints that the worker type and the labor input are not directly observed by the government.

Figure 2 illustrates the properties of the first-best. In panel A, we see that full automation occurs once ϕ falls below 0.31. The real wage rate for both types of workers are the same as in the status-quo equilibrium.[7] The consumption and utility of both agents rise as ϕ falls. Figure 2 also shows that implementing the first-best solution requires large transfers from non-routine to routine workers.

5 Mirrleesian optimal taxation

In this section, we characterize the optimal non-linear income tax in the presence of automation, when the government observes an agent's total income but does not observe the agent's type or labor supply, as in the canonical Mirrlees (1971) problem. As discussed in the introduction, we focus on the case where robot taxes are linear.

In the Mirrlees (1971) model, the productivities of the 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

[7]The reason for this property is as follows. Equations (10) and (11) 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.

decisions by taxing the use of robots to redistribute income from non-routine to routine workers to increase social welfare?

We assume that $\phi < \alpha^{\alpha}(1-\alpha)^{1-\alpha}A$, so that if $\tau_x \leq 0$ non-routine workers earn a higher wage, $w_n > w_r$ in an equilibrium with automation (see equations (10) and (11)). Note that an increase in τ_x raises the wage of routine workers and lowers the wage rate of non-routine agents.

When working with Mirrleesian-style optimal taxation problems, it is useful to express the problem in terms of the income instead of hours worked. The income earned by agent j is

$$Y_j = w_j N_j. \tag{20}$$

The planner's problem is to choose the allocations $\{Y_j, C_j\}$, G, and τ_x to maximize social welfare

$$V = \frac{1}{2}\left[u(C_r) - v\left(Y_r/w_r\right) + g(G)\right] + \frac{1}{2}\left[u(C_n) - v\left(Y_n/w_n\right) + g(G)\right], \tag{21}$$

subject to the resource constraint

$$C_r + C_n + G \leq Y_n \frac{\tau_x + \alpha}{\alpha(1+\tau_x)} + \frac{Y_r}{1+\tau_x}. \tag{22}$$

and two incentive compatibility (IC) constraints

$$u(C_n) - v\left(Y_n/w_n\right) \geq u(C_r) - v\left(Y_r/w_n\right), \tag{23}$$

$$u(C_r) - v\left(Y_r/w_r\right) \geq u(C_n) - v\left(Y_n/w_r\right), \tag{24}$$

The wages of the two types of workers are dictated by production and are given by equations (10) and (11).

Any competitive equilibrium satisfies equations (22), (23), and (24). 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, $\{C_j, Y_j\}$, must be at least as high as the

utility associated with any other bundle $\{C, Y\}$ that satisfies the budget constraint $C \leq Y - T(Y)$:

$$u(C_j) - v(Y_j/w_j) \geq u(C) + v(Y/w_j), \tag{25}$$

In particular, routine workers must prefer their bundle, $\{C_r, Y_r\}$, to the bundle of non-routine workers, $\{C_n, Y_n\}$. Similarly, non-routine workers must prefer their bundle, $\{C_n, Y_n\}$, to the bundle of routine workers, $\{C_r, Y_r\}$. These requirements correspond to the two IC constraints, (23), and (24), so these conditions are necessary.

We show in the Appendix that equation (22) is necessary by combining the first-order conditions to the firms' problems with the resource constraint, (18). In addition, we show that conditions (22), (23), and (24), are also sufficient. To see that (23) and (24) summarize the household problem note that it is possible to choose a tax function such that agents prefer the bundle $\{C_j, Y_j\}$ 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 Y different from Y_j, $j = r, n$. These results are summarized in the following proposition.

Lemma 1. *Equations* (22), (24) *and* (23) *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 compatibility 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 creates a divergence between an agents' productivity and his wage rate. We now describe some useful results for this economy, which are proved in the appendix.

Lemma 2. *The agent with a higher pre-tax wage has weakly higher pre-tax income and consumption, i.e. if $w_i > w_j$ then $Y_i \geq Y_j$ and $C_i \geq C_j$. This result implies that the incentive compatibility constraint for agent i binds.*

Lemma 3. *Let $w_i > w_j$. The incentive compatibility of agent i, holding with equality, and the monotonicity condition, $Y_i \geq Y_j$, are necessary and sufficient conditions for incentive compatibility of the optimal allocation.*

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 the level of social welfare conditional on a value of τ_x, as

$$W(\tau_x) = \max_{C_n, C_r, Y_n, Y_r} \frac{1}{2}\left[u(C_r) - v\left(Y_r/w_r\right) + g(G)\right] + \frac{1}{2}\left[u(C_n) - v\left(Y_n/w_n\right) + g(G)\right]$$

subject to

$$\begin{aligned} & u(C_i) - v\left(Y_i/w_i\right) = u(C_j) - v\left(Y_j/w_i\right), \\ & Y_i \geq Y_j, \\ & C_r + C_n + G \leq Y_n \frac{\tau_x + \alpha}{\alpha(1+\tau_x)} + \frac{Y_r}{1+\tau_x}. \end{aligned}$$

where the index i denotes the agent with the higher pre-tax wage. An optimal choice of τ_x requires that $W'(\tau_x) = 0$.

Proposition 1. *In the optimal plan, if automation is incomplete ($m < 1$ and $Y_r > 0$) and the monotonicity constraint does not bind then robot taxes are strictly positive ($\tau_x > 0$).*

To see the intuition for this result, suppose first that $\tau_x < 0$, then $w_n > w_r$. 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 by rewriting the IC for non-routine workers in terms of hours worked,

$$u(C_n) - v(N_n) \geq u(C_r) - v\left(w_r N_r / w_n\right).$$

The second part of this result establishes that production efficiency is not optimal, so it is optimal to tax robot use. The intuition for this result is as follows. Since a zero tax on robots maximizes output, a marginal increase in that tax produces only second-order output losses. On the other hand, increasing τ_x generates a first-order gain from loosening the informational restriction. Therefore, a planner that chooses $\tau_x = 0$ can always improve its objective 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 ($Y_r > 0$). When full automation is optimal ($m = 1$, $Y_r = 0$) there are no informational gains from taxing robots. Since the robot tax distorts production and does not help loosen the IC of the non-routine agent, the optimal value of τ_x is zero. We prove this result in the appendix.

Now, we turn to the study of the optimal wedges. In what follows, we assume that $Y_n \geq Y_r$ is a non-binding constraint. The optimality conditions imply the following marginal rates of substitution

$$\begin{aligned}
\frac{v'(Y_n/w_n)}{u'(C_n)} &= w_n \frac{\tau_x + \alpha}{\alpha(1+\tau_x)} \\
\frac{v'(Y_r/w_r)}{u'(C_r)} &\geq \frac{0.5 - \eta_n}{0.5 - \eta_n \frac{v'(Y_r/w_n)/w_n}{v'(Y_r/w_r)/w_r}} \frac{w_r}{1+\tau_x}, \quad (= \text{if } \ Y_r > 0),
\end{aligned}$$

where η_n denotes the Lagrange multiplier of the incentive compatibility constraint of the non-routine worker.

One result in the original Mirrlees (1971) model is that 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 the distortion created by the robot tax, which makes the wages of non-routine workers lower than their marginal productivity.

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 their marginal productivity. 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. We see that full automation occurs for values of ϕ lower than 0.26. The robot tax is positive, reaching a maximum value of 10 percent. This tax is used only when automation is incomplete. With complete automation, routine workers supply zero hours. At that point, taxing robots to raise the wages of routine workers does not help reduce income inequality. Since taxing robots distorts production and doesn't help redistribute income, it is optimal to set these taxes to zero.

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 IC constraint. Both types of workers see their consumption rise as ϕ approaches zero. This outcome is achieved through large transfers to the routine workers.

Once there is full automation, the utility of routine and non-routine workers is equalized. This property follows from the IC constraint for non-routine workers (equation (23)). This constraint is binding and when automation is full, the pre-tax income of routine workers is zero, and the right-hand side of the equation is equal to the utility of the routine worker.

6 Optimal policy with simple income taxes

The previous section describes the optimal income tax in the presence of Mirrlees-style information constraints. Despite its natural appeal, this optimal tax schedule can be complex and difficult to implement. For this reason, in this section we characterize the

optimal tax policy when the after-tax income schedule is constrained to have the simple form proposed by Heathcote, Storesletten and Violante (2017) (see equation (12)). We also assume that the utility function takes the form of equation (13).

When the government is restricted to setting income taxes consistent with the functional form (12), the competitive equilibrium in the economy can be summarized by the following conditions:

$$C_n = \lambda(\alpha Y)^{1-\gamma}, \tag{26}$$

$$N_n = \frac{1-\gamma}{1-\gamma+\zeta}, \tag{27}$$

$$C_r = \lambda[(1-\alpha)(1-m)Y]^{1-\gamma}, \tag{28}$$

$$N_r = \frac{1-\gamma}{1-\gamma+\zeta}, \tag{29}$$

$$m = \max\left\{1 - \left(\frac{\phi(1+\tau_x)}{(1-\alpha)A}\right)^{1/\alpha} \frac{N_r}{N_n}, 0\right\}, \tag{30}$$

$$Y = A\left(\frac{N_r}{1-m}\right)^{1-\alpha} N_n^{\alpha}, \tag{31}$$

$$C_n + C_r + G \leq Y - \phi\frac{m}{1-m}N_r, \tag{32}$$

$$x_i = \frac{N_n}{1-m}, \text{for } i \in [0, m], \tag{33}$$

$$n_i = \frac{N_n}{1-m}, \text{for } i \in (m, 1], \tag{34}$$

$$p_i = \phi. \tag{35}$$

Taking the ratio between equations (26) and (28), we can see that a necessary condition is

$$\frac{C_r}{C_n} = \left[\frac{(1-\alpha)(1-m)}{\alpha}\right]^{1-\gamma}. \tag{36}$$

This equation shows that there are two ways to make the ratio C_r/C_n closer to one. The first is to raise τ_x which leads to a fall in the level of automation, m. The second, is to make γ closer to one, i.e. make the tax system more progressive. Both approaches

have drawbacks. Taxing robot use reduces production efficiency and making the tax system more progressive reduces incentives to work. To see the latter effect, note that hours worked are given by equations (27) and (29). As γ approaches one, hours worked approach zero.

We can think of the planner as choosing allocations $\{C_r, C_n, G, m\}$ and progressivity γ, subject to (36) and

$$C_r + C_n + G \leq A \left(\frac{1}{1-m} \right)^{1-\alpha} \frac{1-\gamma}{1-\gamma+\zeta} - \phi \frac{m}{1-m} \frac{1-\gamma}{1-\gamma+\zeta}, \tag{37}$$

which represents the resource condition in the equilibrium definition, where the variables Y, N_n and N_r have been replaced by their equilibrium expressions. These two conditions are necessary and sufficient to describe the set of implementable allocations in terms of $\{C_r, C_n, G, m\}$ and γ. They are necessary because they follow from straightforward manipulations of equilibrium conditions. Furthermore, given any allocations $\{C_r, C_n, G, m\}$ and γ that satisfy these two constraints, we can use equation (26) to find a value for λ; such a value for λ also satisfies equation (28), since equation (36) must be satisfied. In addition, equations (27) and (29) can be used to find solutions for N_n and N_r, respectively. Given an optimal allocation for m, equation (30) can be solved by a choice of τ_x, and equation (31) yields a value for Y. These solutions, together with equation (37), imply that equation (32) is also satisfied. Finally, equation (33) can be used to solve for a value for each x_i, equation (34) for each n_i, and equation (35) for p_i.

Optimality implies the following condition

$$\mu \frac{A(1-\alpha)}{(1-m)(1-\gamma+\zeta)} \left[(1-m)^{\alpha} - \frac{\phi}{A(1-\alpha)} \right] = 0.5 - \frac{C_r}{C_r + C_n}, \tag{38}$$

where μ is the Lagrange multiplier associated with the resource constraint, (37). This expression implies that when automation is incomplete, robot use is always taxed. To see this result, it is useful to note that when $\tau_x = 0$ the equilibrium level of automation

is such that[8]

$$1 - m = \left[\frac{\phi}{(1-\alpha)A}\right]^{1/\alpha}. \tag{39}$$

Equation (39) implies that the left-hand side of equation (38) is equal to zero which is possible only when the $C_r = C_n$. Since the instruments available to the government are distortionary, it is in general not optimal to equalize the consumption of the two agents. Equation (38) implies that when $C_n > C_r$, the level of m is lower that that implied by the competitive equilibrium with $\tau_x = 0$. So, in order for equation (38) to hold, τ_x must be positive.

Figure 4 shows that the form of the tax function constrains heavily the outcomes that can be achieved. Full automation never occurs and robot taxes are used for all values of ϕ 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. The reason is the transfers from non-routine to routine workers are relatively modest.

Optimal policy with lump-sum rebates In both the first-best allocation and the Mirrlees-style allocation, routine workers drop out of the labor force once automation costs are sufficiently low. That property is absent in an equilibrium where taxes take the form proposed by Heathcote, Storesletten and Violante (2017). The reason is simple. Equation (12) implies that when before-tax income is zero, after-tax income is also zero. A worker who drops out of the labor force has zero consumption and $-\infty$ utility. In order to make the outcomes that can be achieved with a simple tax system closer to the Mirrlees-style allocation we now consider the case where the government can use a lump-sum rebate, Ω.

In this specification, the after-tax income of household j is

$$y(w_j N_j) = \lambda(w_j N_j)^{1-\gamma} + \Omega. \tag{40}$$

[8]This result reflects the fact that in equilibrium $N_n = N_r$.

To simplify the algebra, it is useful to think of the planner as choosing ω such that $\Omega = \lambda(Y\omega)^{1-\gamma}$. Consumption levels and working hours for each agent are given by

$$C_r = \lambda Y^{1-\gamma}[(1-\alpha)(1-m)+\omega]^{1-\gamma} \tag{41}$$

$$N_r = \frac{1-\gamma}{1-\gamma+\zeta+\zeta\left(\frac{\omega}{(1-\alpha)(1-m)}\right)^{1-\gamma}} \tag{42}$$

$$C_n = \lambda Y^{1-\gamma}[\alpha+\omega]^{1-\gamma} \tag{43}$$

$$N_n = \frac{1-\gamma}{1-\gamma+\zeta+\zeta(\omega/\alpha)^{1-\gamma}} \tag{44}$$

The planner chooses allocations $\{C_n, C_r, m\}$ and the instruments $\{\gamma, \omega\}$, subject to

$$C_n = \left\{\frac{\alpha+\omega}{(1-\alpha)(1-m)+\omega}\right\}^{1-\gamma} C_r, \tag{45}$$

$$C_r + C_n + G = A\left[\frac{N_r}{(1-m)}\right]^{1-\alpha} N_n^{\alpha} - \phi\frac{m}{1-m}N_r, \tag{46}$$

$$\omega \geq 0, \tag{47}$$

where N_r and N_n are given by equations (42) and (44). Equations (45), (46) and (47) are necessary and sufficient for an equilibrium.

In this economy, the lump-sum transfer plays a similar role to tax progressivity. The consumption ratio is given by,

$$\frac{C_r}{C_n} = \left[\frac{(1-\alpha)(1-m)+\omega}{\alpha+\omega}\right]^{1-\gamma},$$

It is easy to see that when $C_r/C_n < 1$, an increase in ω increases C_r/C_n, reducing consumption inequality. However, lump-sum rebates have to be financed with distortionary income taxes.

Figure 5 illustrates the properties of this allocation. In this equilibrium, income is redistributed through a large lump-sum transfer, in other words, the government guarantees a minimum income to all agents in the economy. 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. When automation is incomplete, robot taxes are used as an additional source of redistribution and τ_x can go as high as 34 percent. Complete automation occurs for values of ϕ lower than 0.19. When automation is incomplete, the income tax is regressive ($\gamma < 0$) to reduce the distortions on the labor supply of the non-routine agents.

7 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 rebates, respectively.

Figure 6 shows the utility of the social planner for values of ϕ in the interval $(0, (1-\alpha)A]$. Recall that $(1-\alpha)A$ is the lowest value of ϕ 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 rebates 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 with $\phi = (1-\alpha)A$. 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 different values of ϕ. 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. In contrast, in the status quo, routine workers become increasingly worse off as ϕ falls. With Mirrleesian optimal taxation, routine workers are made better off once ϕ becomes sufficiently low (lower than 0.34). For simple taxes with and without rebates, routine workers are better off than in the no-automation benchmark for values of ϕ lower than 0.19 and 0.09, respectively.

Figure B of Figure 7 shows that non-routine workers prefer the no-automation benchmark to the first best for high levels of ϕ (higher than 0.23). This preference reflects the large transfers that non-routine workers make to routine workers in the first best. For values of ϕ lower than 0.23, 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. For values of ϕ lower than 0.39 non-routine workers prefer the status quo to all other allocations. This preference results from a combination of high wages and relatively low taxes.

For any level of ϕ lower that $(1-\alpha)A$, routine workers rank the first-best allocation first, Mirrleesian optimal taxation second, simple taxes with rebates third, and simple taxes without rebates fourth and the status quo last. In contrast, non-routine workers rank the status quo first and the first best last. Mirrleesian optimal taxation and simple taxes with and without rebates rank in between the two extremes.

8 Endogenous occupational 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 occupational choices.

Our analysis is related to Saez (2004) and Gomes, Lozachmeur and Pavan (2017). These authors characterize Mirrlees-style optimal tax plans in models with endogenous

occupational 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 agent's 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.

Our approach borrows from both Saez (2004) and Gomes et al. (2017). As in Saez (2004), agents in our model have different preferences for the two occupations. As in Gomes et al. (2017), agents in our model can choose both their occupation and the number of hours worked. We find that production efficiency is generally not optimal in our model.

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

$$u\left(C_\theta\right) - v\left(Y_\theta / w_\theta\right) + g(G) - \mathcal{O}_\theta \theta, \tag{48}$$

where C_θ and Y_θ denote the consumption and income of household θ, respectively. The indicator $\mathcal{O}_\theta$ denotes the household's occupational 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 has a "natural" preference for routine work. All else equal, the household would prefer a routine occupation over a non-routine occupation. If instead $\theta < 0$, the

household prefers non-routine occupations.

The problem of the household is to maximize utility subject to the budget constraint

$$C_\theta \leq Y_\theta - T\left(Y_\theta\right). \tag{49}$$

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 $\Theta_r \equiv \Theta - \Theta_n$.

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

$$\int_\Theta C_\theta f(\theta)d\theta + G \leq \frac{\tau_x + \alpha}{\alpha(1+\tau_x)} \int_{\Theta_n} Y_\theta f(\theta)d\theta + \frac{\int_{\Theta_r} Y_\theta f(\theta)d\theta}{1+\tau_x}. \tag{50}$$

We assume the government moves first and chooses a tax schedule, $T(\cdot)$. Given the tax schedule, the household then chooses the occupation. Finally, the households choose hours of work, earn income, pay taxes, and consume.

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

$$u\left(C_\theta\right) - v\left(Y_\theta/w_\theta\right) \geq u\left(C_{\theta'}\right) - v\left(Y_{\theta'}/w_\theta\right), \quad \forall \theta, \theta' \in \Theta. \tag{51}$$

This *labor-supply incentive constraint* guarantees that the household chooses the assigned allocation, given the occupational choice.

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

$$u\left(C_\theta\right) - v\left(Y_\theta / w_\theta\right) - \mathcal{O}_\theta \theta \geq u\left(C_{\theta'}\right) - v\left(Y_{\theta'} / w_{\theta'}\right) - \mathcal{O}_{\theta'}\theta, \quad \forall \theta, \theta' \in \Theta. \tag{52}$$

This *occupational-choice incentive constraint* ensures that the household chooses the assigned occupation. The other conditions are that the occupational choice, $\mathcal{O}_\theta$, defines the sets Θ_n and Θ_r and that (50) holds.[9]

Proposition 2. *An incentive compatible allocation $\{C_\theta, Y_\theta, \mathcal{O}_\theta\}_{\theta\in\Theta}$ can be implemented by a non-linear income tax schedule $T(y)$ common to all agents.*

To characterize the necessary and sufficient conditions for incentive-compatible occupational choice, it is useful to define utility net of the occupational preference, θ,

$$U_\theta \equiv u\left(C_\theta\right) - v\left(Y_\theta / w_\theta\right). \tag{53}$$

Lemma 4. *An allocation is incentive compatible for occupational 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 $\theta \in \Theta_r$.

[9]These 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.

This proposition follows directly from the incentive compatibility constraints. The first part of the proposition states that all agents who share the same occupation choice must have the same 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 the best allocation.

The second part of the proposition establishes that the incentive compatibility 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.

Proposition 3. *In the optimal plan, if* $\theta, \theta' \in \Theta_i$ *then* $Y_\theta = Y_{\theta'} \equiv Y_i$ *and* $C_\theta = C_{\theta'} \equiv C_i$, *for* $i = r, n$.

Agents who choose the same occupation are essentially identical. They have the same preferences for consumption and leisure, and the same productivity. 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:

$$
\begin{aligned}
W(\tau_x) = & \left[u\left(C_n\right) - v\left(Y_n/w_n\right)\right] F(\theta^*) + \left[u\left(C_r\right) - v\left(Y_r/w_r\right)\right] \left(1 - F(\theta^*)\right) + g(G) - \int_{-\infty}^{\theta^*} \theta f(\theta) d\theta, \\
& u\left(C_n\right) - v\left(Y_n/w_n\right) \geq u\left(C_r\right) - v\left(Y_r/w_n\right), \\
& u\left(C_r\right) - v\left(Y_r/w_r\right) \geq u\left(C_n\right) - v\left(Y_n/w_r\right), \\
& \theta^* = \left[u\left(C_n\right) - v\left(Y_n/w_n\right)\right] - \left[u\left(C_r\right) - v\left(Y_r/w_r\right)\right] \\
& C_n F(\theta^*) + C_r \left(1 - F(\theta^*)\right) + G \leq Y_n F(\theta^*) \frac{\tau_x + \alpha}{\alpha(1 + \tau_x)} + \frac{Y_r \left(1 - F(\theta^*)\right)}{1 + \tau_x}.
\end{aligned}
$$

Optimizing with respect to τ_x requires $W'(\tau_x) = 0$.

The solution to this problem can be decentralized with the mechanism we describe above, in which the government moves first and chooses $\{T(\cdot), \tau_x\}$. In this decentralization, the government sets τ_x to its optimal level and sets income taxes such that

$$T(Y_n) = Y_n - C_n, \quad \text{and} \quad T(Y_r) = Y_r - C_r,$$

and $T(Y) = Y$ for $Y \neq Y_r, Y_n$.

Numerical analysis We now illustrate the properties of the occupational choice model with some numerical examples. We use the same preference and production parameters as in previous sections.[10]

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 2.

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 Mirrlees solution for the same levels of σ. In both cases there is positive taxation of robots when automation costs are sufficiently high. These taxes are higher when dispersion is also higher.

To see the intuition that underlies this result, we need to understand the mechanisms that the planner can use to improve the distribution of utility. First, the planner can

[10]In 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.

use the income tax schedule. We refer to this approach as *direct redistribution*. Second, the planner can induce more agents to become non-routine workers. We refer to this approach as *indirect redistribution*.

The optimal plan balances the costs and benefits from each form of redistribution. For $\sigma = 1$, the costs of indirect redistribution are low, so the planner induces a higher fraction of the population to become non-routine. In fact, for low automation costs the share of non-routine workers becomes higher than in the first-best allocation.

For $\sigma = 2$, 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 work decline faster when σ is higher, automation also advances more rapidly as ϕ falls. We can also see that the share of non-routine workers is lower than in the first-best.

For $\sigma = 1$, the planner engages in less direct redistribution and prefers to induce more agents to switch to non-routine occupations. Since taxes on robots are desirable only so far 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 both values of σ robot taxes reach zero before full automation is implemented. This property reflects the need to provide incentives for agents to choose non-routine occupations. Because of these incentives, the intensive margin incentive compatibility of the non-routine worker no longer binds for values of ϕ that are low but are associated with incomplete automation. For these values of ϕ it is not optimal to distort production.

9 Relation to the public finance literature

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. Another important reference is Atkinson and Stiglitz (1976). These authors argue that in an economy with Mirrleesian income taxes distorting the use of commodities is not optimal, as long as these commodities are separable from leisure in the utility function. Since uniform taxation can be interpreted as production efficiency, those results may also appear to contradict ours. In this section, we discuss the relation between these different results.

Relating our results to Diamond and Mirrlees (1971) It is central to the Diamond-Mirrlees intermediate good theorem that all net trades can be taxed at potentially different (linear) rates. In our model, this property would mean that the labor of the two types of workers can be taxed at different rates. At the heart of the failure of the Diamond and Mirrlees (1971) theorem in our model is the fact that the government cannot discriminate between the two types of workers.

The result in Diamond and Mirrlees holds if we assume that the planner can use different linear taxes for routine and non-routine workers, λ_r and λ_n. In this case, household optimality implies that

$$\frac{v'(N_j)}{u'(C_j)} = \lambda_j w_j, \quad \text{and} \quad C_j = \lambda_j w_j N_j.$$

With the ability to affect each marginal rate of substitution independently, the only constraints faced by the planner are the resource constraint

$$C_r + C_n + G \leq A \left(\frac{N_r}{1-m} \right)^{1-\alpha} N_n^{\alpha} - \phi \frac{m}{1-m} N_r,$$

and the implementability conditions for household optimality

$$u'(C_j)C_j - v'(N_j)N_j = 0, \quad \text{for} j = r, n.$$

These three conditions are necessary and sufficient for an equilibrium.

When the government can use different tax rates for each type of worker, the level of intermediate goods appears only in the resource constraint and not in the implementability condition. Since intermediate goods do not interfere with incentives, they

are 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 both types of worker to pay the same tax rate ($\lambda_r = \lambda_n$), the planner has to distort the labor supply decisions of the two types of workers in the same way, which gives rise to the following additional implementability restriction

$$\frac{v'(N_n)/u'(C_n)}{v'(N_r)/u'(C_r)} = \frac{w_n}{w_r}. \tag{54}$$

The value of τ_x no longer appears only in the resource constraint, it appears in equation (54) because the wage ratio is a function of τ_x. As a result, to relax restriction (54), it might be optimal for the planner 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 different 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. A tax on robots decreases the wage rate of non-routine workers and increase the wage rate of routine workers. This property means that it can be optimal to use robot taxes.

Relating our results to Atkinson and Stiglitz (1976) In the Mirrleesian optimal taxation problem, the planner can choose the entire income tax schedule. One might expect production efficiency to be optimal given that income taxes are very flexible. Indeed, Atkinson and Stiglitz (1976) show that, for preferences that are separable between commodities and leisure, uniform commodity taxation is optimal which implies the optimality of production efficiency. Naito (1999) shows that the Atkinson and

Stiglitz (1976) result relies on the absence of general-equilibrium effects of taxation on wages and prices and that in the presence of these effects uniform commodity taxation may not be optimal.[11] His analysis focuses on 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. A similar form of non-separability is also present in our model, which is why it can be optimal to tax robots.

The intuition for the importance of general-equilibrium effects of taxation on wages and prices is the same we used in discussing proposition 1. Because the government 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 compatibility constraint as follows:

$$u(C_i) - v(N_i) \geq u(C_j) - v\left(w_j N_j / w_i\right).$$

Crucially, this incentive compatibility 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 a Mirrlees (1971) environment. Mirrlees (1971) assumes that production is linear in labor, so taxation does not affect wages through general-equilibrium effects. Stiglitz

[11] Jacobs (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.

(1982) and Stern (1982) show that when production is not linear in labor, the optimal tax schedule is more regressive than in Mirrlees (1971) and the top marginal income tax is negative instead of zero.[12] 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 compatibility constraint easier to satisfy.

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 income tax systems that 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.

Our results and the literature on the optimal taxation of capital For simplicity, we use a static model to study the optimal taxation of robots. We model robots as intermediate goods that are close substitutes to routine labor and complements to non-routine labor. These intermediate goods could be long lived and take time to build. Would our results be different in a model in which robots are a durable intermediate good? In this setting, taxing robots creates intertemporal distortions in addition to static production inefficiency. There are other reasons why it might be optimal to introduce intertemporal distortions, which are well understood from the literature on capital taxation (see Chari, Nicolini and Teles (2016) for a recent overview). First, it might be optimal to use intertemporal distortions to confiscate the initial stock of durable goods, which are in inelastic supply. Second, intertemporal distortions can be optimal when the elasticities of the marginal utility of consumption and labor are time

[12]Rothschild 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.

varying. Third, intertemporal distortions can be optimal to provide insurance in models with idiosyncratic risk. These reasons are orthogonal to the motives for taxing robots studied in this paper.

Most economies tax capital income. Whether or not this tax distorts the accumulation of capital depends on whether investment can be expensed for tax purposes. When this expensing is allowed, as proposed in the current U.S. tax reform, capital income taxes do not distort the accumulation of capital (see Abel (2007)). In this case, when robots are a form of capital their accumulation is not distorted. In so far as taxing robots is optimal, it might desirable to limit the expensing of investment in robots.[13]

10 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 for the reduced level of inequality.

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. Unfortunately, this tax system can be complex and difficult to implement.

An alternative approach is to amend the tax system to include a rebate that is independent of income. In our model, with this rebate in place, it is optimal 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.

[13]This policy was recently implemented in South Korea. See "South Korea introduces world's first 'robot tax'," in *The Telegraph*, August 9, 2017.

Routine workers lose their jobs and live off government transfers, just like in Kurt Vonnegut's "Player Piano."

11 References

Abel, Andrew "Optimal Capital Income Taxation,: manuscript, University of Pennsylvania, 2006.

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.

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., 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 (Nov., 2003), pp. 1279-1333.

Chamley, Christophe,1986 "Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives," Econometrica 54 (3), pp. 607–622.

Chari, V. V., Juan Pablo Nicolini, and Pedro Teles, 2016, "Optimal Capital Taxation Revisited," mimeo, Federal Reserve Bank of Minneapolis.

Cortes, Guido Matias, Nir Jaimovich and Henry E. Siu "Disappearing routine jobs: Who, how, and why? " *Journal of Monetary Economics*, Volume 91, November 2017, Pages 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.

Gomes, Renato, Jean-Marie Lozachmeur and Alessandro Pavan, "Differential Taxation and Occupational Choice." *The Review of Economic Studies* (2017): rdx022.

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.

Judd, Kenneth L., 1985, "Redistributive taxation in a simple perfect foresight model," Journal of Public Economics 28 (1), 59 – 83.

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, 1 May 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.

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.

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

$$V = \max_{\{C_r,N_r,C_n,N_n\},G,m,\{x_i,n_i\}} \frac{1}{2}\left[u(C_r) - v(N_r) + g(G)\right] + \frac{1}{2}\left[u(C_n) - v(N_n) + g(G)\right].$$

$$C_r + C_n + G \leq A\left[\int_0^m x_i^\rho di + \int_m^1 n_i^\rho di\right]^{\frac{1-\alpha}{\rho}} N_n^\alpha - \int_0^m \phi x_i di, \quad [\mu],$$
$$\int_m^1 n_i di = N_r, \quad [\eta].$$

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

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

The first equation implies that the marginal productivity of routine labor should be constant across the activities that use routine labor. This property means that $n_i = N_r/(1-m)$ 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

$$V = \max_{\{C_r,N_r,C_n,N_n\},G,m,x} \sum_{j=r,n} \frac{1}{2}\left[u(C_j) - v(N_j) + g(G)\right].$$

$$C_r + C_n + G \leq A\left[mx^\rho + (1-m)\Big(\frac{N_r}{1-m}\Big)^\rho\right]^{\frac{1-\alpha}{\rho}} N_n^\alpha - \phi m x, \quad [\mu].$$

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

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

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

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

$$V = \max_{\{C_r, N_r, C_n, N_n\}, G, m, x} \sum_{j=r,n} \frac{1}{2}\left[u(C_j) - v(N_j) + g(G)\right].$$

$$C_r + C_n + G \leq A\Big(\frac{N_r}{1-m}\Big)^{1-\alpha} N_n^\alpha - \phi m \frac{N_r}{1-m}, \quad [\mu].$$

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

$$(1-\alpha)A\frac{1}{(1-m)^{2-\alpha}} N_r^{1-\alpha} N_n^\alpha - \phi \frac{N_r}{(1-m)^2} = 0 \Leftrightarrow m = 1 - \Big[\frac{\phi}{A(1-\alpha)}\Big]^{1/\alpha} \frac{N_r}{N_n},$$

provided that m is interior. Then,

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

Furthermore, the first-order conditions with respect to C_r, C_n, N_r, N_n, and G are

$$\frac{1}{2}u'(C_r) = \mu,$$

$$\frac{1}{2}u'(C_n) = \mu,$$

$$\frac{1}{2}v'(N_r) \geq \frac{\mu}{N_r}(1-\alpha)(1-m)Y,$$

$$\frac{1}{2}v'(N_n) = \mu\alpha Y,$$

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

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

$$u'(C_r) = u'(C_n) \Leftrightarrow C_r = C_n.$$

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

$$\frac{v'(N_r)}{u'(C_r)} \geq (1-\alpha)(1-m)\frac{Y}{N_r},$$
$$\frac{v'(N_n)}{u'(C_n)} = \alpha\frac{Y}{N_n}.$$

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

$$g'(G) = \frac{1}{2}u'(C_j). \tag{55}$$

A.2 Proof of Lemma 1

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

$$p_i = \phi. \tag{56}$$

Optimality for final goods producers implies that

$$x_i = \begin{cases} \frac{N_r}{1-m}, & i \in [0, m], \\ 0, & \text{otherwise} \end{cases} \tag{57}$$

$$n_i = \begin{cases} \frac{N_r}{1-m}, & i \in (m, 1], \\ 0, & \text{otherwise} \end{cases} \tag{58}$$

$$m = \max\left\{1 - \left[\frac{(1+\tau_x)\phi}{(1-\alpha)A}\right]^{1/\alpha}\frac{N_r}{N_n}, 0\right\}, \tag{59}$$

$$Y = A\left[\int_0^m x_i^\rho di + \int_m^1 n_i^\rho di\right]^{\frac{1-\alpha}{\rho}} N_n^\alpha, \tag{60}$$

$$w_r = (1-\alpha)(1-m)\frac{Y}{N_r}, \tag{61}$$

$$w_n = \alpha\frac{Y}{N_n}. \tag{62}$$

The resource constraint is

$$C_r + C_n + G = Y - \int_0^m \phi x_i, \tag{63}$$

We can let equation (56) define the price of robots, equation (57) define x_i, equations (58), (59) and (60) determine n_i, m, and Y, respectively. Assuming that m is interior, the wage equations (61) and (62) can be written as (10) and (11). These equations can be used to solve for the equilibrium wage rates. Combining the results above, we can write the resource constraint as

$$C_r + C_n + G = \alpha\frac{A^{1/\alpha}(1-\alpha)^{\frac{1-\alpha}{\alpha}}}{[(1+\tau_x)\phi]^{\frac{1-\alpha}{\alpha}}}\frac{\tau_x+\alpha}{\alpha(1+\tau_x)}N_n + \phi N_r.$$

Replacing $N_j = Y_j/w_j$, the resource constraint can be written as

$$C_r + C_n + G = Y_n\frac{\tau_x+\alpha}{\alpha(1+\tau_x)} + \frac{Y_r}{1+\tau_x}. \tag{64}$$

This derivation makes it clear that the resource constraint (64) summarize the equilibrium conditions of the production side of the economy.

Household optimality requires that

$$u(C_j) - v\left(\frac{Y_j}{w_j}\right) \geq u(C) - v\Big(\frac{Y}{w_j}\Big), \quad \forall (C,Y) : C \leq Y - T(Y).$$

The following incentive compatibility are necessary constraints

$$u(C_n) - v\Big(\frac{Y_n}{w_n}\Big) \geq u(C_r) - v\Big(\frac{Y_r}{w_n}\Big),$$
$$u(C_r) - v\Big(\frac{Y_r}{w_r}\Big) \geq u(C_n) - v\Big(\frac{Y_n}{w_r}\Big).$$

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. $T(Y) = Y$ is an example of such a mechanism.

A.3 Proof of Lemmas 2 and 3

By adding the two incentive compatibility constraints, we obtain

$$v\Big(\frac{Y_i}{w_j}\Big) - v\Big(\frac{Y_j}{w_j}\Big) \geq v\Big(\frac{Y_i}{w_i}\Big) - v\Big(\frac{Y_j}{w_i}\Big).$$

Since $v(\cdot)$ is a convex function and $w_i > w_j$, a necessary and sufficient condition is $Y_i \geq Y_j$.

Then, from the IC for i we can see that

$$u(C_i) - u(C_j) \geq v\Big(\frac{Y_i}{w_i}\Big) - v\Big(\frac{Y_j}{w_i}\Big) \geq 0,$$

which implies that $C_i \geq C_j$.

We now show by contradiction that the incentive compatibility of agent i holds with equality. Suppose not, then

$$u(C_i) - v\left(\frac{Y_i}{w_i}\right) > u(C_j) - v\left(\frac{Y_j}{w_i}\right),$$

which implies that $C_i > C_j$. Suppose then a different allocation where Y_i, Y_j, G and τ_x are kept at the same level but $C_i' = C_i - \epsilon$ and $C_j' = C_j + \epsilon$, for some small $\epsilon > 0$, which preserves the inequality

$$u(C_i') - v\left(\frac{Y_i}{w_i}\right) \geq u(C_j') - v\left(\frac{Y_j}{w_i}\right).$$

This new allocation is still resource feasible and loosens the incentive constraint for agent j.

Concavity of $u(\cdot)$ guarantees that welfare evaluated at this new allocation is strictly greater than the initial candidate solution, which contradicts the premise that the initial allocation was a solution to the problem.

We have already shown that the condition $Y_i \geq Y_j$ is necessary. It remains to be shown that it can replace the incentive constraint of agent j. We now show that since

the IC of i binds with equality, $Y_i \geq Y_j$ implies the IC of j

$$u(C_i) - u(C_j) = v\left(\frac{Y_i}{w_i}\right) - v\left(\frac{Y_j}{w_i}\right) \leq v\left(\frac{Y_i}{w_j}\right) - v\left(\frac{Y_j}{w_j}\right)$$
$$\Rightarrow u(C_i) - v\left(\frac{Y_i}{w_j}\right) \leq u(C_j) - v\left(\frac{Y_j}{w_j}\right).$$

A.4 Proof of Proposition 1

To produce a contradiction suppose not, suppose we find an allocation $\{C_n, Y_n, C_n, Y_n, G, \tau_x\}$ such that $\tau_x \leq 0$. Using the previous results, we know that $w_n > w_r$. This result implies that $Y_n \geq Y_r$, $C_n \geq C_r$, and the IC of the non-routine worker binds. These allocations solve the original optimization problem, or equivalently they solve

$$\begin{aligned} W(\tau_x) = \max_{C_r, C_n, Y_n, Y_n, G} & \sum_j \Big[u(C_j) - v\Big(\frac{Y_j}{w_j}\Big) + g(G)\Big], \quad \text{s.to.} \\ & u(C_n) - v\Big(\frac{Y_n}{w_n}\Big) = u(C_r) - v\Big(\frac{Y_r}{w_n}\Big), \quad [\eta_n], \\ & Y_n \geq Y_r, \quad [\psi], \\ & C_r + C_n + G \leq Y_n \frac{\tau_x + \alpha}{\alpha(1+\tau_x)} + \frac{Y_r}{1+\tau_x}, \quad [\mu]. \end{aligned}$$

The envelope condition is

$$\begin{aligned} W'(\tau_x) = & - v'\Big(\frac{Y_n}{w_n}\Big)(1+\eta_n)\frac{Y_n}{w_n}\frac{1-\alpha}{\alpha(1+\tau_x)} + v'\Big(\frac{Y_r}{w_r}\Big)\frac{Y_r}{w_r}\frac{1}{1+\tau_x}, \\ & + \eta_n v'\Big(\frac{Y_r}{w_n}\Big)\frac{Y_r}{w_n}\frac{1-\alpha}{\alpha(1+\tau_x)} + \frac{\mu}{1+\tau_x}\Big[Y_n\frac{1-\alpha}{\alpha(1+\tau_x)} - \frac{Y_r}{1+\tau_x}\Big]. \end{aligned}$$

The first-order conditions for optimality imply that

$$v'\Big(\frac{Y_n}{w_n}\Big)(1+\eta_n)\frac{1}{w_n} = \mu\frac{\tau_x + \alpha}{\alpha(1+\tau_x)} + \psi.$$
$$v'\Big(\frac{Y_r}{w_r}\Big)\frac{1}{w_r} = \mu\frac{1}{1+\tau_x} + \eta_n v'\Big(\frac{Y_r}{w_n}\Big)\frac{1}{w_n} - \psi.$$

Replacing these expressions in the envelope condition implies

$$W'(\tau_x) = -\mu\tau_x\frac{1-\alpha}{\alpha^2(1+\tau_x)^2} + \eta_n v'\Big(\frac{Y_r}{w_n}\Big)\frac{Y_r}{w_n}\frac{1}{\alpha(1+\tau_x)} - \psi\frac{Y_n}{\alpha(1+\tau_x)}.$$

If the monotonicity constraint doesn't bind then $\psi = 0$. In this case, as long as $\tau_x \leq 0$ then $W'(\tau_x) > 0$.

A.4.1 The full automation case ($m = 1$, $Y_r = 0$)

If the optimal plan has $Y_r = 0$ then $Y_n > 0$, which follows from the Inada conditions on utility. This result implies that $\psi = 0$. From the envelope condition we can see that

$$W'(\tau_x) = -\mu\tau_x \frac{1-\alpha}{\alpha^2(1+\tau_x)^2} = 0 \Leftrightarrow \tau_x = 0. \tag{65}$$

A.5 Proof of Proposition 2

Let $\mathcal{A} = \{C_\theta, Y_\theta, \mathcal{O}_\theta\}_{\theta\in\Theta}$ be an incentive constraint allocation. Define

$$T(Y) = \begin{cases} Y_\theta - C_\theta, & \text{if } (C_\theta, Y_\theta, \mathcal{O}_\theta) \in \mathcal{A}, \\ Y, & \text{otherwise.} \end{cases}$$

For agent θ, $\mathcal{O}_\theta$ is an optimal choice for the bundle (C_θ, Y_θ). To see this result note that since

$$u(C_{\theta'}) - v(Y_{\theta'}/w_{\theta'}) \geq u(C_\theta) - v(Y_\theta/w_{\theta'}),$$

then

$$u(C_\theta) - v(Y_\theta/w_{\theta'}) - \mathcal{O}_{\theta'}\theta \leq u(C_{\theta'}) - v(Y_{\theta'}/w_{\theta'}) - \mathcal{O}_{\theta'}\theta \leq u(C_\theta) - v(Y_\theta/w_\theta) - \mathcal{O}_\theta\theta,$$

which implies that $(C_\theta, Y_\theta, \mathcal{O}_\theta) \succcurlyeq_\theta (C_\theta, Y_\theta, \mathcal{O}_{\theta'})$ for any $\theta' \in \Theta$.

Since $\mathcal{A}$ is incentive compatible, the agent prefers $(C_\theta, Y_\theta, \mathcal{O}_\theta)$ over the consumption and income bundle $(C_{\theta'}, Y_{\theta'})$ for any other θ', with any occupational choice. If $\mathcal{O}_\theta \neq \mathcal{O}_{\theta'}$ then the constraints

$$\begin{aligned} &u(C_\theta) - v(Y_\theta/w_\theta) \geq u(C_{\theta'}) - v(Y_{\theta'}/w_\theta) \\ &u(C_\theta) - v(Y_\theta/w_\theta) - \mathcal{O}_\theta\theta \geq u(C_{\theta'}) - v(Y_{\theta'}/w_{\theta'}) - \mathcal{O}_{\theta'}\theta, \end{aligned}$$

ensure that the agent prefers his assigned bundle over the other with both occupation choices. Suppose, instead, that $\mathcal{O}_\theta = \mathcal{O}_{\theta'}$. Take some $\hat{\theta}$ such that $\mathcal{O}_{\hat{\theta}} \neq \mathcal{O}_\theta$ then

$$u(C_{\theta'}) - v(Y_{\theta'}/w_{\hat{\theta}}) - \mathcal{O}_{\hat{\theta}}\theta \leq u(C_{\hat{\theta}}) - v(Y_{\hat{\theta}}/w_{\hat{\theta}}) - \mathcal{O}_{\hat{\theta}}\theta \leq u(C_\theta) - v(Y_\theta/w_\theta) - \mathcal{O}_\theta\theta,$$

where the first inequality follows from the intensive margin constraint for $\hat{\theta}$ and the second from the extensive margin constraint for θ. Furthermore, for any other allocation $(C, Y, \mathcal{O}) \notin \mathcal{A}$, we have that $C = Y - T(Y) = 0$, which implies that such a choice is suboptimal.

This result implies that announcing a tax function common to all agents implements the allocation $\mathcal{A}$, with agents choosing their assigned bundles optimally.

A.6 Proof of Lemma 4

With the definition U_θ we can write the incentive compatibility 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 compatibility 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 3

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

$$W(\tau_x) = \max U_n F(\theta^*) + U_r\big(1 - F(\theta^*)\big) - \int_{-\infty}^{\theta^*} \theta f(\theta) d(\theta) + g(G), \quad \text{s.to.} \tag{66}$$

$$[\eta_r(\theta)] \quad U_n \geq U_r + v\Big(\frac{Y_\theta}{w_r}\Big) - v\Big(\frac{Y_\theta}{w_n}\Big), \quad \forall \theta \in \Theta_r \tag{67}$$

$$[\eta_n(\theta)] \quad U_r \geq U_n + v\Big(\frac{Y_\theta}{w_n}\Big) - v\Big(\frac{Y_\theta}{w_r}\Big), \quad \forall \theta \in \Theta_n \tag{68}$$

$$[\eta] \quad \theta^* = U_n - U_r, \tag{69}$$

$$[\psi_n(\theta)] \quad U_n = u\big(C_\theta\big) - v\big(Y_\theta / w_n\big), \quad \forall \theta \in \Theta_n \tag{70}$$

$$[\psi_r(\theta)] \quad U_r = u\big(C_\theta\big) - v\big(Y_\theta / w_r\big), \quad \forall \theta \in \Theta_r \tag{71}$$

$$[\lambda] \quad \int_\Theta C_\theta f(\theta) d\theta + G(\theta) \leq Y_n \frac{\tau_x + \alpha}{\alpha(1+\tau_x)} + \frac{Y_r}{1+\tau_x}. \tag{72}$$

$$[\lambda_n] \quad Y_n = \int_{\Theta_n} Y_\theta f(\theta) d\theta \tag{73}$$

$$[\lambda_r] \quad Y_r = \int_{\Theta_r} Y_\theta f(\theta) d\theta. \tag{74}$$

The variables inside squared parenthesis define the lagrange multipliers. Take $w_i > w_j$.

Part 1: To show that $Y_\theta = Y_j$ for all $\theta \in \Theta_j$. Suppose, towards a contradiction, that $Y_\theta \neq Y_{\theta'}$ for $\theta, \theta' \in \Theta_j$. Then, take

$$\hat{Y} \equiv \sup_{\theta \in \Theta_j} Y_\theta.$$

Note that since $\hat{Y}$ is the supremum of Y_θ then it must be that

$$U_i \geq U_j + v\Big(\frac{\hat{Y}}{w_j}\Big) - v\Big(\frac{\hat{Y}}{w_i}\Big),$$

and furthermore, since $w_j < w_i$, then

$$v\Big(\frac{\hat{Y}}{w_j}\Big) - v\Big(\frac{\hat{Y}}{w_i}\Big) > v\Big(\frac{Y_\theta}{w_j}\Big) - v\Big(\frac{Y_\theta}{w_i}\Big), \quad \forall\, Y_\theta < \hat{Y}.$$

Which implies that the incentive compatibilities for all other θ with lower income are not binding.

$$U_i > U_j + v\Big(\frac{Y_\theta}{w_j}\Big) - v\Big(\frac{Y_\theta}{w_i}\Big) \Rightarrow \eta_j(\theta) = 0.$$

The first order conditions with respect to C_θ and Y_θ are

$$\begin{aligned}
&\psi_j(\theta)u'\big(C_\theta\big) = \lambda f(\theta),\\
&\psi_j(\theta)v'\big(Y_\theta/w_j\big) = \lambda_j f(\theta), \quad \text{if } Y_\theta < \hat{Y}.
\end{aligned}$$

Suppose first that there does not exist θ such that $Y_\theta = \hat{Y}$. In this case, optimality requires that the following two conditions hold for all $\theta \in \Theta_j$

$$\begin{aligned}
&\frac{v'\big(Y_\theta/w_j\big)}{u'\big(C_\theta\big)} = \frac{\lambda_j}{\lambda},\\
&U_j = u\big(C_\theta\big) - v\big(Y_\theta/w_j\big).
\end{aligned}$$

This two conditions implicitly define a solution C_θ, Y_θ which is constant across $\theta \in \Theta_j$, thus reaching a contradiction.

Suppose instead that there exists $\hat{\theta}$ such that $Y_{\hat{\theta}} = \hat{Y}$. In this case the two conditions above must continue to hold for all other θ. For $\hat{\theta}$ it is true that

$$\begin{aligned}
&\psi_j(\hat{\theta})u'\big(C_{\hat{\theta}}\big) = \lambda f(\hat{\theta}),\\
&\psi_j(\hat{\theta})v'\big(Y_{\hat{\theta}}/w_j\big) + \eta_j(\hat{\theta})\Big[v'\big(Y_{\hat{\theta}}/w_j\big) - v'\big(Y(\hat{\theta})/w_i\big)\Big] = \lambda_j f(\hat{\theta}) \Leftrightarrow \psi_j(\hat{\theta})v'\big(Y_{\hat{\theta}}/w_j\big) < \lambda_j f(\hat{\theta}).
\end{aligned}$$

This implies that

$$\frac{v'\big(Y_{\hat{\theta}}/w_j\big)}{u'\big(C_{\hat{\theta}}\big)} < \frac{v'\big(Y_\theta/w_j\big)}{u'\big(C_\theta\big)} \Leftrightarrow \frac{v'\big(Y_{\hat{\theta}}/w_j\big)}{v'\big(Y_\theta/w_j\big)} < \frac{u'\big(C_{\hat{\theta}}\big)}{u'\big(C_\theta\big)} \Rightarrow \frac{u'\big(C_{\hat{\theta}}\big)}{u'\big(C_\theta\big)} > 1 \Rightarrow C_{\hat{\theta}} < C_\theta,$$

where the last inequality follows from the fact that $Y_{\hat{\theta}} > Y_\theta$. This is clearly a contradiction because we must have that

$$U_j = u\big(C_\theta\big) - v\big(Y_\theta/w_i\big) = u\big(C_{\hat{\theta}}\big) - v\big(Y_{\hat{\theta}}/w_i\big).$$

Part 2: To show that $Y_\theta = Y_i$ for all $\theta \in \Theta_i$. Suppose, towards a contradiction, that $Y_\theta \neq Y(\theta')$ for $\theta, \theta' \in \Theta_i$. Then, take

$$\hat{Y} \equiv \inf_{\theta \in \Theta_i} Y_\theta.$$

Note that since $\hat{Y}$ is the supremum of Y_θ then it must be that

$$U_j \geq U_i + v\Big(\frac{\hat{Y}}{w_i}\Big) - v\Big(\frac{\hat{Y}}{w_j}\Big),$$

and furthermore, since $w_j < w_i$, then

$$v\Big(\frac{\hat{Y}}{w_j}\Big) - v\Big(\frac{\hat{Y}}{w_i}\Big) < v\Big(\frac{Y_\theta}{w_j}\Big) - v\Big(\frac{Y_\theta}{w_i}\Big), \quad \forall\, Y_\theta > \hat{Y}.$$

Which implies that the incentive compatibilities for all other θ with higher income are not binding.

$$U_j > U_i + v\Big(\frac{Y_\theta}{w_i}\Big) - v\Big(\frac{Y_\theta}{w_j}\Big) \Rightarrow \eta_i(\theta) = 0.$$

The first order conditions with respect to C_θ and Y_θ are

$$\begin{aligned}
&\psi_i(\theta) u'\big(C_\theta\big) = \lambda f(\theta), \\
&\psi_i(\theta) v'\big(Y_\theta / w_i\big) = \lambda_i f(\theta), \quad \text{if } Y_\theta < \hat{Y}.
\end{aligned}$$

Suppose first that there does not exist θ such that $Y_\theta = \hat{Y}$. In this case, optimality requires that the following two conditions hold for all $\theta \in \Theta_i$

$$\begin{aligned}
&\frac{v'\big(Y_\theta / w_i\big)}{u'\big(C_\theta\big)} = \frac{\lambda_i}{\lambda}, \\
&U_i = u\big(C_\theta\big) - v\big(Y_\theta / w_i\big).
\end{aligned}$$

This two conditions implicitly define a solution C_θ, Y_θ which is constant across $\theta \in \Theta_i$, thus reaching a contradiction.

Suppose instead that there exists $\hat{\theta}$ such that $Y_{\hat{\theta}} = \hat{Y}$. In this case the two conditions above must continue to hold for all other θ. For $\hat{\theta}$ it is true that

$$\begin{aligned}
&\psi_i(\hat{\theta}) u'\big(C_{\hat{\theta}}\big) = \lambda f(\hat{\theta}), \\
&\psi_i(\hat{\theta}) v'\big(Y_{\hat{\theta}} / w_i\big) + \eta_i(\hat{\theta}) \Big[v'\big(Y_{\hat{\theta}} / w_i\big) - v'\big(Y(\hat{\theta}) / w_j\big) \Big] = \lambda_i f(\hat{\theta}) \Leftrightarrow \psi_j(\hat{\theta}) v'\big(Y_{\hat{\theta}} / w_i\big) > \lambda_i f(\hat{\theta}).
\end{aligned}$$

This implies that

$$\frac{v'\big(Y_{\hat{\theta}}/w_i\big)}{u'\big(C_{\hat{\theta}}\big)} > \frac{v'\big(Y_{\theta}/w_i\big)}{u'\big(C_{\theta}\big)} \Leftrightarrow \frac{v'\big(Y_{\hat{\theta}}/w_i\big)}{v'\big(Y_{\theta}/w_i\big)} > \frac{u'\big(C_{\hat{\theta}}\big)}{u'\big(C_{\theta}\big)} \Rightarrow \frac{u'\big(C_{\hat{\theta}}\big)}{u'\big(C_{\theta}\big)} < 1 \Rightarrow C_{\hat{\theta}} > C_{\theta},$$

where the last inequality follows from the fact that $Y_{\hat{\theta}} < Y_{\theta}$. This is clearly a contradiction because we must have that

$$U_i = u\big(C_{\theta}\big) - v\big(Y_{\theta}/w_i\big) = u\big(C_{\hat{\theta}}\big) - v\big(Y_{\hat{\theta}}/w_i\big).$$

A.8 Endogenous occupational choice - The first-best allocation

The first-best planner in this environment solves

$$\max_{\{C_\theta, N_\theta, \mathcal{O}_\theta\}} \int_{\Theta} \big[u\big(C_\theta\big) - v\big(N_\theta\big) - \mathcal{O}_\theta \theta\big] f(\theta) d\theta + g(G), \quad \text{s.to.}$$

$$[\lambda] \quad \int_{\Theta} C_\theta f(\theta) d\theta + G = A\left(\int_0^m x_i^\rho di + \int_m^1 n_i^\rho di\right)^{\frac{1-\alpha}{\rho}} N_n^\alpha - \phi \int_0^m x_i di$$

$$[\lambda_n] \quad N_n = \int_{\Theta_n} N_\theta f(\theta) d\theta, \quad \text{and} \quad [\lambda_r] \quad \int_m^1 n_i di = \int_{\Theta_r} N_\theta f(\theta) d\theta \equiv N_r.$$

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

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

$$n_i = \begin{cases} 0, & i \in [0, m] \\ \frac{N_r}{1-m}, & \text{otherwise} \end{cases}$$

$$m = 1 - \Big[\frac{\phi}{(1-\alpha)A}\Big]^{1/\alpha} \frac{N_r}{N_n}.$$

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

$$\max_{\{C_n,N_n,C_r,N_r,\mathcal{O}_\theta\}} \int_{\Theta_n} \big[u\big(C_n\big)-v\big(N_n\big)-\theta\big] f(\theta)d\theta+\int_{\Theta_r} \big[u\big(C_r\big)-v\big(N_r\big)\big] f(\theta)d\theta+g(G), \quad \text{s.to.}$$

$$[\lambda] \quad \int_{\Theta_n} C_n f(\theta)d\theta + \int_{\Theta_r} C_r f(\theta)d\theta + G = w_n \int_{\Theta_n} N_n f(\theta)d\theta + w_r \int_{\Theta_r} N_r f(\theta)d\theta$$

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

$$u\big(C_n\big) - v\big(N_n\big) - \theta + \lambda\big(w_n N_n - C_n\big).$$

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

$$u\big(C_r\big) - v\big(N_r\big) + \lambda\big(w_r N_r - C_r\big).$$

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

$$\theta \leq \lambda\Big[(w_n N_n - C_n) - (w_r N_r - C_r)\Big] + \Big[\big(u\big(C_n\big) - v\big(N_n\big)\big) - \big(u\big(C_r\big) - v\big(N_r\big)\big)\Big] \equiv \theta^*.$$

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,

$$u\big(C_r\big) - v\big(N_r\big) - \Big[u\big(C_n\big) - v\big(N_n\big) - \theta\Big],$$

with the social benefit of generating more consumption goods

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

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

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

$$[\lambda] \quad C_n F(\theta^*) + C_r(1 - F(\theta^*)) + G = w_n N_n F(\theta^*) + w_r N_r(1 - F(\theta^*)).$$

The solution to this optimization problem satisfies the following optimality conditions

$$\begin{aligned}
\frac{v'\big(N_n\big)}{u'\big(C_n\big)} &= w_n, \\
\frac{v'\big(N_r\big)}{u'\big(C_r\big)} &\geq w_r, \\
u'\big(C_n\big) = u'\big(C_r\big) &\Leftrightarrow C_r = C_n,
\end{aligned}$$

$$\theta^* = \lambda\Big[(w_n N_n - C_n) - (w_r N_r - C_r)\Big] + \Big[\big(u\big(C_n\big) - v\big(N_n\big)\big) - \big(u\big(C_r\big) - v\big(N_r\big)\big)\Big].$$

Using the fact that at the optimum $C_n = C_r$ and that $\lambda w_i = v'(N_i)$, we can solve for the threshold rule as a function of N_n and N_r

$$\theta^* = \Big[v'\big(N_n\big)N_n - v\big(N_n\big)\Big] - \Big[v'\big(N_r\big)N_r - v(N_r)\Big].$$

When $w_n > w_r$, the optimal plan features $N_n > N_r$ which implies that $\theta^* > 0$.

B Figures

Figure 1: Status-Quo Equilibrium

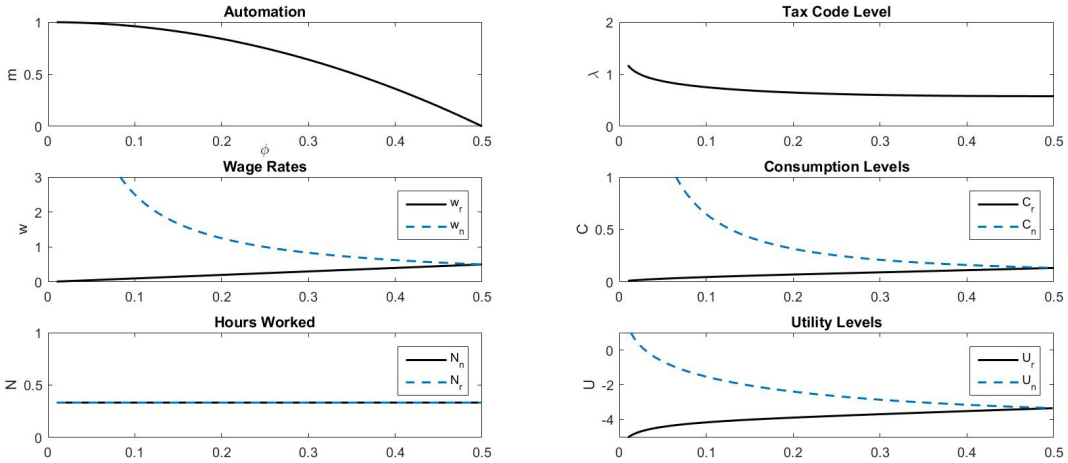


Figure 2: First Best

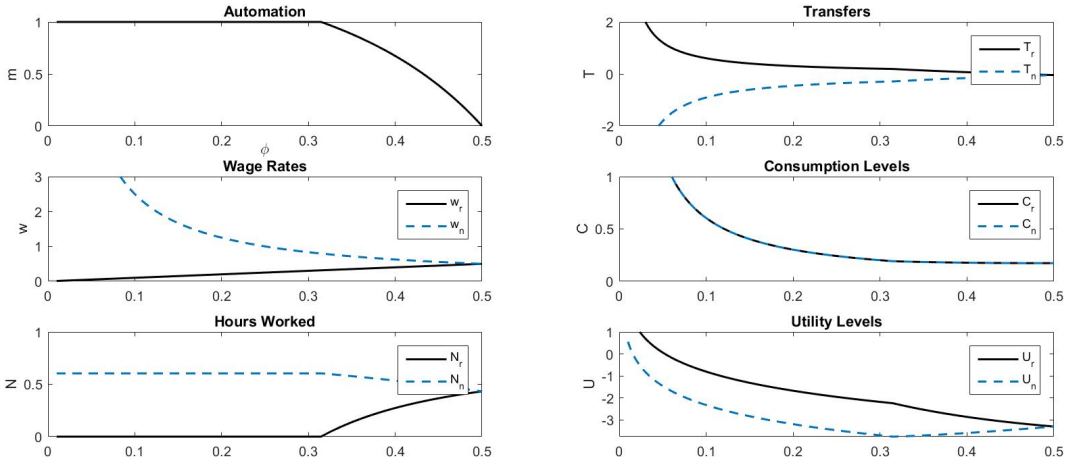


Figure 3: Mirrleesian Optimal Taxation

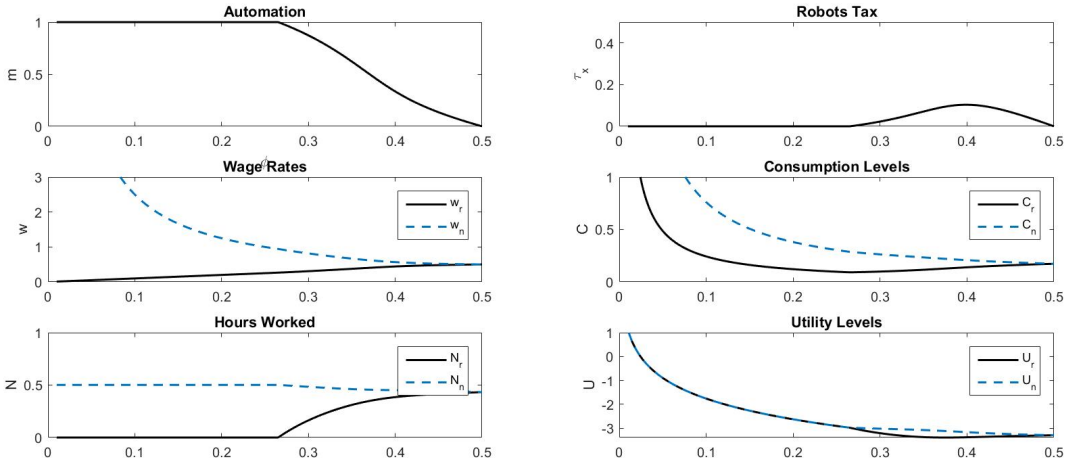


Figure 4: Simple Taxes - Panel A

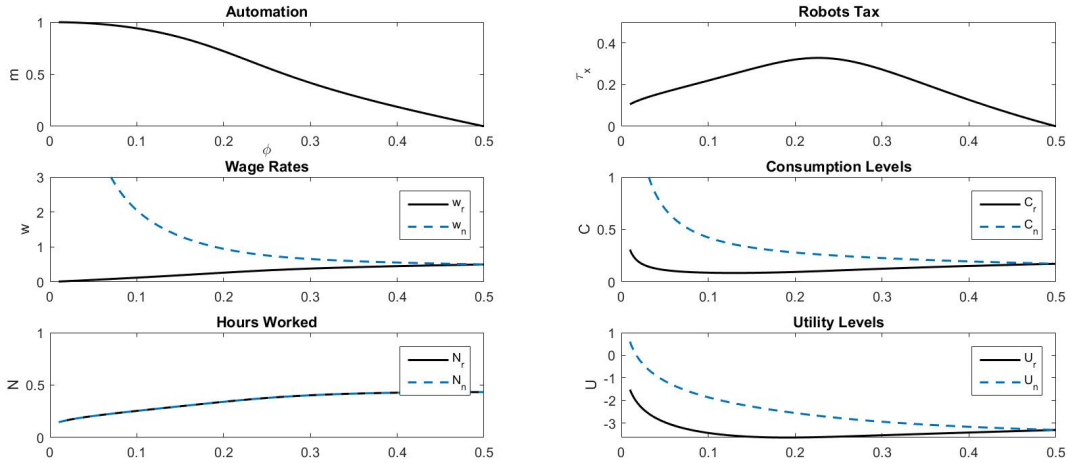


Figure 4: Simple Taxes - Panel B

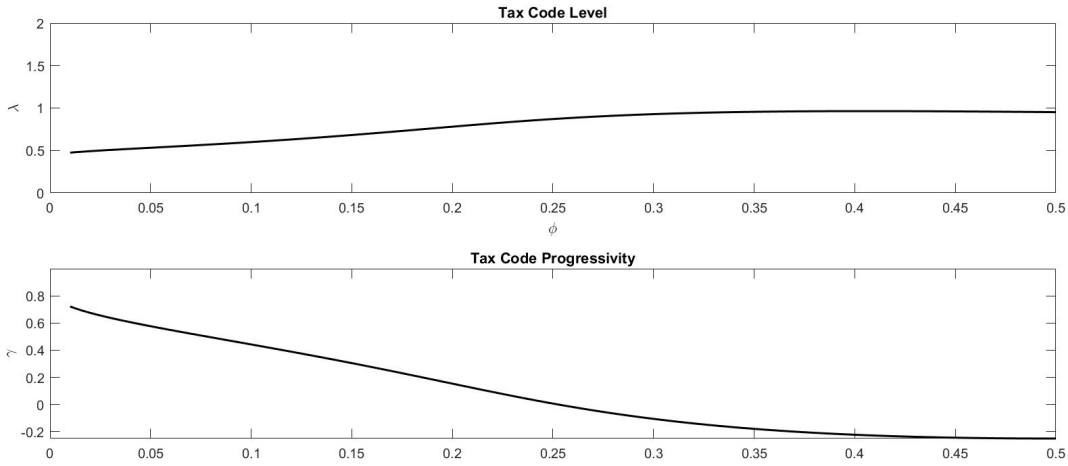


Figure 5: Simple Taxes & Lump Sum Rebate - Panel A

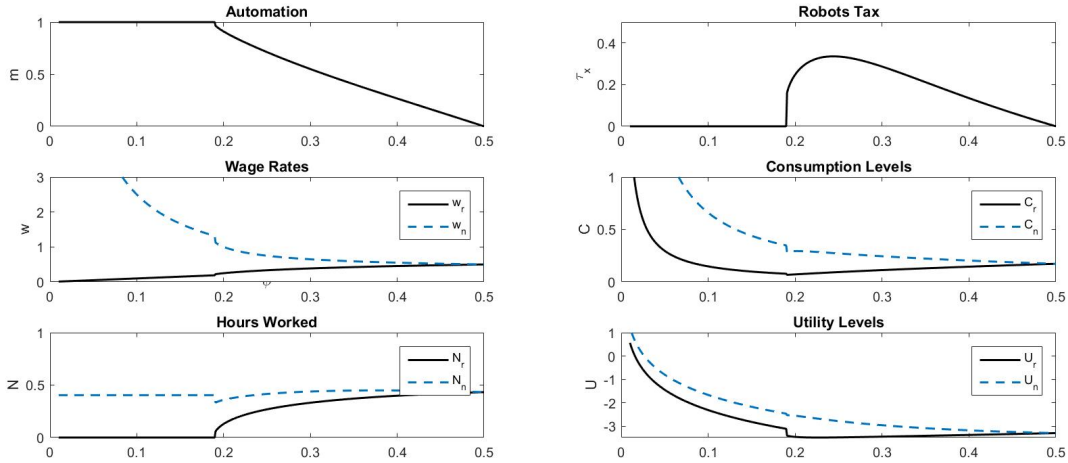


Figure 5: Simple Taxes & Lump Sum Rebate - Panel B

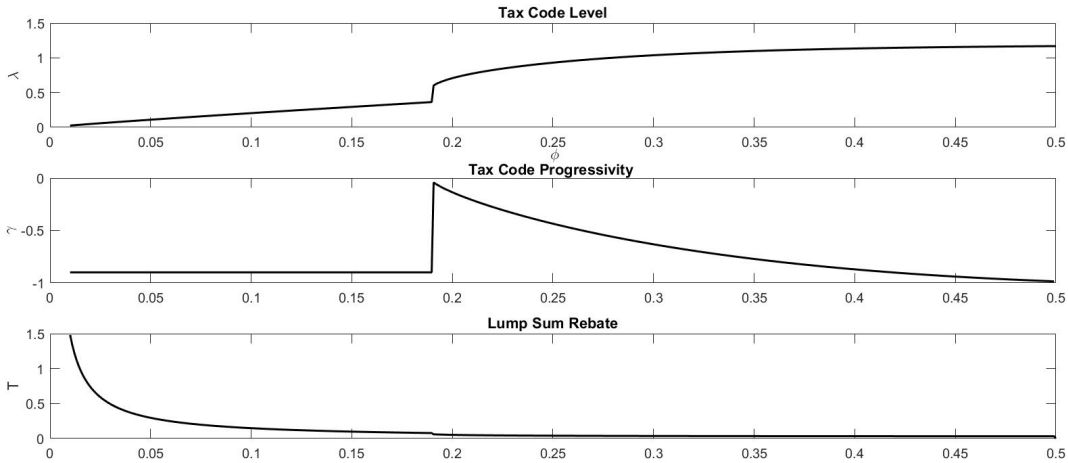


Figure 6: Welfare

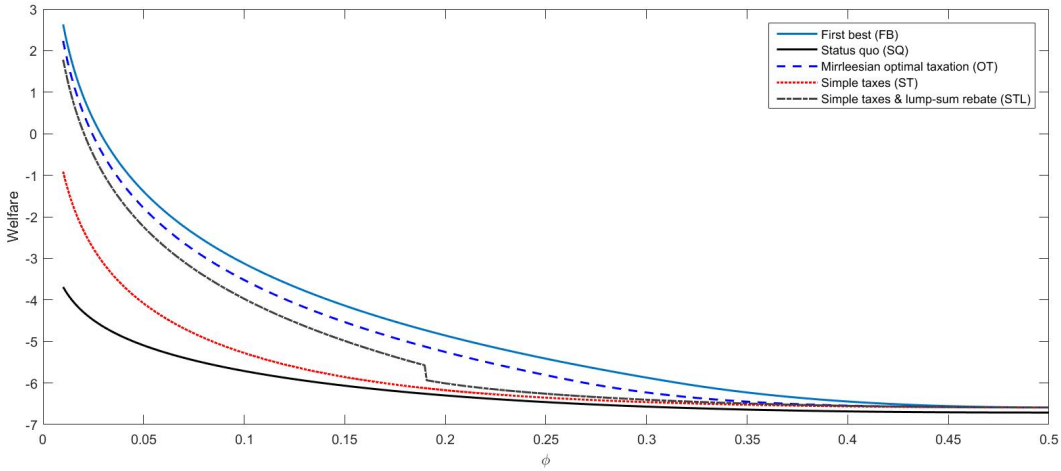


Figure 7: Consumption Equivalent - Panel A

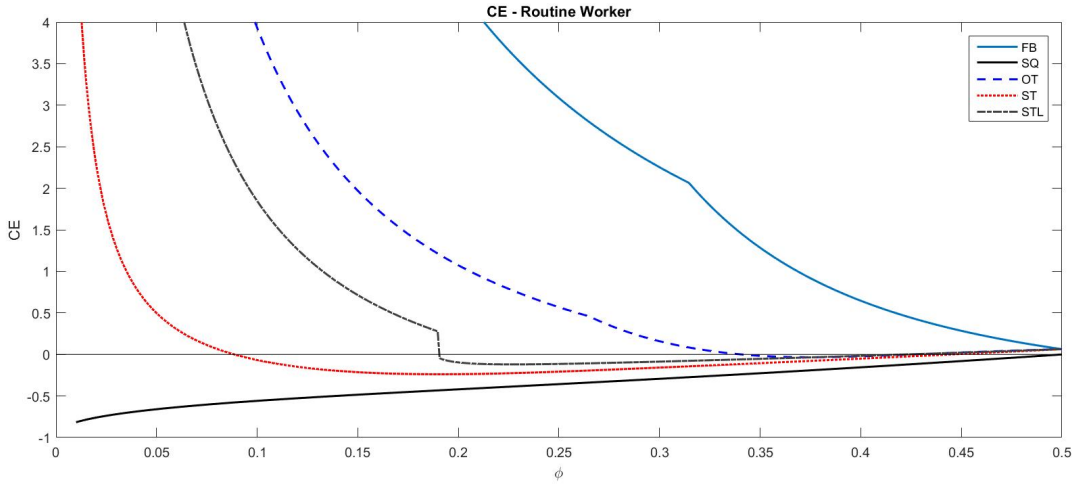


Figure 7: Consumption Equivalent - Panel B

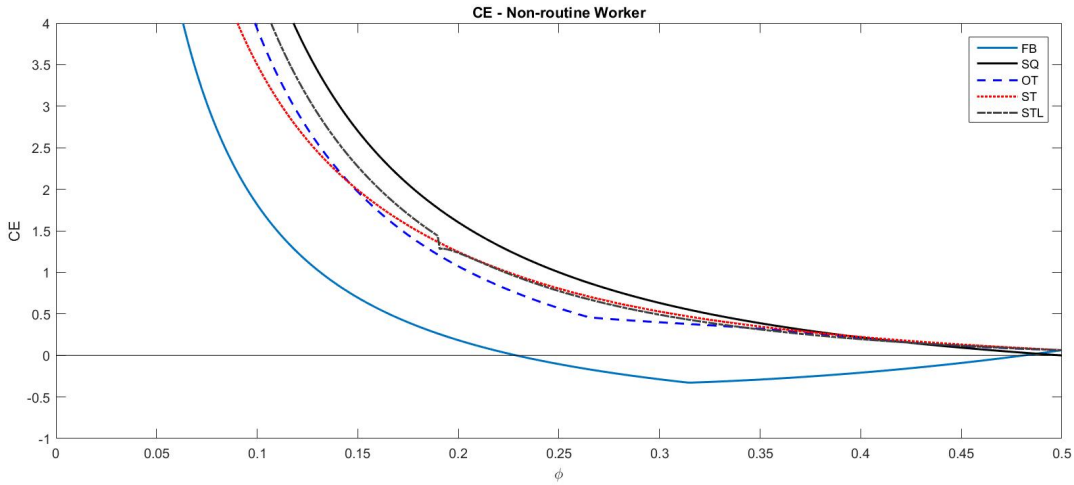


Figure 8: First Best with Occupational Choice

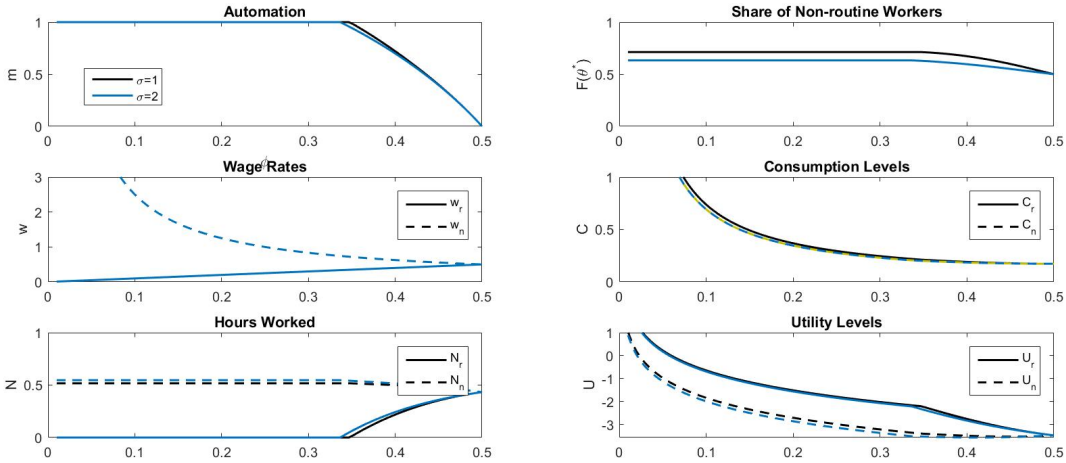


Figure 9: Mirrlees Second Best with Occupational Choice (Panel A)

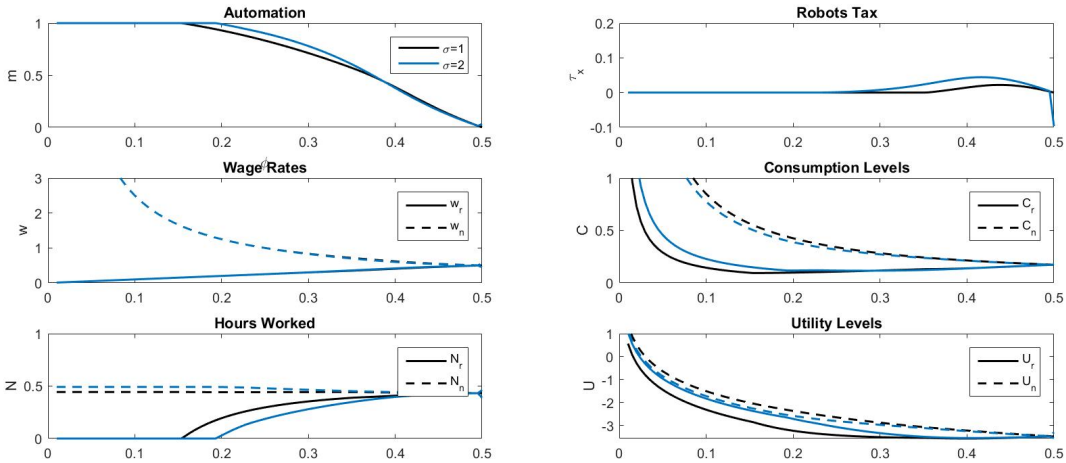


Figure 9: Mirrlees Second Best with Occupational Choice (Panel B)

