Automation and Optimal Taxation: A Task-Based Theory

Henrik Kleven & Owen Zidar*

Princeton University and NBER

September 2026

Abstract

We characterize optimal labor and capital income taxation in a task-based model of automation. Workers and machines are perfect substitutes in automatable tasks, which run from the bottom of the skill distribution up to a threshold. Workers supply labor on the extensive margin. Capitalists supply machines at a finite elasticity. In this general-equilibrium automation economy, optimal tax formulas take canonical partial-equilibrium forms. Yet automation substantially changes optimal tax rates because wages and capital returns are endogenous. We calibrate the model to the US wage distribution and automation exposure. In equilibrium, the middle class is the most automated, with machine intensity peaking around the 40th wage percentile. Relative to a no-automation benchmark, the optimal labor tax schedule is more progressive: a larger EITC subsidy at the bottom, lower taxes in the middle, and higher taxes at the top. The optimal capital tax is sizable but unaffected by automation.

*We thank Ben Moll, Florian Scheuer, and Damián Vergara for helpful comments, and Niklas Hein for outstanding research assistance.

1 Introduction

The rise of automation has important implications for labor markets and inequality. Acemoglu and Restrepo (2020) estimate that, in US local labor markets, one additional robot per thousand workers reduces the employment-to-population ratio by 0.2 percentage points and wages by 0.42 percent. Acemoglu and Restrepo (2022) find that task displacement can account for 50–70 percent of the changes in the US wage structure between 1980 and 2016, driven by relative wage declines for workers specializing in routine tasks. Muro *et al.* (2019) document that roughly a quarter of US employment is highly exposed to automation—and another third moderately exposed—with the exposure concentrated in low-wage occupations. Recent innovations in artificial intelligence (AI) have heightened these concerns.[1]

How should the tax-transfer system respond? The policy debate has focused largely on taxing the machines themselves. Bill Gates has proposed a robot tax, the European Parliament has debated one, and South Korea has scaled back its tax incentives for automation investments. Yet the literature evaluating these proposals finds that optimal machine taxes are small or transitional, and that they produce only modest welfare gains (Guerreiro *et al.* 2022; Thuemmel 2023; Costinot and Werning 2023). In this literature, income taxation serves as a flexible background instrument rather than the central object of study. Our paper focuses directly on the implications of automation for optimal labor and capital income taxation in a general-equilibrium economy. We find that automation calls for more progressive labor taxation—larger subsidies at the bottom combined with higher taxes at the top—along with a sizable tax on capital income.

Because the policy debate about automation centers on job displacement and employment losses, we adopt an extensive margin model of labor supply where workers in different occupations choose between working and not working (Saez 2002). We combine this worker-side specification with a task-based production technology in the tradition of Zeira (1998) and Moll (2025). Output is produced from a large number of tasks, ordered by the skill of the workers who perform them. Each task corresponds to an occupation. Tasks are automatable from the bottom of the skill distribution up to a threshold: their services can be produced by either workers

[1]About half of US workers report being worried about the effects of AI on their jobs (Acemoglu *et al.* 2026), and a large majority of US adults believe that AI will lead to fewer jobs (Autor 2024).

or machines, which are perfect substitutes.[2] Tasks above the threshold can be produced only by workers. The threshold can take any value, potentially allowing for middle- and high-skill tasks to be automated in equilibrium. We abstract from the possibility that workers retrain and switch between automatable and non-automatable occupations.[3] The machines used in automatable tasks are supplied by a class of capitalists, whose capital supply responds to the net-of-tax return with a finite elasticity. The government levies a nonlinear tax on labor income and a linear tax on capital income. Because wages are distinct across occupations, the labor tax corresponds to an occupation-specific tax.

Wage determination differs across tasks. In automatable tasks, firm indifference between workers and machines ties the wage to the cost of performing the task with machines: the rental rate of capital per efficiency unit. Because automation is adopted only where it lowers costs, machines depress wages in automated tasks. In non-automatable tasks, wages clear occupation-specific labor markets without any direct competition from machines. The framework distinguishes two forms of technological progress: displacement, whereby more tasks become automatable, and capital-augmenting progress, whereby machines become more productive.

In our setting, tax changes affect the wage rate in every occupation and the rental rate for capital, so their welfare effects depend on equilibrium adjustments across the entire economy. To characterize the optimal tax system, we adopt a *compensated* perturbation approach for general-equilibrium settings, building on the standard partial-equilibrium perturbation approach (Saez 2001, 2002).[4] Specifically, we perturb taxes on workers in one occupation and make compensating adjustments that switch off general-equilibrium spillovers by construction: every other occupation's tax liability moves one-for-one with the wage change induced by the reform, while the capital tax moves with the rental rate to hold the net-of-tax return fixed. The capital-tax perturbation mirrors this construction. Because factor-price movements only reshuffle a fixed aggregate income across workers and capitalists, the compensations redirect the reform's entire incidence into the government budget, leaving a partial-equilibrium trade-off in the reformed market alone. Our approach delivers the full optimal tax system in a few simple steps, despite the rich general-equilibrium environment.

[2] The within-task perfect substitution between workers and machines is the standard technology of the task-based macro literature (Zeira 1998; Acemoglu and Autor 2011; Acemoglu and Restrepo 2018; Moll *et al.* 2022; Hémous and Olsen 2022; Jones and Tonetti 2026).

[3] As empirical support for this assumption, the occupation-level estimates of Acemoglu and Restrepo (2020) show little evidence of offsetting employment gains in any occupation in response to robot adoption, suggesting that workers displaced by robots exit employment rather than switch occupations.

[4] Jacquet and Lehmann (2025) and Vergara (2026) adopt related compensated perturbation approaches.

The resulting tax formulas correspond to canonical partial-equilibrium formulas. In every task—automatable and non-automatable alike—the optimal participation tax rate satisfies the formula of Saez (2002), depending only on the workers' welfare weight and the participation elasticity. Similarly, the optimal capital tax rate depends only on the capitalists' welfare weight and the capital supply elasticity. Own-wage incidence, cross-occupation wage spillovers, and rental-rate adjustments leave no imprint on the formulas. This neutrality reflects the targeting logic of Diamond and Mirrlees (1971). Because each occupation and the capitalists face a dedicated tax instrument, every agent's net income is under the government's direct control. Whatever redistribution equilibrium prices perform, the tax system could have performed on its own.

Importantly, the fact that the *reduced-form* tax formulas maintain their familiar structure does not imply that optimal tax rates are unaffected by automation. The endogeneity of wages and the capital returns feeds into the terms of the optimal tax formulas and has substantial policy implications. We quantify these implications in simulations calibrated to the US wage distribution and automation exposure, setting labor supply and capital supply elasticities based on the empirical literature. Automation rotates the optimal labor tax schedule relative to the partial-equilibrium benchmark of Saez (2002), which is a fixed-wage, labor-only economy built from the same primitives and solved under the same formula. At the bottom, where machines depress wages and raise welfare weights, the optimal work subsidy—an Earned Income Tax Credit (EITC)—is far more generous than in the benchmark. It reaches 77 percent of earnings at the very bottom (vs 36 percent in the benchmark), and the subsidy region spans the bottom quarter of the wage distribution (vs the bottom 7 percent in the benchmark). Above the automation threshold, optimal participation taxes are around 40 percent and exceed the benchmark through most of the upper half. The optimum is thus substantially more progressive than its partial-equilibrium counterpart.

Our simulations also reveal who is hurt the most by automation in equilibrium. Machine intensity tends to peak around the 40th percentile of the wage distribution and remains substantial everywhere except at the very bottom and the very top. Consequently, there is a hollowing out of the middle class, who experience the largest losses in jobs and pay.

Technological progress amplifies these effects. Broader automation—a greater share of tasks being automatable—sharply lowers optimal participation taxes on newly affected workers, with minimal effects on taxes outside the affected region. Deeper automation—cheaper capital within automatable tasks—raises the optimal subsidy rate at the bottom and widens the region over which it applies. Meanwhile, assuming that capitalists are high-income earners and therefore have a zero welfare weight at the margin, the capital income tax rate is unaffected by automation. It depends only on the elasticity of capital supply. At a long-run elasticity of one—consistent with the estimates in Jakobsen *et al.* (2020)—the optimal rate is 50 percent.

Our paper clarifies how automation does and does not matter for optimal redistribution. Its general-equilibrium effects do not change optimal tax rules, but these effects nevertheless change optimal tax rates substantially. Our contribution speaks to three literatures. The first is the theory of optimal labor income taxation. The canonical Mirrleesian framework assumes exogenous wages (Mirrlees 1971; Diamond 1998; Saez 2001; see Piketty and Saez 2013 for a survey), and the same assumption underlies its extensive margin counterpart (Saez 2002; Kleven *et al.* 2009).

A separate branch endogenizes wages and shows that general-equilibrium incidence typically calls for lower top tax rates through trickle-down (Stiglitz 1982; Rothschild and Scheuer 2013; Sachs *et al.* 2020). Our finding that partial-equilibrium formulas hold exactly in a general-equilibrium automation economy runs against these results. The reason for the discrepancy is that this literature models intensive margin responses, whereby workers of different types can earn the same income. Earnings-based taxation cannot be type-specific, and wage compression becomes a valuable redistributive channel. In our framework, each occupation corresponds to a unique earnings level. Taxation is effectively occupation-specific conditional on working, and pecuniary spillovers between occupations or between capital and labor lose their policy relevance. This incidence-irrelevance property would carry over to a model with occupational switching—including upgrading from the automatable to the non-automatable sector—because the occupation-specific instruments still allow the government to neutralize equilibrium spillovers. The property would not carry over to a model with intensive margin responses *within occupations*.

Our theoretical results have close precedents in the literature. Saez (2004) shows that partial-equilibrium results remain valid in general equilibrium in a long-run model of occupational choice. Scheuer and Werning (2018) show that optimal income taxation is independent of production technology by reformulating Mirrlees (1971) as a special case of Diamond and Mirrlees (1971). Jacquet and Lehmann (2025) show that the production efficiency theorem of Diamond and Mirrlees (1971) hinges only on the flexibility of the tax system to target each factor's income, using a GE-neutralizing reform approach related to ours.[5] We apply these insights to the implications of

[5] Also related, Saez and Zucman (2026) argue that tax incidence is irrelevant for tax reform analysis because any

automation for redistribution, extending them to an economy with capitalists and an endogenous rental rate. Our numerical analysis shows that optimal tax schedules change substantially.

The second literature is the theory of optimal capital income taxation, including classic zero-tax results (Atkinson and Stiglitz 1976; Judd 1985; Chamley 1986) and sufficient statistics formulas expressed in terms of the capital supply elasticity (Saez and Stantcheva 2018). Our analysis shows that the sufficient statistics formula of Saez and Stantcheva (2018) survives intact in general equilibrium. With workers compensated, the capital tax problem collapses to a partial-equilibrium trade-off between revenue extraction and the capitalists' welfare weight. At the long-run capital supply elasticities estimated by Jakobsen *et al.* (2020), between 0.5 and 2 across specifications, the optimal tax on capitalists is large.

The final literature concerns the policy implications of automation. This work has focused on the taxation of machines. Guerreiro *et al.* (2022) find that robots should be taxed during the transition but not in the steady state; Thuemmel (2023) finds an optimal robot tax that is small and of ambiguous sign; Costinot and Werning (2023) derive an optimal robot tax of roughly 4 percent that tends to decline as robots become cheaper; Beraja and Zorzi (2025) show that automation can be inefficiently fast when displaced workers reallocate slowly and face borrowing constraints; and Acemoglu *et al.* (2020) argue that the US tax code favors capital over labor, generating excessive automation.[6] In our model, the linear capital income tax is formally equivalent to a uniform tax on machine inputs in production. But the normative rationale for our tax is different, resting on revenue extraction from imperfectly elastic capitalists rather than production inefficiencies. At empirically reasonable elasticities, our optimal capital tax is an order of magnitude larger than the robot taxes found in the prior literature.[7]

The paper is organized as follows. Section 2 sets out the task-based model of automation. Section 3 characterizes the optimal taxation of labor and capital income. Section 4 quantifies how much automation changes optimal tax rates in simulations calibrated to the US wage distribution and automation exposure. Section 5 concludes.

---

effects on pre-tax prices can be offset with additional tax adjustments.

[6]The non-zero robot taxes in Guerreiro *et al.* (2022), Thuemmel (2023), and Costinot and Werning (2023) rely on a failure of the production efficiency theorem (Diamond and Mirrlees 1971). The production-efficiency failure in these papers relates to Naito (1999) and arises because their income taxes cannot condition on a worker's occupation. This force for machine taxes is not present in our setting, where labor taxes condition on occupation-specific wages.

[7]To our knowledge, our paper is the first to make optimal income taxation the central object of study in a task-based model of automation. In related work on technology and income taxation, Ales *et al.* (2015) study optimal taxation in an assignment model of tasks and talents in which technical change reshapes the optimal schedule, and Loebbing (2022) analyzes optimal income taxation under directed technical change, whereby the tax system influences the direction of innovation, a channel we abstract from.

2 A Task-Based Model of Automation

The model has four building blocks, described below. The production technology is task-based: tasks up to a skill threshold are automatable—they can be performed by either workers or machines—while tasks above the threshold are not. Workers take their occupation (task) as given and choose whether or not to work. Capitalists supply machines to the automatable tasks at a finite elasticity. The government levies a nonlinear labor tax and a linear capital tax.

2.1 Production Technology

The production side of the model is a modified version of Zeira (1998) and Moll (2025). Aggregate output Y is generated by combining tasks $i = 1, ..., N$ using a Cobb-Douglas technology:

$$Y = AX_1^{1/N} X_2^{1/N} \cdot ... \cdot X_N^{1/N}, \tag{1}$$

where A is total factor productivity and X_i denotes the input of task services of type i.

Each task is either automatable or not. Automatable tasks can be produced by either capital or labor, while non-automatable tasks can be produced only by labor. We order tasks in terms of increasing sophistication: the skill level of workers required to perform a given task (or equivalently, a given occupation) is increasing in $i = 1, ..., N$. We assume that tasks $i = 1, ..., n$ can be automated while tasks $i = n + 1, ..., N$ cannot. The automatable share of tasks, n/N, can take any value, including one. We have

$$X_i = \begin{cases} \omega_i L_i + \rho_i K_i & \text{for } i \leq n \\ \omega_i L_i & \text{for } i > n, \end{cases} \tag{2}$$

where L_i and K_i denote labor and capital input in task i. This specification implies that, in tasks where automation is possible, workers and machines are perfect substitutes. The productivity of workers in task i equals ω_i, where $\omega_i < \omega_{i+1}$ for all i. The productivity of machines in an automatable task i equals ρ_i.

In this model, technological innovation takes two distinct forms: (1) an increase in the share of automatable tasks n/N ("displacement effect"), and (2) an increase in the productivity of machines in automatable tasks ρ_i ("capital-augmenting technological change"). We examine the policy implications of each type of innovation.

A representative firm chooses labor inputs $\{L_i\}_{i=1}^N$ and capital inputs $\{K_i\}_{i=1}^n$ to maximize profits, defined as

$$Y - \sum_{i=1}^{N} \left(w_i L_i + R K_i \cdot \mathbf{1}\left[i \le n \right] \right), \tag{3}$$

where Y is specified in equations (1)-(2), w_i denotes the wage rate of workers in task i, and R denotes the rental rate of capital. Both wages w_i and the rental rate R are determined in equilibrium by clearing of labor and capital markets.

Labor demand in non-automatable tasks is characterized by the following first-order condition:

$$\frac{w_i L_i}{Y} = \frac{1}{N} \quad \forall i > n. \tag{4}$$

That is, the labor share of skill type i equals the Cobb-Douglas weight $1/N$ of the associated task.

Consider next automatable tasks. We restrict attention to interior solutions where firms use both workers and machines. This is the realistic case when tasks correspond to skill levels (or occupation ranks) rather than within-occupation "micro-tasks." At our level of aggregation, automation does not eliminate the need for human workers entirely, but it will affect their wages and employment in equilibrium. At an interior solution, we have

$$\frac{w_i L_i + R K_i}{Y} = \frac{1}{N} \quad \text{and} \quad \frac{w_i}{\omega_i} = \frac{R}{\rho_i} \quad \forall i \le n. \tag{5}$$

The first condition in (5) implies that the value-added share of task i equals its Cobb-Douglas weight $1/N$. The second condition ensures that the marginal cost of task i is the same regardless of whether the last unit is produced from labor ($= w_i/\omega_i$) or from capital ($= R/\rho_i$). As a result, the firm is indifferent between workers and machines in equilibrium. This condition ties equilibrium wages in automatable tasks to the rental rate, $w_i = R\omega_i/\rho_i$. Technological innovation through a higher ρ_i depresses equilibrium wages in the affected tasks.

A similar logic determines when automation is actually introduced. Firms employ machines in task i only if doing so lowers the cost of performing the task, which requires that the capital cost per efficiency unit, R/ρ_i, is smaller than the labor cost per efficiency unit under labor-only production, w_i^L/ω_i, where w_i^L denotes the equilibrium wage without machines. Equivalently, automation is adopted only if the resulting wage does not exceed the labor-only wage, i.e. if $w_i = R\omega_i/\rho_i \le w_i^L$. Assuming that the introduction of automation is associated with fixed costs, the automation wage must be *strictly* smaller than the labor-only wage for automation to take place.

2.2 Workers

There is a discrete distribution of workers with heterogeneous skills $\omega_1 < \omega_2 < ... < \omega_N$. Workers draw skill from a probability mass function denoted by $h(\omega_i) = h_i$. With the worker population normalized to one, h_i also corresponds to the number of workers with skill ω_i. Each skill level is associated with a given task (or occupation) in the production technology specified in equations (1)-(2). As described above, the skill parameter ω_i maps onto a realized wage rate w_i. A worker's occupation is fixed. We abstract from the possibility of retraining to switch between automatable and non-automatable occupations.[8]

We consider an extensive margin model of labor supply similar to Saez (2002). Workers decide whether or not to work, but hours worked conditional on working are normalized to one. To capture extensive margin responses in terms of smooth elasticities, we incorporate fixed costs of work q and allow these to be heterogeneous in the population. Within each skill level, there is a continuous distribution of fixed work costs given by the distribution function $P(q|\omega_i)$ and density function $p(q|\omega_i)$.

The government imposes taxes and provides transfers based on earnings z_i, where $z_i = w_i$ if working and $z_i = 0$ if not working. Taxes and transfers are embodied in the function $T(z_i)$, which will be negative for a worker receiving net transfers. The utility of a worker with wage rate w_i when working and not working, respectively, is given by

$$u_i = \begin{cases} w_i - T(w_i) - q & \text{if working} \\ -T(0) & \text{if not working.} \end{cases} \tag{6}$$

We denote the utility level of an employed worker by $u(w_i, q)$ and of an unemployed worker by $u(0)$. Each individual chooses to work if and only if $u(w_i, q) \geq u(0)$, i.e.

$$w_i - T(w_i) - q \geq -T(0), \tag{7}$$

which implies

$$q \leq w_i - [T(w_i) - T(0)] = w_i(1 - \tau_i) \equiv \bar{q}_i, \tag{8}$$

where $\tau_i \equiv (T(w_i) - T(0))/w_i$ denotes the *participation tax rate.* This is an average tax rate that

[8] This assumption is consistent with recent evidence suggesting that workers displaced by robots exit employment rather than switch occupations (Acemoglu and Restrepo 2020).

accounts for both taxes paid and benefits lost when entering the labor market. In a given task, everyone with a fixed cost below the threshold $\bar{q}_i$ will be employed and everyone with a fixed cost above $\bar{q}_i$ will be unemployed. Hence, the employment rate among workers in task i is given by $P(\bar{q}_i|\omega_i)$. Aggregate labor supply in task i is given by $P(\bar{q}_i|\omega_i) \cdot h_i$, which in equilibrium must equal labor demand L_i characterized in the previous section.

Finally, for characterizing optimal tax policy, it will be useful to define an extensive margin labor supply elasticity:

$$\eta_i \equiv \frac{d \ln P(\bar{q}_i|\omega_i)}{d \ln \bar{q}_i}. \tag{9}$$

From the definition of $\bar{q}_i$ in equation (8), η_i is the elasticity of extensive margin labor supply with respect to the (average) net-of-tax wage rate. Note that η_i is a *micro* elasticity: it captures workers' supply responses to changes in the net-of-tax wage, but not the general-equilibrium effects on wages (operating through labor demand) that will feed back into labor supply decisions and thus equilibrium employment.

2.3 Capitalists

The economy includes a class of capitalists who own and supply the capital used in automatable tasks. There are h_K identical capitalists. Capital income is taxed at the linear rate t_K, such that the net-of-tax return to capital is given by $(1 - t_K) R$. Capitalists have a reduced-form utility function of the form $u_K = u_K((1 - t_K) RK, K)$, where the first argument reflects the utility of consumption and the second argument reflects the opportunity cost of supplying capital to domestic production. We assume that capitalists' utility is quasi-linear and iso-elastic:

$$u_K = (1 - t_K) RK - \frac{\chi}{1 + 1/\zeta_K} \left(\frac{K}{\chi}\right)^{1+1/\zeta_K}. \tag{10}$$

The first-order condition for capital supply gives

$$K = \chi \left[(1 - t_K) R\right]^{\zeta_K}, \tag{11}$$

where $\chi > 0$ is a scale parameter and ζ_K is the elasticity of capital supply with respect to the net-of-tax return $(1 - t_K) R$. The limit of $\zeta_K \to \infty$ corresponds to perfectly elastic capital supply, as would arise in a small open economy facing a global capital market.

Our reduced-form modeling of capitalists follows the wealth-in-utility approach of Saez and Stantcheva (2018). Rather than modeling the dynamics of saving and wealth accumulation explicitly, it captures their policy-relevant content in a single parameter ζ_K: the long-run elasticity of capital supply with respect to the net-of-tax return. We calibrate this elasticity to empirical estimates in Section 4.

2.4 Equilibrium

An equilibrium of this model satisfies firms' optimal factor demands in each task, characterized in equations (4)-(5). The equilibrium also satisfies workers' optimal labor supply in each task, which equals $P(\bar{q}_i|\omega_i) \cdot h_i$ where $\bar{q}_i$ is characterized in equation (8). Labor market clearing implies $P(\bar{q}_i|\omega_i) \cdot h_i = L_i \;\; \forall i$. In automatable tasks ($i \leq n$), partial-equilibrium labor demand is perfectly elastic: at a given rental rate, firms absorb any amount of labor at the wage rate $w_i = R\omega_i/\rho_i$, with capital satisfying the residual input requirements to produce the desired amount of task services X_i. General-equilibrium labor demand, by contrast, is not perfectly elastic. With imperfectly elastic capital supply, changes in employment move the equilibrium rental rate and therefore wages in all automatable tasks. Finally, the rental rate R clears the capital market, $\sum_{i\leq n} K_i = h_K K$, where capital supply K is given by equation (11).

2.5 Government

The government's preferences are described by a concave social welfare function defined over the individual utilities of workers and capitalists. We express social welfare as

$$W = \sum_{i=1}^{N} \left[\int_0^{\bar{q}_i} \Psi \left[u \left(w_i, q \right) \right] dP \left(q|\omega_i \right) + \int_{\bar{q}_i}^{\infty} \Psi \left[u \left(0 \right) \right] dP \left(q|\omega_i \right) \right] h_i + h_K \Psi \left[u_K \right], \tag{12}$$

where Ψ is an increasing and concave transformation of individual utilities ($\Psi' > 0, \Psi'' < 0$). The first two terms capture social welfare from employed and unemployed workers, respectively, while the last term captures social welfare from the capitalists. For workers in automatable tasks ($i \leq n$), the wage is tied to the rental rate by $w_i = R\omega_i/\rho_i$, so their utility level is $u(R\omega_i/\rho_i, q)$.

The government sets labor taxes at each earnings level $\{T(0), T(w_1), ..., T(w_N)\}$ and the capital tax rate t_K to maximize social welfare in (12) subject to the equilibrium conditions derived in previous sections and a government budget constraint given by

$$\sum_{i=1}^{N} \left[T \left(w_i \right) P \left(\bar{q}_i|\omega_i \right) + T \left(0 \right) \left(1 - P \left(\bar{q}_i|\omega_i \right) \right) \right] h_i + t_K R h_K K \geq G, \tag{13}$$

where G is an exogenous revenue requirement.

The government's redistributive preferences can be captured by social marginal welfare weights. Denoting the Lagrange multiplier associated with the government budget constraint by μ (corresponding to the marginal value of public funds), we define welfare weights on employed workers, unemployed workers, and capitalists as follows:

$$g_{i\le n}\left(R\omega_i/\rho_i\right) = \frac{\int_0^{\bar{q}_i}\Psi'\left[u\left(R\omega_i/\rho_i,q\right)\right]dP\left(q|\omega_i\right)}{\mu P\left(\bar{q}_i|\omega_i\right)} \tag{14}$$

$$g_{i>n}\left(w_i\right) = \frac{\int_0^{\bar{q}_i}\Psi'\left[u\left(w_i,q\right)\right]dP\left(q|\omega_i\right)}{\mu P\left(\bar{q}_i|\omega_i\right)} \tag{15}$$

$$g_0 = \frac{\Psi'\left[u\left(0\right)\right]}{\mu} \tag{16}$$

$$g_K = \frac{\Psi'\left[u_K\right]}{\mu}. \tag{17}$$

All welfare weights are equilibrium objects. They depend on wages, the rental rate, employment rates, and the tax system. As we shall see, the endogeneity of welfare weights to the general equilibrium is the key mechanism through which automation affects optimal taxation.

3 Optimal Tax Design

We solve for the optimal tax system—the participation taxes on workers and the capital tax on capitalists—using a perturbation approach. The idea of this approach is that, at the optimum, small changes in the tax system have no first-order effects on social welfare. The approach was developed in partial-equilibrium models (e.g., Saez 2001, 2002). We propose a version of the approach tailored to general-equilibrium settings. The key is to choose the perturbations wisely. A change in any single tax sets off equilibrium adjustments in all wages and the rental rate, entangling the optimality condition for one instrument with spillovers to every agent in the economy. For each instrument, we therefore consider a *compensated reform* that switches off these spillovers by construction, isolating the agents targeted by the instrument.[9]

[9] Our compensated perturbation approach is related to prior uses of compensating tax adjustments. Recent papers by Vergara (2026) and Jacquet and Lehmann (2025) are closest to our approach. Vergara (2026) pairs a minimum-wage increase with an equal tax increase on affected workers, holding their labor supply and utility fixed, and a corporate-tax cut that holds the capitalist's utility fixed. Jacquet and Lehmann (2025) develop a general version: assuming that the tax system can target each factor's income, they consider GE-neutralizing tax reforms that offset the effect of factor-price changes on every agent's budget, holding their behavior and utility fixed. A related literature bundles tax perturbations with income-tax adjustments that compensate workers in utility terms: Kaplow (2006) in partial equilibrium, Schulz *et al.* (2023) and Costinot and Werning (2023) in general equilibrium, building on the trade-policy analysis of Dixit

3.1 Labor Income Taxation

Consider a compensated reform of the tax liability in task i. Specifically, suppose the government increases the liability in task i by dT_i and makes two compensating adjustments. First, in every other task $j \neq i$, it changes tax liability one-for-one with the wage change induced by the reform:

$$dT_j = dw_j \qquad \forall j \neq i. \tag{18}$$

Second, it changes the capital tax rate to hold the net-of-tax rental rate fixed, $d\left[(1-t_K)\,R\right] = 0$:

$$dt_K = (1-t_K)\,\frac{dR}{R}. \tag{19}$$

The first adjustment holds every other task's net gain from working, $\bar{q}_j = w_j - [T(w_j) - T(0)]$, and their net income, $w_j - T(w_j)$, fixed. Hence, outside task i, employment and worker consumption are unchanged. The second adjustment holds the capitalists' net return fixed, leaving their capital supply and utility unaffected. Both adjustments are needed regardless of which task is reformed: a tax change in any task moves employment, output, and the demand for capital, and hence every non-automatable wage $w_j = Y/(NL_j)$, every automatable wage $w_j = R\omega_j/\rho_j$, and the rental rate R.

In general, a tax reform has four types of welfare effects: a mechanical revenue effect (dM_i), a behavioral revenue effect (dB_i), a direct welfare effect (dW_i), and a cross effect (dC_i) on all other agents, namely workers in other tasks and capitalists. The compensated reform is constructed exactly to ensure that the last effect vanishes, $dC_i = 0$. We derive the three remaining effects below and characterize optimal labor income taxation from the condition $dM_i + dB_i + dW_i = 0$ in every task i. Despite the rich general-equilibrium setting, our approach delivers the optimal tax schedule in a few simple steps, and the implied schedule corresponds to the canonical partial-equilibrium rule of Saez (2002). As we describe below, this finding relates to classic results by Diamond and Mirrlees (1971).

Consider first the *mechanical revenue effect*. Let $E_i \equiv P(\bar{q}_i|\omega_i)\,h_i$ denote the number of employed workers in task i. Holding behavior fixed, each of the E_i employed workers in task i pays an additional dT_i in taxes; the compensating adjustments change the tax liability of each of the E_j

and Norman (1986). Because these utility compensations leave behavioral responses in place, they do not eliminate general-equilibrium effects. Our compensation holds net incomes fixed, so that no behavioral response occurs outside the perturbed margin.

workers in every other task by $dT_j = dw_j$; and the capital-tax adjustment collects the entire change in gross capital income from the capitalists, $h_K K \cdot d(t_K R) = h_K K \cdot dR$.[10] Differentiating the zero-profit identity $Y = \sum_j w_j L_j + R \sum_{j \leq n} K_j$, subtracting the production-side condition $dY = \sum_j w_j \cdot dL_j + R \cdot d\left(\sum_{j \leq n} K_j\right)$ implied by profit maximization (envelope theorem), and using market clearing ($L_j = E_j$ and $\sum_{j \leq n} K_j = h_K K$), we obtain

$$\sum_{j=1}^{N} E_j \cdot dw_j + h_K K \cdot dR = 0. \tag{20}$$

Given zero profits and firm optimization, tax-induced factor-price adjustments are pure transfers between workers and capitalists. They create no aggregate income. The mechanical revenue effect of the compensated reform is therefore

$$dM_i = E_i \cdot dT_i + \sum_{j \neq i} E_j \cdot dw_j + h_K K \cdot dR = E_i \cdot (dT_i - dw_i). \tag{21}$$

Second, there is a *behavioral revenue effect* as the reform changes the net gain from working by $d\bar{q}_i = dw_i - dT_i$, accounting for both the direct tax effect and any indirect wage effect in task i. Using the definition of the extensive margin elasticity in equation (9), this creates an employment effect equal to $dE_i = \eta_i E_i d\bar{q}_i / \bar{q}_i$. Because each worker pays a participation tax $T(w_i) - T(0)$, the fiscal externality of employment responses can be written as

$$dB_i = [T(w_i) - T(0)] \cdot dE_i = -\frac{\tau_i}{1-\tau_i} \cdot \eta_i \cdot E_i \cdot (dT_i - dw_i). \tag{22}$$

Third, the *direct welfare effect* captures the welfare-weighted utility implications of the reform. Importantly, because marginal workers who exit the labor force are indifferent between working and not working, their extensive margin responses have no first-order utility effects (envelope theorem). The only first-order utility effects come from the change in disposable income among *inframarginal* workers in task i. We obtain

$$dW_i = -g_i \cdot E_i \cdot (dT_i - dw_i). \tag{23}$$

All three effects are proportional to the factor $dT_i - dw_i$. What matters is the part of the tax increase that is not passed through to the worker's own wage. The optimality condition is

[10] Under equation (19), we have $d(t_K R) = R dt_K + t_K dR = dR$. Hence, with the net-of-tax return held fixed, the entire change in the gross rental rate accrues to the government.

therefore informative only if this pass-through is incomplete, which is always the case. Under the compensated reform, employment in all other tasks ($j \neq i$) and the aggregate capital stock are held fixed. The wage in task i itself, whether automatable or non-automatable, responds to the reform through a downward-sloping labor demand curve. In non-automatable tasks, the labor demand elasticity can be derived as $\varepsilon_i^D \equiv -d\ln L_i / d\ln w_i = N/(N-1)$; in automatable tasks, the labor demand elasticity equals $\varepsilon_i^D = nY/[(N-n) w_i L_i]$.[11] In equilibrium, the change in labor supply, $d\ln E_i = \eta_i \cdot d\ln \bar{q}_i = \eta_i (dw_i - dT_i)/\bar{q}_i$, must be equal to the change in labor demand, $d\ln L_i = -\varepsilon_i^D \cdot d\ln w_i$. Equating the two and using the definition $\bar{q}_i = w_i(1-\tau_i)$ implies

$$\frac{dw_i}{dT_i} = \frac{\eta_i}{\eta_i + \varepsilon_i^D (1-\tau_i)} \in [0,1). \tag{24}$$

Hence, a tax increase on workers in a given task (compensating all other workers and capitalists for general-equilibrium effects) lowers the net gain from working: $dw_i - dT_i < 0$.[12]

We characterize optimal labor taxes as follows:

Proposition 1 (Optimal Participation Taxes). *In every task i, automatable and non-automatable alike, the optimal participation tax rate satisfies*

$$\tau_i = \frac{1-g_i}{1-g_i+\eta_i}, \tag{25}$$

where g_i is the social marginal welfare weight on employed workers in task i, equal to $g_{i\leq n}(R\omega_i/\rho_i)$ in automatable tasks and $g_{i>n}(w_i)$ in non-automatable tasks, and η_i is the extensive margin labor supply elasticity.

Proof. This follows from the optimality condition $dM_i + dB_i + dW_i = 0$ in every task i, with each term characterized in equations (21)-(23). This implies $E_i(dT_i - dw_i)\left[1 - g_i - \frac{\tau_i}{1-\tau_i}\eta_i\right] = 0$. Because $dT_i - dw_i \neq 0$ by equation (24), optimality requires $1 - g_i - \frac{\tau_i}{1-\tau_i}\eta_i = 0$ and rearranging gives equation (25). □

[11] With the net-of-tax return held fixed, aggregate capital supply $h_K K$ is fixed. Consider first a non-automatable task i. With employment fixed in every other task $j \neq i$, the capital market clearing condition of Section 2.4—which may be rewritten as $nY/(NR) = h_K K + \sum_{j\leq n} \omega_j L_j/\rho_j$ by using equation (5)—implies that Y/R stays constant, so that every automatable task input $X_j = \rho_j Y/(NR)$ is also constant (as are non-automatable task inputs $j \neq i$). Hence, output varies only through task i: $d\ln Y = d\ln L_i/N$ by equation (1). Equation (4) ($w_i = Y/(NL_i)$) then gives $d\ln w_i = d\ln Y - d\ln L_i = -[(N-1)/N]\, d\ln L_i$, which implies $\varepsilon_i^D = N/(N-1)$. Consider next an automatable task i, whose wage moves with the rental rate, $d\ln w_i = d\ln R$. With non-automatable inputs fixed, the production function implies $d\ln Y = (n/N)(d\ln Y - d\ln R)$, and differentiating the capital market clearing condition (with L_i now entering its right-hand side) yields $d\ln w_i = d\ln R = -[(N-n) w_i L_i/(nY)]\, d\ln L_i$, which implies $\varepsilon_i^D = nY/[(N-n) w_i L_i]$.

[12] The upper bound in equation (24) uses $\tau_i < 1$, which holds at any optimum (see the proof of Corollary 1 below).

Equation (25) corresponds to the canonical partial-equilibrium tax formula of Saez (2002). The proposition establishes that this formula holds even in our general-equilibrium model, where taxes in any task affect the task's own wage, spill over to the wages of every other task, and shift the return to capital. Whatever form this incidence takes, it does not affect the optimal tax formula. This does not imply that the optimal tax *schedule* is unaffected, however. The endogeneity of wages and the capital return feeds into the social marginal welfare weights and, as we show through calibrated simulations, has substantial policy implications.

The reason that the partial-equilibrium tax rule survives rests on the targeting logic of Diamond and Mirrlees (1971). The government has one tax instrument per occupation (wage level) plus a tax on capital income, so the tax system reaches every agent's consumption directly. Whatever equilibrium prices turn out to be, labor taxes can deliver any desired net income in each occupation, and the capital tax any desired net return to capitalists. Tax-induced price movements therefore give the government nothing it does not already control: any redistribution implied by wages or the rental rate can be accomplished directly through the tax system.

Proposition 1 runs against a literature on optimal taxation with general-equilibrium wage incidence (Stiglitz 1982; Rothschild and Scheuer 2013; Sachs *et al.* 2020), in which incidence effects provide an argument for lower marginal tax rates on high-skill workers through trickle-down. The key difference is that these papers focus on intensive-margin rather than extensive-margin responses. Workers adjust earnings conditional on working—by varying hours worked (Stiglitz 1982; Sachs *et al.* 2020) or by varying effort and switching sectors (Rothschild and Scheuer 2013)—so that workers of different types can earn the same income. Taxes based on earnings therefore cannot replicate type-specific taxation, and redistribution is constrained by incentive compatibility: high-skill workers must be deterred from mimicking low-skill workers, and the cost of mimicking depends on relative pre-tax wages.[13] Tax-induced wage compression therefore relaxes the binding incentive constraints. Lower marginal tax rates at the top stimulate high-skill labor supply, compress the pre-tax wage distribution, and redistribute indirectly toward low-skill workers.

[13]In the mechanism-design formulations of Stiglitz (1982) and Rothschild and Scheuer (2013), tax-induced wage changes affect the optimum conditions only through the multipliers on the self-selection constraints, and these constraints depend on before-tax wages: given a menu of consumption-earnings bundles, the labor input required to attain any given earnings level is determined by the pre-tax wage. Sachs *et al.* (2020) work instead with variational methods and no explicit incentive constraints; there, the same friction appears as tax-induced wage changes that reshuffle pre-tax incomes across workers facing different marginal tax rates and carrying different welfare weights. Their trickle-down prescription characterizes the optimum; for reforms of a suboptimal progressive tax schedule, they show that general-equilibrium effects make raising top marginal tax rates more, not less, desirable.

In our economy, by contrast, each occupation corresponds to a unique earnings level. Hence, labor taxation is effectively occupation-specific conditional on working and, combined with a separate capital tax, this allows the government to neutralize general-equilibrium incidence with offsetting tax adjustments. Our compensated perturbation approach exploits precisely this feature. Accordingly, incidence effects place no restraint on optimal tax schedules—not because such effects are absent, but because complete taxation of occupations and capital income eliminates their policy relevance. The incidence-irrelevance property would carry over to a model with occupational switching—including upgrading from the automatable to the non-automatable sector—because the occupation-specific instruments still allow the government to neutralize equilibrium spillovers. The property would not carry over to a model with intensive margin responses *within occupations*.

The key role of the labor supply margin has close precedents in the literature. Saez (2004) contrasts the model of Stiglitz (1982), in which workers of fixed skill types choose hours of work, with a long-run model of occupational choice, in which labor supply adjusts through entry into jobs. In the occupational-choice model, each wage level corresponds to a unique occupation—just as in our model—and the results of Diamond and Mirrlees (1971) are restored: optimal tax formulas take the same form whether factor prices are fixed or determined in general equilibrium, production efficiency holds, and the production-inefficiency results of Naito (1999) break down.[14] Scheuer and Werning (2018) establish a general version of this logic for the Mirrlees model, embedding it into the Diamond-Mirrlees framework by treating each earnings level as a distinct commodity taxed at its own linear rate. Within a class of preferences in which workers choose among different earnings levels rather than the intensity of a given labor type, the optimal income tax formula is independent of production technology, and the general-equilibrium corrections of the trickle-down literature can be interpreted as a consequence of missing occupation- or sector-specific tax instruments. Our model is a member of the Saez-Scheuer-Werning family, where the income tax instruments are complete and any incidence effects may be offset through compensating adjustments.

We now turn to the implications of the optimal tax formula for the taxation of workers in low-wage, automatable occupations. Negative participation taxes—corresponding to an EITC—are generally optimal at the bottom. We have:

[14] The Roy model of Rothschild and Scheuer (2013) contains both margins: workers vary effort within sectors and switch sectors. Consistent with the margin logic, they find that occupational switching dampens the trickle-down force: the optimal tax schedule is less progressive than with exogenous wages, but more progressive than in the fixed-occupation model of Stiglitz (1982).

Corollary 1 (EITC in Automatable Tasks). *Whenever $g_{i\leq n}\left(R\omega_i/\rho_i\right) > 1$, a participation subsidy, $\tau_i < 0$, is socially optimal in task i. In this case, a higher marginal productivity of capital ρ_i (capital-augmenting technological progress) implies a larger participation subsidy in that task.*

Proof. This result follows from equation (25) provided that $1 - g_i + \eta_i > 0$. To see that this is satisfied, suppose instead $1 - g_i + \eta_i \leq 0$ (which requires $g_i > 1$). Then equation (25) admits no participation tax consistent with positive employment: it implies $\tau_i > 1$ when the inequality is strict and has no solution when it binds, so employment in task i is zero. In such a situation, let the government reduce τ_i until employment E_i turns marginally positive. The welfare effect of such a reform equals $dM_i + dB_i + dW_i = dB_i = \left[T\left(w_i\right) - T\left(0\right)\right] \cdot dE_i > 0$. Hence, an equilibrium with $1 - g_i + \eta_i \leq 0$ cannot be optimal. The comparative static in ρ_i follows from equation (25): the optimal tax is strictly decreasing in the welfare weight ($d\tau/dg = -\eta/\left(1 - g + \eta\right)^2 < 0$); a higher ρ_i lowers the wage $w_i = R\omega_i/\rho_i$;[15] this raises g_i due to the concavity of social welfare and therefore lowers τ_i. Given the tax is negative, a higher ρ_i enlarges the subsidy. □

The EITC case described above is the typical outcome. Because welfare weights are measured in terms of the marginal value of public funds μ, and given the absence of income effects in the model, welfare weights average to one in the worker population.[16] Hence, if the government's preferences are associated with declining welfare weights across occupations, it is generally the case that $g_{i\leq n}\left(R\omega_i/\rho_i\right) > 1$ in the lowest-wage, automatable occupations.

The consequences of technological progress for optimal taxation follow immediately. Recall from Section 2.1 that innovation takes two forms: an increase in the share of automatable tasks n/N (displacement effect) and an increase in machine productivity ρ_i in automatable tasks (capital-augmenting technological change). Because the optimal tax in every task is governed by the same formula, both forms of progress transmit to the optimal labor income tax schedule exclusively through equilibrium wages and the welfare weights evaluated at those wages.[17]

[15] The wage falls even though the rental rate is endogenous: log-linearizing the equilibrium at fixed taxes shows that $d\ln R/d\ln\rho_i \in (0,1)$ whenever machines are in use. A productivity gain in one automated task raises the demand for capital and hence the rental rate, but less than proportionally, so $w_i = R\omega_i/\rho_i$ declines.

[16] This property follows from optimality with respect to a uniform shift dT of the entire tax schedule $\{T\left(0\right), T\left(w_1\right), ..., T\left(w_N\right)\}$. This leaves every net gain from working unchanged and therefore triggers no behavioral responses, creating only a mechanical revenue effect dM and a direct welfare effect dW. Optimality requires $dM + dW = 0$, which implies $dT - dT \cdot \sum_{i=1}^{N} h_i \left[P\left(\bar{q}_i|\omega_i\right) g_i + \left(1 - P\left(\bar{q}_i|\omega_i\right)\right) g_0\right] = 0$ and therefore $\mathbb{E}\left[g_i\right] = 1$. Capitalists, who are reached by the capital income tax rather than the labor income tax, are not part of this average.

[17] This statement holds the extensive margin elasticities fixed, as with iso-elastic labor supply, and treats the welfare weights as functions of the wage alone; more generally, technological progress also moves η_i and g_i through the equilibrium participation threshold $\bar{q}_i$ and the marginal value of public funds μ, and it moves the optimal capital tax through the capitalists' weight g_K.

Consider first the displacement effect. When task $n+1$ becomes automatable, its wage falls below the labor-only wage w_{n+1}^{L}. At the adoption threshold, where machine productivity just makes automation cost-neutral, no machines are used and the automation wage coincides with w_{n+1}^{L}; any further increase in ρ_{n+1} strictly lowers the wage, as established in the proof of Corollary 1; and adoption requires strictly cost-reducing machines (Section 2.1). With concave social welfare, the welfare weight of the newly automated task therefore rises and its optimal participation tax falls. Hence, displacement unambiguously lowers the optimal tax in affected tasks, and it enlarges the participation subsidy whenever the welfare weight exceeds one. Capital-augmenting progress works the same way within the set of already-automated tasks (Corollary 1).

Beyond the directly affected tasks, technological progress spills over through two prices. Aggregate output moves every non-automatable wage $w_j = Y / (NL_j)$, and the rental rate moves every automatable wage. A capital-productivity gain in one automated task, for instance, raises the wages of all other automated tasks. These price movements shift welfare weights—and thereby optimal taxes—throughout the economy. That is, even though the general-equilibrium structure does not explicitly enter the reduced-form tax formula, the endogeneity of wages and the rental rate nevertheless shift the entire tax structure.

3.2 Capital Income Taxation

Consider next the taxation of capitalists. As above, we characterize the optimal tax formula using a compensated reform perturbation. The government raises the capital tax rate by dt_K and adjusts every task's liability one-for-one with the induced wage change, $dT_j = dw_j$ for all $j = 1, ..., N$. The compensations hold every task's net gain from working and net income fixed. Hence, employment and worker consumption are unchanged throughout the economy, and non-workers are also unaffected. The reform reaches only the capitalists and the government budget.

The welfare effects parallel equations (21)–(23). Consider first the mechanical revenue effect: holding behavior fixed, capitalists pay an additional $h_K K \cdot d(t_K R)$ in taxes, and the compensating adjustments change collections from workers by $\sum_j E_j dw_j = -h_K K \cdot dR$, using the zero-sum identity (20). The mechanical revenue effect is therefore

$$dM_K = h_K K \cdot d(t_K R) + \sum_{j=1}^{N} E_j \, dw_j = -h_K K \cdot d\left[(1-t_K) R\right]. \tag{26}$$

Second, the supply of capital moves with the net-of-tax return according to equation (11), and each unit of capital withdrawn costs the government $t_K R$ in revenue:

$$dB_K = t_K R \cdot h_K dK = \frac{t_K \zeta_K}{1-t_K} \cdot h_K K \cdot d\left[\left(1-t_K\right) R\right]. \tag{27}$$

Third, because capitalists choose K optimally, the utility effect of the reform equals the capital stock times the change in the net-of-tax return (envelope theorem), so the direct welfare effect is

$$dW_K = g_K \cdot h_K K \cdot d\left[\left(1-t_K\right) R\right]. \tag{28}$$

All three effects are proportional to the common factor $d\left[\left(1-t_K\right) R\right]$: what matters is the part of the tax increase that is not shifted into the gross rental rate. This pass-through is again incomplete. With employment fixed in every task by the compensations, capital demand has elasticity $\varepsilon_K^D \equiv -d\ln K / d\ln R = nY / \left[\left(N-n\right) R h_K K\right]$—the exact counterpart of the labor demand elasticity in automatable tasks, with aggregate capital income $R h_K K$ in place of the task wage bill $w_i L_i$.[18] In equilibrium, the change in capital supply, $d\ln K = \zeta_K \left[d\ln R - dt_K / \left(1-t_K\right)\right]$, must be equal to the change in capital demand, $d\ln K = -\varepsilon_K^D \cdot d\ln R$. Equating the two, this implies

$$\frac{d\ln R}{dt_K / \left(1-t_K\right)} = \frac{\zeta_K}{\zeta_K + \varepsilon_K^D} \in [0,1). \tag{29}$$

A tax increase therefore always lowers the net-of-tax return: $d\left[\left(1-t_K\right) R\right] < 0$. We obtain:

Proposition 2 (Optimal Capital Tax). *The optimal capital income tax rate satisfies*

$$t_K = \frac{1-g_K}{1-g_K+\zeta_K}, \tag{30}$$

where g_K is the social marginal welfare weight on capitalists, and ζ_K is the elasticity of capital supply.

Proof. This follows from the optimality condition $dM_K + dB_K + dW_K = 0$, with each term characterized in equations (26)-(28). This implies $-h_K K \cdot d\left[\left(1-t_K\right) R\right] \left[1 - g_K - \frac{t_K}{1-t_K}\zeta_K\right] = 0$. Because $d\left[\left(1-t_K\right) R\right] \neq 0$ by equation (29), optimality requires $1 - g_K - \frac{t_K}{1-t_K}\zeta_K = 0$, which rearranges to equation (30). □

Proposition 2 is a mirror image of Proposition 1. The optimal capital tax takes the same form

[18] With employment fixed in every task, non-automatable inputs are fixed and each automatable input satisfies $X_j = \rho_j Y / \left(NR\right)$ by equation (5), so the production function implies $d\ln Y = \left(n/N\right)\left(d\ln Y - d\ln R\right)$, that is, $d\ln Y - d\ln R = -\left[N / \left(N-n\right)\right] d\ln R$. Differentiating capital demand $h_K K = nY / \left(NR\right) - \sum_{j\leq n} \omega_j L_j / \rho_j$ at fixed employment gives $d\ln K = \left[nY / \left(NRh_K K\right)\right]\left(d\ln Y - d\ln R\right) = -\frac{nY}{\left(N-n\right)Rh_K K} d\ln R$.

as the optimal labor tax, with the capital supply elasticity ζ_K in place of the participation elasticity η_i and the capitalists' welfare weight g_K in place of the workers' weight g_i. Once workers are compensated, taxing capitalists is a partial-equilibrium problem: the pass-through of the tax into the rental rate, the capital demand elasticity, and every other feature of the general-equilibrium environment drop out of the formula. At the joint optimum, every dollar of factor income—wage or capital—has a social value of exactly one: $g_i + \tau_i\eta_i / (1 - \tau_i) = 1$ in every task and $g_K + t_K\zeta_K / (1 - t_K) = 1$.

Two limit cases are instructive. If the government places a zero marginal welfare weight on capitalists ($g_K = 0$), the optimal capital tax is $t_K = 1/(1 + \zeta_K)$: the rate that maximizes the revenue extracted from capitalists. If capital supply is perfectly elastic ($\zeta_K \to \infty$)—the case of a small open economy facing a global capital market—the optimal capital tax is zero. In this case, machines are an intermediate input available at a fixed world price, and taxing them would distort production without extracting any revenue from capitalists. This is an instance of the production-efficiency theorem of Diamond and Mirrlees (1971).

4 Quantifying the Effects of Automation

We now quantify how much automation changes optimal tax rates. We proceed in three steps. First, we calibrate the model to the US wage distribution and to survey-based measures of automation exposure, which place the automation threshold just below the median worker. Second, we compare the optimal labor income tax in the automation economy to the partial-equilibrium benchmark of Saez (2002), and show that automation rotates the schedule toward greater progressivity. Third, we show how optimal taxes respond to the two forms of technological progress: broader automation, whereby more tasks become automatable, and deeper automation, whereby machines become cheaper.

4.1 Calibration Details

We calibrate the primitives of the model and then solve for its general equilibrium and optimal tax rules. The primitives are the occupation-specific skill levels $\omega_1, ..., \omega_N$, the fraction of workers in each occupation $h_1, ..., h_N$, the productivity of capital $\rho_1, ..., \rho_n$, total factor productivity A, the automatable share n/N, the capital supply parameters (h_K, χ, ζ_K), and the parameters governing labor supply and social preferences. The cross-sectional dispersion of wages is an equilibrium object, generated jointly by the dispersion in skills, the distribution of workers across occupations, labor supply decisions, and the technology of automation. The rental rate of capital is likewise an equilibrium object, determined jointly with wages and employment by capital market clearing.[19]

We assume that skills are log-normally distributed, with a median normalized to one and a standard deviation of log skill σ_ω. We divide the population into $N = 1,000$ occupations that partition this distribution in order of skill, $\omega_1 < \cdots < \omega_N$, where ω_i is the median within bin i of the underlying continuous distribution. We calibrate bins of unequal width such that occupation i holds a population share $h_i \propto i^{-\alpha}$, with $\sum_i h_i = 1$. The shape parameter α determines how fast the population of workers declines across occupation ranks. Given the number of automatable occupations n (along with other parameters of the model), α governs the fraction of workers exposed to automation. We describe the baseline calibration of n, α, and σ_ω below.

The skill levels determine wages in automatable (bottom) occupations via $w_i = R\omega_i/\rho_i$ in condition (5), while population sizes h_i govern wages in non-automatable (top) occupations. Three considerations motivate the power-law decline in occupation size. First, it reflects that high skill is scarce, so that the higher-skill occupations are the smaller ones. Second, this scarcity is what produces wage dispersion at the top: in non-automatable occupations, the equilibrium wage satisfies $w_i L_i = Y/N$ in condition (4), so that $w_i = Y/(NL_i)$ is inversely proportional to occupation size and the scarcer, higher-skill occupations command higher wages. It is thus the falling occupation size, not the skill levels per se, that generates the upper tail of the wage distribution. Third, the decline cannot be too steep: because the wage is inversely proportional to occupation size, an occupation made very small would command a very high wage and inflate dispersion at the top, so a modest exponent is required.

Table 1 summarizes the calibrated parameters and their empirical sources. We set $n = 256$ (i.e., $n/N = 0.256$) and calibrate the degree of skill dispersion σ_ω and the occupation-size exponent α to ensure that the model generates a standard deviation of log wages equal to 0.92 and a fraction of the workforce exposed to automation equal to 46%.[20] A standard deviation of log wages of 0.92 corresponds to the most recent measure in Song *et al.* (2019).[21] A fraction of exposed workers of

[19] When calibrating primitives to match empirical moments—as described in detail below—participation tax rates τ_i and the capital tax rate t_K are set at their approximate values in the current US tax code. The calibrated skill and technology parameters are invariant to the initial (unoptimized) tax system; the labor- and capital-supply scales (κ_i and $h_K\chi$) adjust mechanically to it.

[20] The calibrated parameter values are $\sigma_\omega = 0.70$ and $\alpha = 0.44$.

[21] They report a variance of log annual earnings among US wage earners rising from 0.652 in 1981 to 0.846 in 2013, corresponding to a standard deviation rising from 0.81 to 0.92. Because hours worked conditional on working are fixed in our model, the standard deviation of log wages among the employed coincides with that of log annual earnings.

46% matches the average automation exposure estimated in Muro *et al.* (2019).[22] The automatable share n/N can be grounded in the same data: occupations highly exposed to automation—defined as those with more than 70 percent of their task content automatable—comprise roughly one quarter of US occupations, consistent with our value of n/N.[23] We also report comparative statics with respect to the size of n/N.

The capital side of the model introduces three additional primitives: the number of capitalists h_K, the scale of capital supply χ, and the capital supply elasticity ζ_K. We set $\zeta_K = 1$ at baseline, following the evidence in Jakobsen *et al.* (2020), who estimate long-run elasticities of capital supply between 0.58 and 1.91.[24] We report results for capital supply elasticities between 0.5 and 5. The parameters h_K and χ matter for the equilibrium only through their product: capitalists receive no transfers and—as discussed below—carry no social welfare weight, so capitalists matter for optimal policy only through aggregate capital tax revenue and hence aggregate capital supply $h_K K = h_K \chi \left[(1-t_K)\, R\right]^{\zeta_K}$. The product $h_K\chi$ therefore adds a single scale to the calibration.

We take the productivity of capital to be homogeneous across automatable tasks, $\rho_i = \rho$, so that the capital cost per efficiency unit R/ρ—the relative price of capital—is a single ratio. Because the rental rate clears the capital market, this ratio is an equilibrium object rather than a fixed parameter. Automation is adopted in task i only if it lowers the cost of producing task services, that is, only if $R\omega_i/\rho \leq w_i^L$, where w_i^L is the wage that would prevail under labor-only production. Because automatable occupations are populous, the labor-only wage $Y/\left(NL_i\right)$ is relatively low, and adoption requires the equilibrium rental rate to be low enough that capital still undercuts labor even at that low wage. We therefore calibrate the capital supply scale $h_K\chi$ and total factor productivity A jointly, such that automation is adopted throughout the automatable range, with the most labor-intensive automatable task retaining a small positive capital share,[25] and such that

[22] Averaging across occupations (weighted by employment), their Table 1 reports an aggregate automation exposure share of 46 percent.

[23] Muro *et al.* (2019) classify occupations with more than 70 percent of their task content automatable as highly exposed; such occupations hold 25 percent of US employment. The report does not directly state the corresponding fraction of occupations, but the distribution of automation exposure across detailed occupation codes within each occupational family (their Figure 4), combined with the number of detailed occupations per family in the Standard Occupational Classification, implies that highly exposed occupations make up roughly one quarter of the more than 800 detailed US occupations.

[24] These estimates refer to the simulated long-run capital supply elasticities in Table III of Jakobsen *et al.* (2020). For the top 1% wealthiest households (and assuming a gross capital return of 5 percent), they estimate a 30-year elasticity of taxable wealth with respect to the net-of-tax return equal to 1.15.

[25] We require that the capital share in every automatable occupation is above a minimum value k^*, i.e. $RK_i/\left(w_iL_i + RK_i\right) \geq k^*\ \forall i \leq n$. Specifically, we set $k^* = 0.1$ in our calibration. The requirement of a capital share of at least 10% is motivated by the idea that, with fixed costs of automation adoption, firms will introduce machines only if their optimal factor share is sufficiently above zero.

the capital cost per efficiency unit equals $R/\rho = 0.9$. The implied capital share of automatable value added is approximately 0.54.[26] Adoption pushes the automatable wages below the non-automatable range and opens an endogenous gap in the wage distribution at the automation threshold. A lower relative price of capital R/ρ deepens automation, raising the capital share of the automatable tasks and widening this gap, and we vary this ratio below, through the scale of capital supply, to trace out the consequences of more aggressive automation.

The extensive margin labor supply elasticities η_i are governed by the distributions of fixed work costs q across occupations. We assume that, in every occupation i, fixed costs are distributed as a power function with the CDF $P_i(q) = (q/\kappa_i)^{\eta}$ on the support $[0, \kappa_i]$, where κ_i is an occupation-specific scale. The employment rate is then $P_i(\bar{q}_i) = ((1-\tau_i)\, w_i/\kappa_i)^{\eta}$. We calibrate the scales κ_i such that the equilibrium employment rate equals 85 percent in every occupation, broadly in line with the employment-population ratio among prime-age US workers (currently 80.2 percent). The extensive margin elasticity is assumed to be homogeneous across tasks: $\eta_i = \eta\ \forall i$. We set $\eta = 0.5$ at baseline, which falls within the range of extensive margin elasticities surveyed by Chetty *et al.* (2013). It is straightforward to allow for heterogeneous elasticities—realistically, declining η_i across skill levels—by specifying $P_i(q) = (q/\kappa_i)^{\eta_i}$.

Social preferences take the CRRA form $\Psi[u] = u^{1-\gamma}/(1-\gamma)$, where the coefficient of relative risk aversion γ governs the planner's taste for redistribution: $\gamma = 0$ is laissez-faire, $\gamma \to \infty$ is Rawlsian, and $\gamma = 1$ is the logarithmic case. We take $\gamma = 1$ as our baseline. The social marginal welfare weights are obtained from $\Psi'(u)/\mu = u^{-\gamma}/\mu$, averaged over employed individuals in each occupation to yield the weights in automatable and non-automatable tasks ($g_{i\leq n}$ and $g_{i>n}$, respectively) that enter the optimal tax rules. The average welfare weight across all workers—employed and unemployed together—equals one, which pins down the demogrant $T(0)$. We set the social marginal welfare weight on capitalists equal to zero, $g_K = 0$, which is a reasonable approximation provided capitalists have sufficiently high incomes.[27] We set the aggregate revenue requirement from labor and capital taxes equal to zero ($G = 0$) such that the exercise is one of pure redistribution.

Given this calibration, we solve the model as a full general-equilibrium fixed point: the

[26] Because only automatable tasks use capital, the economy-wide capital share is n/N times this number, roughly a seventh of value added. Capital here is the machine stock that substitutes for labor in automatable tasks rather than the aggregate capital stock, so the model's aggregate factor shares are not comparable to national-accounts shares and are not calibration targets.

[27] Because the calibration described above fixes only the product $h_K\chi$, the number of capitalists h_K can be used as a free parameter to ensure that capital incomes are high and $g_K \approx 0$. This requires h_K to be sufficiently small.

participation taxes of Proposition 1, the capital tax of Proposition 2, the demogrant, the social marginal welfare weights, wages, employment, the rental rate, capital supply, and aggregate output are determined jointly. We compare the optimal tax system of our general-equilibrium automation economy with the partial-equilibrium benchmark of Saez (2002). We implement this benchmark as an economy with a linear labor-only technology in which wages equal skills, $w_i = \omega_i$, in every task.

4.2 Simulation Results

Figure 1 presents a central result of the paper. Panel A shows the optimal participation tax schedule in our general-equilibrium automation economy, alongside the partial-equilibrium benchmark of Saez (2002). The schedule in our model has three defining features. First, the optimal policy at the bottom is a large EITC. Optimal participation tax rates are negative through roughly the bottom quarter of the wage distribution, with a work subsidy reaching 77 percent of earnings in the lowest-wage occupations. Second, optimal participation tax rates at the top are high. Through the upper half of the distribution, tax rates rise gradually from about 38 to 47 percent. Third, the automatable and non-automatable regions are separated by a sharp discontinuity in the tax rate, located just below the median of the wage distribution. The discontinuity results from the discrete jump in wages at the automation threshold. The tax formula in Proposition 1 is the same on each side of the threshold, but the social welfare weight falls discretely as the wage jumps, and the optimal tax rate jumps with it. As discussed below, the wage discontinuity is larger than the tax discontinuity such that the net-of-tax wage rate $(1 - \tau)\, w$ increases discretely at the threshold.

We compare this schedule to the Saez benchmark, also shown in Panel A. As described above, this benchmark is a labor-only economy with linear technology in which wages equal skills in all tasks. Apart from the difference in production technologies, the two economies share the same primitives and are solved under the same optimal tax formula. The two schedules therefore differ only through the equilibrium wage distributions and welfare weights that enter the formula. Yet the differences are substantial. At the bottom, the automation economy features a much larger EITC: a work subsidy that reaches 77 percent of earnings vs 36 percent in the benchmark, and a subsidy region that covers the bottom quarter of the distribution vs 7 percent in the benchmark. Below the automation threshold, our tax rates are everywhere lower than in the benchmark. Above the threshold, the opposite applies. Our tax rates on workers sheltered from automation exceed the benchmark rates, apart from at the very top of the distribution. Automation thus rotates the optimal tax schedule—lower taxes on the exposed part of the workforce, higher taxes on the sheltered part—implying substantially stronger tax progressivity than in the partial-equilibrium benchmark.[28]

To further illustrate the different implications of the two economies, let us consider the profile of net-of-tax wages $(1-\tau)\,w$ in each of them. Panel B of Figure 1 plots these profiles at the optimal tax schedules, normalizing by the average net-of-tax wage within each economy. In the benchmark economy, the net-of-tax wage is smooth across the entire distribution. In the automation economy, by contrast, the net-of-tax wage jumps discretely at the automation threshold. This jump reflects offsetting discontinuities in the pre-tax wage and the optimal tax rate, where the latter was shown in Panel A. Absent switching between the automatable and non-automatable sectors, this discontinuity is not a notch in any worker's budget set, and it therefore does not create the behavioral responses typically associated with tax notches (Kleven and Waseem 2013; Kleven 2016). If workers could switch, the discontinuity would constitute a real notch—one that rewards retraining and occupational upgrading.[29] The empirical evidence suggests that this margin is inelastic (Acemoglu and Restrepo 2020), indicating that our fixed-occupation assumption provides a reasonable, albeit stylized, approximation of the real world.

Figure 2 shows where automation takes place in equilibrium. It plots the labor share of task production, $w_iL_i/\left(w_iL_i+RK_i\right)$, across wage percentiles at the optimal tax system. The automation threshold is located just below the median of the wage distribution, but the intensity of automation below this threshold is an equilibrium outcome. The figure shows that the labor share falls from 100 percent in the lowest-wage occupations to about 45 percent at the threshold. The reason why the lowest-wage occupations feature a large labor share is that these occupations are relatively populous and heavily subsidized, both of which drive up labor supply and employment in equilibrium. At the very bottom of the distribution, these forces are strong

[28]The comparison in Figure 1 bundles two differences between the economies: general equilibrium vs fixed wages, and automation versus a labor-only technology. The pure effect of automation will be isolated in Figure 4 presented below, in which we compare the GE automation economy to a GE labor-only economy built from the same primitives. Both comparisons hold primitives, rather than the observed wage distribution, fixed across economies (see Scheuer and Werning 2017 on this distinction). Finally, with a constant participation elasticity, equation (25) implies $g_i = 1-\eta\tau_i/\left(1-\tau_i\right)$: any difference between two optimal tax schedules translates directly into a difference in welfare weights.

[29]Allowing for occupational upgrading, the model would predict excess bunching of workers in the marginal non-automatable occupation $n+1$, combined with missing mass in automatable occupations close to n. A general result from optimal tax theory is that, with positive behavioral elasticities, notches are not a feature of optimal policies (e.g., Mirrlees 1971; Kleven 2016). We therefore conjecture that a more general model with occupational upgrading would smooth the net-of-tax wage—as well as pre-tax wages and taxes separately—around the automation threshold, replacing the discontinuity with a steep but continuous segment. In this sense, our stylized model captures the optimal tax schedule of the general model in stark form.

enough to produce a corner solution: labor is sufficiently abundant that the equilibrium wage falls below the cost of performing the task with machines (at the corner), and the lowest automatable occupations end up not being automated at all under the optimal tax system. Medium-wage occupations near the threshold, by contrast, are less populous and more heavily taxed, leaving a large residual for machines to fill. Machine production is thus most intense around the middle of the wage distribution, peaking at the 41st percentile. Our model's equilibrium is therefore consistent with the argument that recent innovations in automation and AI pose a greater threat to middle-class, white-collar jobs than to low-skill jobs.

Figure 3 compares the taxation of workers and capitalists under different capital supply elasticities. We simulate the model for capital supply elasticities $\zeta_K \in \{0.5, 1, 2, 5\}$, calibrating each economy using the same approach as the baseline. The taxation of workers is relatively insensitive to the capital supply elasticity. In Panel A, the four participation tax schedules remain close to each other: above the automation threshold they differ by no more than about 3 percentage points, while below the threshold they fan out moderately—by up to about 11 percentage points. Every schedule retains the three defining features of Figure 1: a large EITC at the bottom, high tax rates at the top, and a discontinuity at the automation threshold. These patterns reflect the structure of Proposition 1: the optimal labor tax depends only on the welfare weights and the participation elasticity. The capital supply elasticity affects these weights mainly through the equilibrium wage effects of the optimal capital tax shown in Panel B: a more elastic capital supply implies a lower capital tax and hence cheaper capital, deeper automation, and somewhat deeper work subsidies in the automatable range.

The taxation of capitalists, on the other hand, depends strongly on the capital supply elasticity. Panel B shows the optimal capital tax rate: it falls from 67 percent at $\zeta_K = 0.5$ to 17 percent at $\zeta_K = 5$. With a zero welfare weight on capitalists, Proposition 2 prescribes $t_K = 1/\left(1 + \zeta_K\right)$, the rate that maximizes the revenue extracted from capitalists. When capital supply is inelastic, the tax base erodes slowly as the capital tax rate rises and the optimal rate is therefore high. As the elasticity rises, the tax base erodes more quickly and the optimal rate declines, vanishing in the limit of perfectly elastic capital supply. Capital tax revenue ranges from 21 percent of total revenue at $\zeta_K = 0.5$ to 8 percent at $\zeta_K = 5$.

Figures 4 and 5 examine the implications of technological progress in its two forms: broader automation, whereby more tasks become automatable (a larger n/N), and deeper automation, whereby capital becomes cheaper relative to labor (a lower R/ρ). Panel A of Figure 4 compares the baseline economy ($n/N = 0.26$) to an economy without automation ($n/N = 0$), an economy in which half of all tasks are automatable ($n/N = 0.50$), and an economy in which three quarters of all tasks are automatable ($n/N = 0.75$). The economy without automation is a general-equilibrium labor-only economy in which every wage clears its own labor market. Its optimal schedule features a subsidy region confined to the bottom tenth of the distribution and participation taxes that rise quickly, reaching 38 percent at the median. Automation reshapes this schedule. Machines depress the wages of exposed workers, raise their welfare weights, and pull down participation taxes throughout the automatable range. The reductions are large. At the baseline of $n/N = 0.26$, the subsidy region widens from the bottom tenth to the bottom quarter of the distribution, and the tax rate at the 10th percentile falls from $+3$ percent to -19 percent.[30] Broader automation extends the low-tax range further up the wage distribution: with half of all tasks automatable, the threshold discontinuity slides from just below the median to the 69th percentile, and with three quarters automatable, to the 86th percentile—placing a growing majority of the workforce under the lower automatable-range taxes.

Panel B of Figure 4 shows where automation takes place in these economies. It plots the equilibrium labor share of task production—as in Figure 2—for the three automation economies. Broader automation moves the threshold far to the right, but it barely moves the location of peak automation: the labor-share trough sits at the 41st percentile at baseline and moves only to the 39th and 38th percentiles as the automatable share rises to 0.50 and 0.75. Moreover, the peak machine share stays at roughly 56 percent throughout. Equilibrium automation thus peaks around the middle of the wage distribution regardless of how far up the skill distribution automatability extends. The location of the peak reflects two opposing forces: moving up the wage distribution, occupations become scarcer—leaving more of the task to machines—while workers become more skilled and cover more of their task. The first force dominates through the lower middle of the distribution, the second one higher up.

Figure 5 varies the depth of automation, lowering the capital cost per efficiency unit R/ρ from 0.90 to 0.75 and 0.60. Because the rental rate is endogenous, we implement this experiment by expanding the scale of capital supply at fixed total factor productivity: more abundant capital depresses the equilibrium rental rate and deepens automation. Cheaper capital depresses the technology-pinned wages at the bottom proportionally, and the optimal response focuses

[30] Above the automation threshold, by contrast, the baseline and labor-only schedules agree within about one percentage point: the pure effect of automation on optimal taxes is confined to the automatable range.

there. The work subsidy at the 5th wage percentile deepens from 32 to 45 to 61 percent of earnings, and at $R/\rho = 0.60$ the subsidy region extends through the bottom half of the wage distribution—the entire automatable range—while taxes above the threshold barely move. Consistent with Corollary 1, cheaper automation is optimally met with more generous support for low-wage work. Note also that the relative price of capital governs the size of the tax discontinuity: more expensive capital raises automatable wages toward the labor-only level, narrowing the wage gap and compressing the jump in the tax rate at the threshold.

The comparative statics in Figures 4 and 5 deliver one of the key lessons of our analysis. By the targeting logic of Diamond and Mirrlees (1971) discussed in Section 3, general-equilibrium effects leave no imprint on the optimal tax formulas: Propositions 1 and 2 retain their reduced-form, partial-equilibrium structure in full general equilibrium. But it would be wrong to conclude that the general-equilibrium effects of automation are therefore irrelevant for optimal taxation. Technological progress, whether broader or deeper automation, reshapes the equilibrium wage distribution and thereby the welfare weights that enter the formulas. Automation does not change the optimal tax formulas, but it changes the optimal tax rates substantially.

5 Conclusion

This paper has studied optimal income taxation in an economy shaped by automation. The model joins task-based production with extensive margin labor supply in the tradition of Saez (2002). Tasks are automatable from the bottom of the skill distribution up to a threshold, with workers and machines being perfect substitutes in automatable tasks. The threshold can take any value, potentially allowing for middle- and high-skill tasks to be automatable as well. The technology splits the labor market in two: below the automation threshold, wages are pinned by machines; above the threshold, wages clear occupation-specific labor markets without competition from machines. A class of capitalists supply the machines used in automatable tasks at a finite elasticity. Tax policy consists of occupation-specific participation taxes—the form a nonlinear earnings tax takes when each occupation pays a distinct wage—together with a linear tax on capital income.

Although tax changes reverberate through every labor market and the capital market, the optimal tax system takes a simple form: the participation tax in every occupation and the capital tax are given by canonical partial-equilibrium formulas in which equilibrium incidence plays no explicit role. We establish this result through a compensated perturbation approach in which each reform is packaged with compensating tax adjustments that absorb its incidence in every other market, so that only the reformed market enters the optimality condition. The economics behind the result is the targeting logic of Diamond and Mirrlees (1971): with a dedicated instrument for every occupation and for capitalists, and with labor responding only along the extensive margin, the government can replicate directly any redistribution that equilibrium prices carry out.

The invariance of the tax *formulas* does not translate into invariance of the tax *rates*: automation shapes employment, wages, and the rental rate, which feed into the welfare weights that enter the formulas. We quantify the effects in a calibration to the US wage distribution and automation exposure. Machines push down wages—and thereby raise welfare weights—throughout the bottom half of the distribution, and the optimal work subsidy (an EITC) responds strongly: relative to the partial-equilibrium benchmark, the subsidy rate at the bottom deepens from roughly 36 to 77 percent of earnings, and its coverage expands from the bottom 7 percent to the bottom quarter of the distribution. Above the threshold, participation taxes rise from about 38 percent to 47 percent, exceeding the partial-equilibrium benchmark except at the very top. These prescriptions strengthen as automation advances: broader automation lowers taxes on the newly automatable occupations, while deeper automation—cheaper machines—makes the earnings subsidies larger and more widespread.

The simulations also reveal who is hurt the most by automation in equilibrium. These are not the lowest-skill workers, but rather the middle class. At the optimal tax system, machine intensity peaks around the 40th percentile of the wage distribution—with machines commanding a value-added share of roughly 56 percent—and remains substantial everywhere except at the very bottom and the very top. Although our model assumes that automatability runs from the bottom of the skill distribution up to a threshold, the equilibrium is consistent with the observation that recent innovations in automation and AI threaten middle-class jobs at least as much as low-skill jobs.

The optimal capital income tax, by contrast, is unaffected by automation. The sufficient statistics formula of Saez and Stantcheva (2018) carries over to general equilibrium, and because high-income capitalists carry zero welfare weight, the optimal rate depends only on the capital supply elasticity. Although taxing capital income in our model amounts to taxing machine inputs at a uniform rate, the rationale differs from that of the robot-tax literature—surplus extraction from capitalists rather than a failure of production efficiency—and the optimal rate far exceeds the robot taxes this literature prescribes.

Our results contrast with the endogenous-wage literature, where trickle-down effects typically temper optimal top tax rates (Stiglitz 1982; Rothschild and Scheuer 2013; Sachs *et al.* 2020). That mechanism relies on intensive margin responses that pool different worker types at the same earnings level, so that compressing pre-tax wages becomes an indirect tool for redistribution. In our extensive margin economy no such pooling occurs—no two occupations share the same earnings level—and the case for distorting taxes to manipulate equilibrium wages disappears. The general lesson is that automation changes optimal taxation not by overturning canonical tax rules, but by changing what those rules prescribe.

References

Acemoglu, Daron, and David Autor (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," in *Handbook of Labor Economics*, ed. by D. Card, and O. Ashenfelter, vol. 4, chap. 12. Elsevier B.V. 2

Acemoglu, Daron, David Autor, and Simon Johnson (2026): "Building Pro-Worker Artificial Intelligence," NBER Working Paper No. 34854. 1

Acemoglu, Daron, Andrea Manera, and Pascual Restrepo (2020): "Does the US Tax Code Favor Automation?," *Brookings Papers on Economic Activity*, Spring, 231–285. 5

Acemoglu, Daron, and Pascual Restrepo (2018): "The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review*, 108(6), 1488–1542. 2

——— (2020): "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 128(6), 2188–2244. 1, 2, 8, 25

——— (2022): "Tasks, Automation, and the Rise in U.S. Wage Inequality," *Econometrica*, 90(5), 1973–2016. 1

Ales, Laurence, Musab Kurnaz, and Christopher Sleet (2015): "Technical Change, Wage Inequality, and Taxes," *American Economic Review*, 105(10), 3061–3101. 5

Atkinson, Anthony, and Joseph Stiglitz (1976): "The Design of Tax Structure: Direct versus Indirect Taxation," *Journal of Public Economics*, 6(1-2), 55–75. 5

Autor, David (2024): "Applying AI to Rebuild Middle Class Jobs," NBER Working Paper No. 32140. 1

Beraja, Martin, and Nathan Zorzi (2025): "Inefficient Automation," *Review of Economic Studies*, 92, 69–96. 5

Chamley, Christophe (1986): "Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives," *Econometrica*, 54(3), 607–622. 5

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber (2013): "Does Indivisible Labor Explain the Difference Between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities," *NBER Macroeconomics Annual*, 27, 1–56. 23, 35

COSTINOT, ARNAUD, AND IVÁN WERNING (2023): "Robots, Trade, and Luddism: A Sufficient Statistic Approach to Optimal Technology Regulation," *Review of Economic Studies*, 90, 2261–2291. 1, 5, 11

DIAMOND, PETER (1998): "Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates," *American Economic Review*, 88(1), 83–95. 4

DIAMOND, PETER, AND JAMES MIRRLEES (1971): "Optimal Taxation and Public Production I: Production Efficiency," *American Economic Review*, 61(1), 8–27. 3, 4, 5, 12, 15, 16, 20, 28, 29

DIXIT, AVINASH, AND VICTOR NORMAN (1986): "Gains from Trade without Lump-Sum Compensation," *Journal of International Economics*, 21(1-2), 111–122. 11

GUERREIRO, JOÃO, SERGIO REBELO, AND PEDRO TELES (2022): "Should Robots Be Taxed?," *Review of Economic Studies*, 89, 279–311. 1, 5

HÉMOUS, DAVID, AND MORTEN OLSEN (2022): "The Rise of the Machines: Automation, Horizontal Innovation, and Income Inequality," *American Economic Journal: Macroeconomics*, 14(1), 179–223. 2

JACQUET, LAURENCE, AND ETIENNE LEHMANN (2025): "Generalized Production Efficiency," THEMA Working Paper No. 2025-08. 2, 4, 11

JAKOBSEN, KATRINE, KRISTIAN JAKOBSEN, HENRIK KLEVEN, AND GABRIEL ZUCMAN (2020): "Wealth Taxation and Wealth Accumulation: Theory and Evidence from Denmark," *The Quarterly Journal of Economics*, 135(1), 329–388. 4, 5, 22, 35

JONES, CHARLES, AND CHRISTOPHER TONETTI (2026): "Past Automation and Future A.I.: How Weak Links Tame the Growth Explosion," Working Paper. 2

JUDD, KENNETH (1985): "Redistributive Taxation in a Simple Perfect Foresight Model," *Journal of Public Economics*, 28(1), 59–83. 5

KAPLOW, LOUIS (2006): "On the Undesirability of Commodity Taxation Even When Income Taxation is not Optimal," *Journal of Public Economics*, 90, 1235–1250. 11

KLEVEN, HENRIK (2016): "Bunching," *Annual Review of Economics*, 8, 435–464. 25

KLEVEN, HENRIK, CLAUS KREINER, AND EMMANUEL SAEZ (2009): "The Optimal Income Taxation of Couples," *Econometrica*, 77(2), 537–560. 4

KLEVEN, HENRIK, AND MAZHAR WASEEM (2013): "Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence From Pakistan," *The Quarterly Journal of Economics*, 128(2), 669–723. 25

LOEBBING, JONAS (2022): "Redistributive Income Taxation with Directed Technical Change," Working Paper. 5

MIRRLEES, JAMES (1971): "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 38(2), 175–208. 4, 25

MOLL, BENJAMIN (2025): "Applications: Automation & AI, Germany without Russian Gas," Lecture Slides 2025/26. 1, 6

MOLL, BENJAMIN, LUKASZ RACHEL, AND PASCUAL RESTREPO (2022): "Uneven Growth: Automation's Impact on Income and Wealth Inequality," *Econometrica*, 90(6), 2645–2683. 2

MURO, MARK, ROBERT MAXIM, AND JACOB WHITON (2019): "Automation and Artificial Intelligence," Brookings Report. 1, 22, 35

NAITO, HISAHIRO (1999): "Re-examination of Uniform Commodity Taxes under a Non-linear Income Tax System and Its Implication for Production Efficiency," *Journal of Public Economics*, 71(2), 165–188. 5, 16

PIKETTY, THOMAS, AND EMMANUEL SAEZ (2013): "Optimal Labor Income Taxation," in *Handbook of Public Economics*, ed. by A. Auerbach, R. Chetty, M. Feldstein, and E. Saez, vol. 5, chap. 7. Elsevier B.V. 4

ROTHSCHILD, CASEY, AND FLORIAN SCHEUER (2013): "Redistributive Taxation in the Roy Model," *The Quarterly Journal of Economics*, 128(2), 623–668. 4, 15, 16, 30

SACHS, DOMINIK, ALEH TSYVINSKI, AND NICOLAS WERQUIN (2020): "Nonlinear Tax Incidence and Optimal Taxation in General Equilibrium," *Econometrica*, 88(2), 469–493. 4, 15, 30

SAEZ, EMMANUEL (2001): "Using Elasticities to Derive Optimal Tax Rates," *Review of Economic Studies*, 68(1), 205–229. 2, 4, 11

——— (2002): "Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses," *Quarterly Journal of Economics*, 117(3), 1039–1073. 1, 2, 3, 4, 8, 11, 12, 15, 20, 24, 28, 36

——— (2004): "Direct or Indirect Tax Instruments for Redistribution: Short-Run versus Long-Run," *Journal of Public Economics*, 88(3-4), 503–518. 4, 16

SAEZ, EMMANUEL, AND STEFANIE STANTCHEVA (2018): "A Simpler Theory of Optimal Capital Taxation," *Journal of Public Economics*, 162, 120–142. 5, 9, 29

SAEZ, EMMANUEL, AND GABRIEL ZUCMAN (2026): "Distributional Tax Analysis in Theory and Practice: Harberger Meets Diamond-Mirrlees," *Journal of Public Economics*, 259. 4

SCHEUER, FLORIAN, AND IVÁN WERNING (2017): "The Taxation of Superstars," *The Quarterly Journal of Economics*, 132(1), 211–270. 25

——— (2018): "Mirrlees Meets Diamond-Mirrlees: Simplifying Nonlinear Income Taxation," Working Paper. 4, 16

SCHULZ, KARL, ALEH TSYVINSKI, AND NICOLAS WERQUIN (2023): "Generalized Compensation Principle," *Theoretical Economics*, 18(4), 1665–1710. 11

SONG, JAE, DAVID PRICE, FATIH GUVENEN, NICHOLAS BLOOM, AND TILL VON WACHTER (2019): "Firming Up Inequality," *The Quarterly Journal of Economics*, 134(1), 1–50. 21, 35

STIGLITZ, JOSEPH (1982): "Self-selection and Pareto Efficient Taxation," *Journal of Public Economics*, 17(2), 213–240. 4, 15, 16, 30

THUEMMEL, UWE (2023): "Optimal Taxation of Robots," *Journal of the European Economic Association*, 21(3), 1154–1190. 1, 5

VERGARA, DAMIÁN (2026): "Minimum Wages and Optimal Redistribution: The Role of Firm Profits," Working Paper. 2, 11

ZEIRA, JOSEPH (1998): "Workers, Machines, and Economic Growth," *Quarterly Journal of Economics*, 113(4), 1091–1117. 1, 2, 6

TABLE 1: CALIBRATED PARAMETERS

Parameters	Values	Sources and Targets
Panel A. External Evidence or Calibration		
Labor supply elasticity, η	0.50	Chetty *et al.* (2013)
Capital supply elasticity, ζ_K	1.00 (0.5, 2, 5)	Jakobsen *et al.* (2020)
Skill dispersion, σ_ω	0.70	SD of log wages $= 0.92$ (Song *et al.* 2019)
Occupation-size exponent, α	0.44	Automation exposure $= 46\%$ (Muro *et al.* 2019)
Fixed-cost scales, κ_i	Occupation-specific	Employment rate (EPOP) $= 85\%$
Automatable share, n/N	0.26 (0, 0.50, 0.75)	Highly automatable tasks (Muro *et al.* 2019)
Capital cost per efficiency unit, R/ρ	0.90 (0.75, 0.60)	Capital share in automatable sector $= 0.54$
Panel B. Normalizations and Settings		
Curvature of social welfare, γ	1	Logarithmic social welfare function
Welfare weight on capitalists, g_K	0	Typical top-of-distribution assumption
Number of occupations, N	1,000	Normalization
Median skill	1	Normalization
Revenue requirement, G	0	Pure redistribution
Minimum capital share, k^*	0.10	Threshold for machine adoption

Notes: This table summarizes the parameter values used in the numerical simulations. Panel A lists parameters that are either set from empirical evidence (behavioral elasticities) or calibrated to match US moments. When calibrating parameters to match US moments, tax rates are set equal to their approximate values in the current US tax code. Panel B lists other parameters that are normalized or assumed. Values displayed in parentheses are examined in comparative statics exercises.

FIGURE 1: OPTIMAL LABOR TAXES

A: Participation Taxes

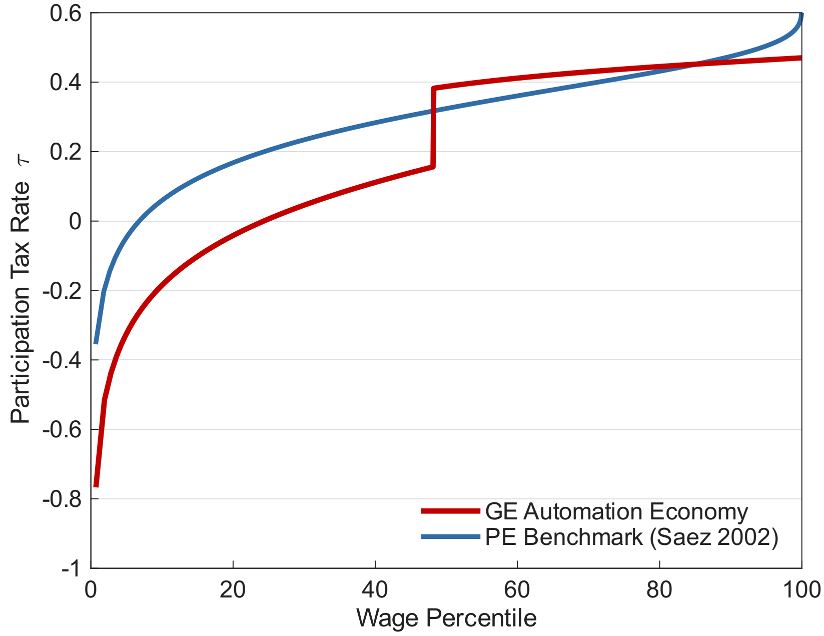

B: Net-of-Tax Wages

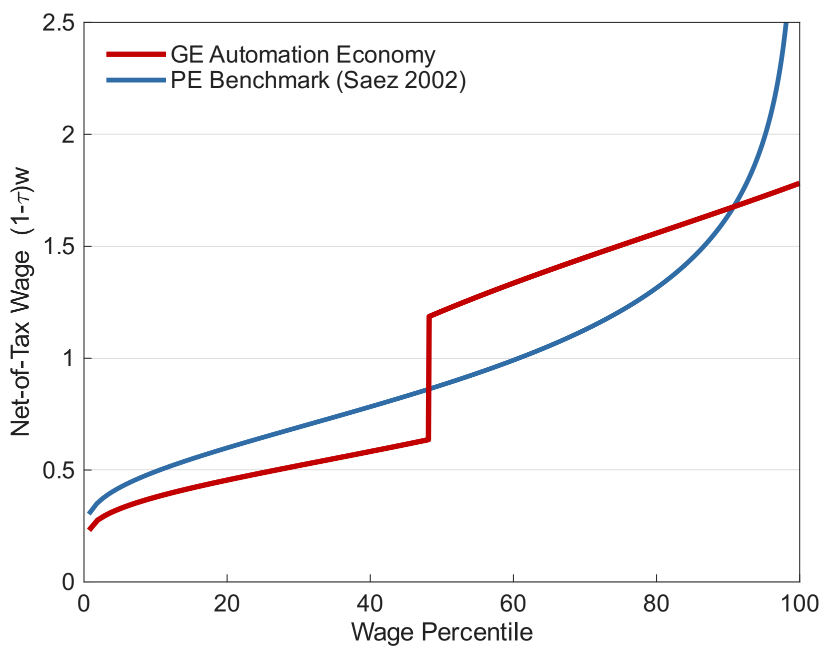

Notes: This figure compares optimal labor taxes in our general-equilibrium automation economy to those in the partial-equilibrium model of Saez (2002) without automation. The PE benchmark is based on a linear labor-only technology in which wages equal skills. Apart from the difference in production technologies, the two economies share the same primitives and are solved under the same optimal tax formula. Panel A plots the optimal participation tax rate τ across wage percentiles. Panel B plots the implied net-of-tax wage $(1-\tau)\,w$, normalizing each economy by its own average net-of-tax wage. The automation economy features a more progressive tax schedule: a larger earnings subsidy (EITC) at the bottom and higher rates through most of the upper half.

FIGURE 2: AUTOMATION IN EQUILIBRIUM

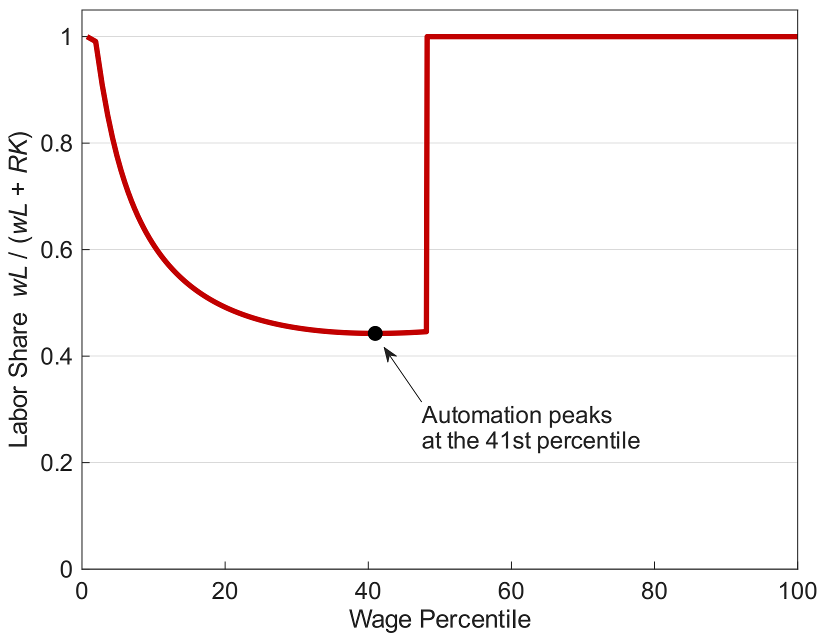

Notes: This figure plots the labor share of task production, $w_i L_i / (w_i L_i + RK_i)$, across wage percentiles at the optimal tax system. The labor share falls from 100 percent in the lowest-wage occupations to about 45 percent at the automation threshold, before jumping back to 100 percent above the threshold. The lowest-wage occupations, while potentially automatable, are at a corner solution without automation in equilibrium. Middle-class occupations are the most automated, with machine intensity peaking at the 41st percentile of the wage distribution.

FIGURE 3: OPTIMAL LABOR VS CAPITAL TAXES

EFFECT OF CAPITAL SUPPLY ELASTICITY ζ_K

A: Participation Taxes

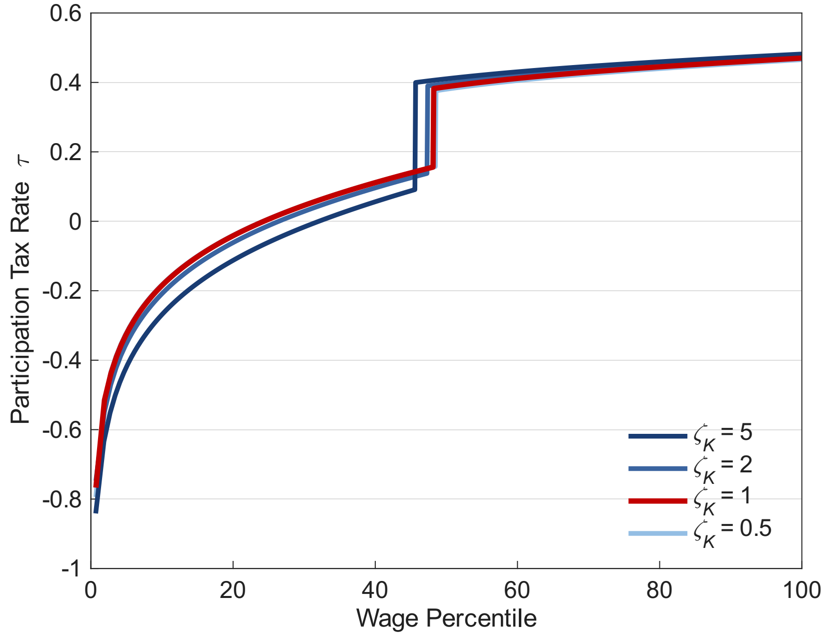

B: Capital Tax

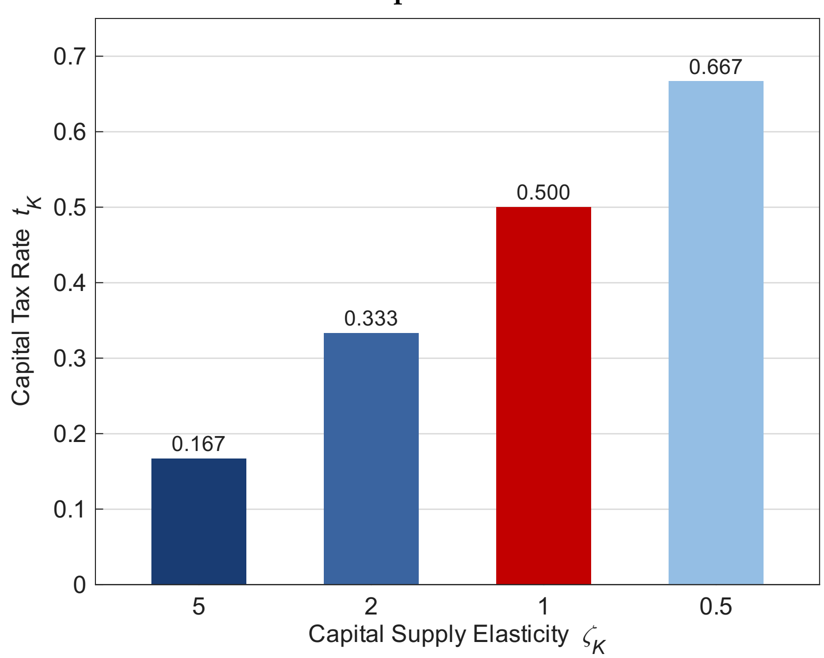

Notes: This figure shows the optimal tax system at capital supply elasticities $\zeta_K = 5, 2, 1, 0.5$, where $\zeta_K = 1$ corresponds to the baseline calibration. Each economy is calibrated using the same approach as the baseline economy. Panel A plots the optimal participation tax schedules across elasticity scenarios: these differ mainly in the automatable range, through the effect of the optimal capital tax on the equilibrium rental rate. Panel B plots the optimal capital income tax rate across the same scenarios, with shades matching the corresponding schedules in Panel A. Given the inverse elasticity rule in Proposition 2, the optimal capital tax rate falls sharply with the capital supply elasticity.

FIGURE 4: BROADER AUTOMATION

EFFECT OF AUTOMATABLE SHARE n/N

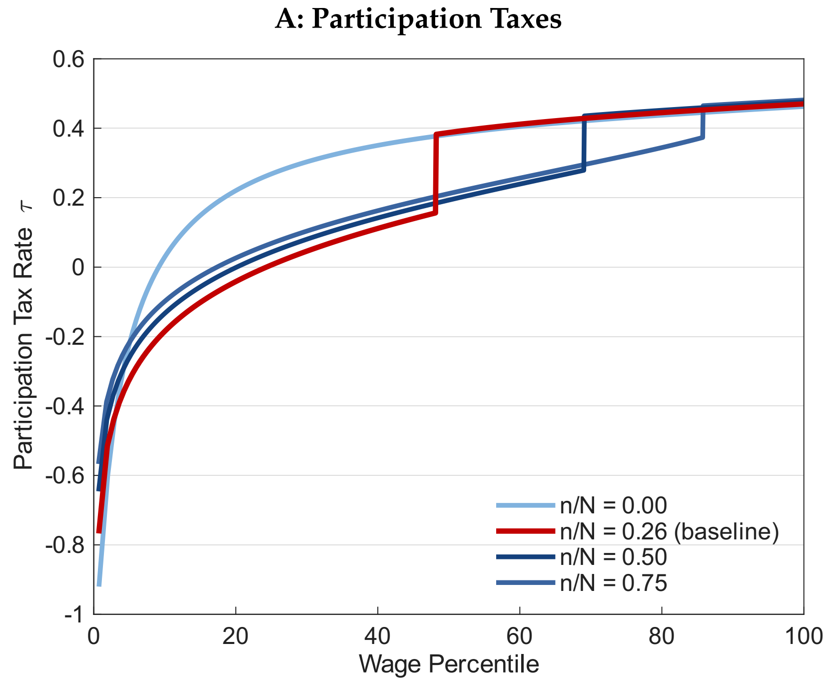

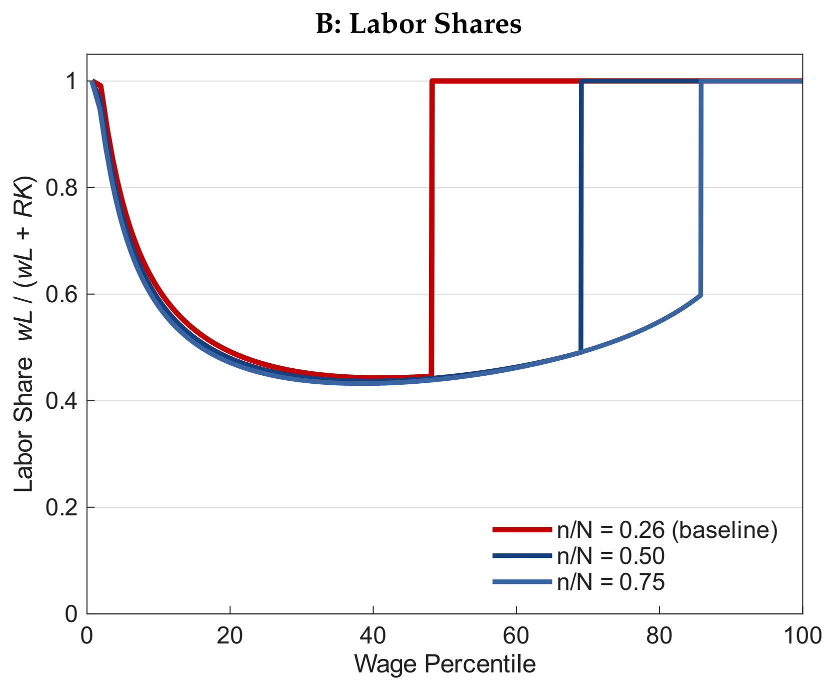

Notes: This figure shows the effect of broader automation—a larger automatable share n/N—on the optimal tax system. Panel A plots the optimal participation tax schedules for automatable shares $n/N = 0$, 0.26 (baseline), 0.50, and 0.75. The economy without automation is a general-equilibrium labor-only economy in which every wage clears its own labor market. Panel B plots the equilibrium labor share of task production for the three automation economies. In each automation economy, the capital supply scale $h_K\chi$ is recalibrated to maintain our adoption-margin target ($k^* = 0.1$). The fixed-cost scales κ_i are recalibrated in every economy to match our employment-rate target of 85 percent. A larger automatable share extends the low-tax automatable range further up the wage distribution, while equilibrium automation continues to peak around the middle of the distribution.

FIGURE 5: DEEPER AUTOMATION

EFFECT OF CAPITAL COST PER EFFICIENCY UNIT R/ρ

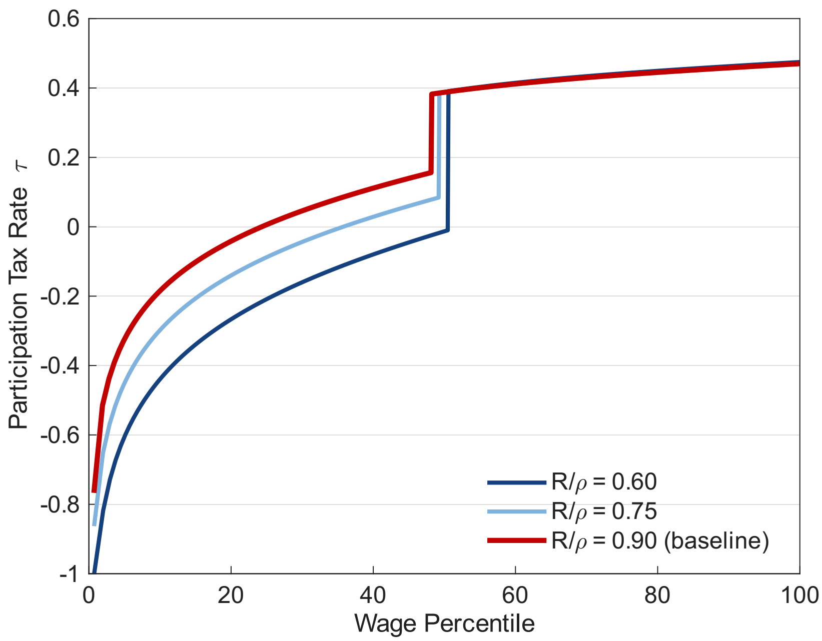

Notes: This figure shows the effect of deeper automation—cheaper capital in automatable tasks—on the optimal participation tax schedule. The capital cost per efficiency unit R/ρ is lowered from its baseline value of 0.90 to 0.75 and 0.60 by expanding the capital supply scale $h_K\chi$ at fixed total factor productivity. In each economy, the fixed-cost scales κ_i are recalibrated to the 85 percent employment-rate target. Cheaper capital raises optimal earnings subsidies (EITC) at the bottom and widens their reach, while taxes above the automation threshold barely move.