DARON ACEMOGLU
Massachusetts Institute of Technology

ANDREA MANERA
Massachusetts Institute of Technology

PASCUAL RESTREPO
Boston University

Does the US Tax Code Favor Automation?

ABSTRACT We argue that the US tax system is biased against labor and in favor of capital and has become more so in recent years. As a consequence, it has promoted levels of automation beyond what is socially desirable. Moving from the US tax system in the 2010s to optimal taxation of capital and labor would raise employment by 4.02 percent and the labor share by 0.78 percentage point and restore the optimal level of automation. If moving to optimal taxes is infeasible, more modest reforms can still increase employment by 1.14–1.96 percent, but in this case it is also beneficial to impose an additional automation tax to reduce the equilibrium level of automation. This is because marginal automated tasks do not bring much productivity gains but displace workers, reducing employment below its socially optimal level. We additionally show that reducing labor taxes or combining lower capital taxes with automation taxes can increase employment much more than the uniform reductions in capital taxes enacted between 2000 and 2018.

Conflict of Interest Disclosure: Daron Acemoglu is Institute Professor at the Massachusetts Institute of Technology (MIT) and a member of the Temporary Nomination Committee for the National Academy of Sciences; Andrea Manera is a graduate student in economics at MIT; Pascual Restrepo is an assistant professor of economics at Boston University and recently served as guest editor for the Journal of Labor Economics. Beyond these affiliations, the authors did not receive financial support from any firm or person for this paper or from any firm or person with a financial or political interest in this paper. They are currently not officers, directors, or board members of any organization with an interest in this paper. No outside party had the right to review this paper before circulation. The views expressed in this paper are those of the authors, and do not necessarily reflect those of MIT, Boston University, or the National Academy of Sciences.

The last three decades have witnessed a declining share of labor in national income, stagnant median real wages, and lower real wages for low-skill workers in the US economy (Elsby, Hobijn, and Şahin 2013; Acemoglu and Autor 2011; Karabarbounis and Neiman 2014). The labor share in nonfarm private businesses declined from 63 percent in 1980 to 56 percent in 2017, while median real wages grew only by 16 percent (as compared to GDP per capita which doubled during the same period), and the real wages of male workers with a high school diploma fell by 6 percent between 1980 and 2017. In the meantime, production processes have become increasingly automated, as computerized numerical control machines, industrial robotics, specialized software, and, lately, artificial intelligence technologies have spread rapidly throughout the economy. For instance, the US economy had a total of 2.5 industrial robots per thousand workers in manufacturing in 1993, and this number rose to 20 by 2019 (Acemoglu and Restrepo 2020). From a base of essentially zero in the mid-2000s, the share of vacancies posted for artificial intelligence–related activities increased to 0.75 percent by 2018 (Acemoglu and others 2020).

A common perspective among economists is that even if automation is contributing to the decline in the labor share and the stagnation of wages, the adoption of these technologies is beneficial, and any adverse consequences should be dealt with using redistributive policies and investments in education and training. But could it be that the extent of automation is excessive, meaning that businesses are adopting automation technologies beyond the socially optimal level? If this is the case, the policy responses to this trend need to be rethought.

In this paper, we show that the US tax system is biased against labor and as a result generates excessive automation and suboptimally low levels of employment and labor share. We first introduce a task-based model of automation, building on Acemoglu and Restrepo (2018, 2019a, 2019b) and Zeira (1998), to study the interplay between taxes and automation. Our first theoretical result establishes that optimal capital and labor taxes depend on the inverse supply elasticities of these factors and labor market frictions. Consistent with Diamond and Mirrlees (1971), once capital and labor taxes are set optimally, there is no reason to distort equilibrium automation decisions. Intuitively, optimal taxes undo any distortions and ensure that market prices reflect the social values of capital and labor. Automation decisions based on these prices are therefore optimal.[1]

---

[1] We assume that the labor market friction is common across tasks. When labor market frictions affect tasks differentially, there is an additional reason for excessive automation,

Yet this result does not imply that equilibrium automation decisions are optimal at arbitrary capital and labor taxes. Our second theoretical result shows that if a tax system is biased against labor and in favor of capital—that is, taxes on labor are too high and taxes on capital are too low—then reducing automation at the margin improves welfare. We show that this reduction can be achieved with an *automation tax*, which is an additional tax on the use of capital in tasks where labor has a comparative advantage. An automation tax is beneficial because reducing automation below its equilibrium level has second-order costs and first-order benefits. The costs are second-order as the productivity gains from automating marginal tasks are small, or, equivalently, the automation of marginal tasks corresponds to "so-so automation" in the terminology of Acemoglu and Restrepo (2019a, 2019b). But when the tax system is biased against labor and thus the level of employment is below the social optimum, limiting automation and avoiding the resulting displacement of labor has first-order benefits.

A common intuition is that if taxes are distorted, then the best policy remedy is to correct these distortions. Hence, if a tax system treats capital too favorably, we should directly tackle this distortion and increase capital taxes. We demonstrate that this intuition does not always apply in the presence of other constraints—for example, a lower bound on labor taxes. Our third theoretical result shows that a tax system distorted in favor of capital may call for reducing equilibrium automation even if raising capital taxes is possible. In fact, when moving to the unconstrained optimum is not feasible, constrained optimal policy may involve lower capital taxes in addition to a reduced level of automation because this combination avoids the displacement of workers from marginal tasks while ensuring that capital gets used intensively in tasks that are (and should be) automated. Both of these margins contribute to raising employment and welfare. This result underscores the importance of distinguishing between the choice of capital intensity in tasks where capital has a comparative advantage and automation, which involves the substitution of capital for labor in additional tasks. An automation tax is beneficial precisely because it does not reduce capital intensity uniformly but discourages the automation of marginal tasks.

Armed with these theoretical results, we turn to measuring effective taxes on capital and labor in the United States and comparing them to their optimal counterparts. We find that labor is much more heavily taxed than capital, and this difference has increased in recent years. Effective labor

---

as shown in our companion paper (Acemoglu, Manera, and Restrepo in progress), and in that case, distorting automation may be beneficial even when taxes are set optimally.

taxes in the United States are in the range of 25.5–33.5 percent. Effective capital taxes on software and equipment, on the other hand, are much lower, 10 percent in the 2010s and 5 percent after the 2017 tax reforms, though they used to be about 20 percent in 2000.[2] About half of this decline is due to the greater generosity of depreciation allowances.

Using plausible ranges for the elasticities of the capital and labor supply and estimates of labor market distortions, we find that the US tax system is biased against labor. In fact, our baseline estimates suggest that optimal labor taxes are lower than capital taxes—an 18.22 percent labor tax compared to a 26.65 percent capital tax. Optimal taxes are lower for labor than for capital because empirically plausible ranges of supply elasticities for capital and labor are similar, but employment is further distorted by labor market imperfections. Moving from the current tax system to optimal taxes would reduce the range of automated tasks by 4.1 percent and increase employment by 4.02 percent and the labor share by 0.78 percentage point.[3]

Our quantitative results show that, as in our theory, reducing automation is socially beneficial. Specifically, with no changes in capital and labor taxes, an automation tax of 10.15 percent—which implies that only tasks where the substitution of labor for capital reduces unit costs by more than 10.15 percent are automated—maximizes welfare and raises employment by 1.14 percent and the labor share by 1.93 percentage points. If capital taxes can be reduced as well, then a 12.9 percent automation tax combined with a reduction in capital taxes from 10 percent to 8.39 percent would achieve even higher welfare gains and increase employment by 1.59 percent and the labor share by 2.44 percentage points. We further show that tax reforms that involve lower labor taxes or combine lower capital taxes with an automation tax would have increased welfare and expanded employment much more than the uniform capital tax reductions enacted between 2000 and 2018.

We conclude with two extensions. First, we show that if human capital is endogenous, the asymmetric treatment of labor becomes more costly as it distorts human capital investments, leading to even lower optimal taxes on labor and more excessive automation under the current system. Second, we consider endogenous development of automation technologies,

---

[2] Acemoglu and Restrepo (2019b) document that technological changes in the four decades after World War II involved less automation and more rapid advances in technologies that increased human productivity (such as the creation of new tasks for workers) than has been the case recently. Though there are other reasons for why the direction of technology altered, the lower taxation of equipment and software capital may have also played a role.

[3] Despite these large changes in employment, the increase in welfare is given by a Harberger's triangle and is thus smaller—0.38 percent in consumption-equivalent terms.

which come at the expense of other types of innovations that are more beneficial for labor. In this case, there are reasons for not just preventing excessive adoption of automation technologies but also redirecting technological change away from further automation (and this is true even with optimal taxes on capital and labor).

Our paper is related to several classic and recent articles, though, to the best of our knowledge, no other paper investigates whether the US tax system favors automation.

First, there is an emerging literature on redistribution and taxation of automation technologies (Guerreiro, Rebelo, and Teles 2017; Thuemmel 2018; Costinot and Werning 2018). This literature studies whether adverse distributional effects of automation call for taxes on automation technologies. Our paper is complementary to this literature, as it focuses on situations in which the tax system is biased against labor and the key policy objective is to raise employment (not to redistribute income).

Second, our paper is related to the literature on optimal capital taxation (e.g., Atkinson and Stiglitz 1972, 1976; Judd 1985; Chamley 1986; Straub and Werning 2020). Our contribution is to show that in both two-period and infinite-horizon settings, provided that the government must run a balanced budget at each date, optimal taxes are given by the same inverse-elasticity formulas (with an additional term adjusting for labor market frictions). In contrast, this literature typically assumes that the government can freely accumulate assets and concludes that zero capital taxation is optimal in the long run. Straub and Werning (2020) show that if the supply of capital is not perfectly elastic (which means utility is not time-additive), then the government accumulates sufficient assets so that both capital and labor face zero taxes in the long run. We demonstrate in the online appendix that in the presence of labor market frictions, the same reasoning leads to a subsidy to labor. Thus, in the empirically relevant case of a finite supply elasticity of capital, even without the balanced budget assumption, the US tax system with low capital taxes and high labor taxes is far from optimal.

Third, our paper relates to the literature on the effects of tax reforms on investment and labor market outcomes. A branch of this literature estimates the differential responses of investment across firms facing different taxes (Goolsbee 1998; Hassett and Hubbard 2002; Edgerton 2010; Yagan 2015).[4]

---

[4] Modal results in this literature find investment elasticities with respect to the keep rate (one minus the tax rate) between 0.5 and 1. More recent work by House and Shapiro (2008) documents a larger investment response and argues that this was due to the temporary nature of the bonus, while Zwick and Mahon (2017) estimate investment elasticities with respect to the keep rate that are around 1.5 for most firms.

However, these estimates are informative about firms' demand for capital, not about the (long-run) elasticity of the supply of capital, which is the relevant object for optimal taxes. We discuss below estimates of this elasticity based on the response of the supply of capital to wealth and capital income taxes (see Kleven and Schultz 2014; Zoutman 2018; Brülhart and others 2016; Jakobsen and others 2020; Durán-Cabré, Esteller-Moré, and Mas-Montserrat 2019). More closely connected to our work is a branch of this literature on the labor market implications of tax reforms. Suárez Serrato and Zidar (2016) exploit the incidence of tax changes across US counties and estimate that a 1 percent increase in the keep rate of corporate taxes raises employment by 3.5 percent and wages by 0.8 percent and that workers bear 35 percent of the incidence. Garrett, Ohrn, and Suárez Serrato (2020) compare counties at the 75th percentile of exposure to bonus depreciation allowances to those at the 25th percentile and find a 2 percent increase in employment, no changes in wages, and a 3.3 percent increase in investment in response to the reform. These estimates point to a fairly elastic response of employment and a less than perfectly elastic response of capital in local labor markets (a perfectly elastic response of capital would cause workers to bear the full incidence).

Finally, our modeling of automation builds on Zeira (1998), Autor, Levy, and Murnane (2003), Acemoglu and Autor (2011), and most closely, Acemoglu and Restrepo (2018, 2019a, 2019b). The task-based framework is useful in our setting because it shows how automation (substituting capital for labor in tasks previously performed by humans) creates a displacement effect while automating marginal tasks generates limited productivity gains (because firms are approximately indifferent between automating these tasks or producing with labor). This combination of displacement effects and small productivity gains is at the root of our result that the planner would like to reduce automation at the margin when the tax system is biased against labor. Our framework also clarifies how policy can affect the level of automation and why taxing automation is not the same as taxing capital.

The rest of the paper is organized as follows. Section I introduces our conceptual framework and derives our theoretical results. Section II provides a detailed discussion of the US tax system and maps the complex US tax code into effective capital and labor income taxes. Section III then explores whether these taxes are biased and how they compare against optimal taxes. Section IV discusses two extensions of our framework, while section V concludes. The online appendix contains proofs of the results stated in the text, various theoretical generalizations, and further details for and robustness checks on our empirical work.

I. Conceptual Framework

This section presents our conceptual framework for evaluating the optimality of capital and labor taxes and the extent of automation. To facilitate the exposition, we focus on a two-period model and generalize our main results to an infinite-horizon setting in the online appendix.

I.A. Environment

There is a unique final good, produced at time t = 1 by combining a unit measure of tasks:

$$\left(\int_0^1 y(x)^{\frac{\lambda-1}{\lambda}} dx\right)^{\frac{\lambda}{\lambda-1}}.$$

Tasks are allocated between capital and labor and performed with the following task-level production function:

$$y(x) = \psi^{\ell}(x) \cdot \ell(x) + \psi^{k}(x) \cdot k(x), \tag{1}$$

where $\ell(x)$ is labor employed in task x, $k(x)$ is capital used in the production of task x, and $\psi^{\ell}(x)$ and $\psi^{k}(x)$ denote, respectively, the productivities of labor and capital in task x. We order tasks such that $\psi^{\ell}(x)/\psi^{k}(x)$ is nondecreasing and simplify the exposition by assuming that it is strictly increasing. We also suppose that when indifferent between producing a task with capital or labor, firms produce with capital. Therefore, there exists a threshold task θ such that tasks in $[0, \theta]$ are produced with capital and tasks in $[\theta, 1]$ are produced with labor. For now, there is no distinction between the adoption and the development of such technologies. We explore the implications of this distinction in section IV.B.

The household side is inhabited by a representative household that lives for two periods, $t = 0$ and $t = 1$. There is no production in period 0, but the representative household is endowed with $\bar{y}$ units of output. Out of this, it consumes c_0 and saves the remaining $k = \bar{y} - c_0$ units, which are allocated to producing capital. Capital is used during period 1, is subject to depreciation at the rate δ, and is rented to firms at the rental rate R, so that households earn an after-tax return of $(R - \delta) \cdot (1 - \tau^k)$. The period 1 budget constraint facing the household is

$$c \leq (1 + (R - \delta) \cdot (1 - \tau^k)) \cdot k + w \cdot (1 - \tau^{\ell}) \cdot \ell,$$

where R is the rental rate on capital paid by firms and w is the wage rate. Tax revenues are used for financing a fixed level of government expenditure, denoted by g.

The household chooses consumption and the supply of capital and hours to maximize

$$u(\bar{y} - k) + c - v(\ell).$$

Here, $u(\bar{y} - k)$ is a concave function representing the utility from consuming $\bar{y} - k$ units of output in period 0; c denotes the utility from consumption in period 1; and $v(\ell)$ is a convex function representing the disutility from working. Quasi linearity in period 1 is imposed for simplicity (see the online appendix for more general preferences).

We allow for various types of frictions in the labor market, modeled as introducing a wedge between the market wage and the representative household's marginal cost of supplying labor. We denote this wedge by $\varrho \geq 0$.[5]

Market clearing for capital and labor requires $k = \int_0^1 k(x)dx$ and $\ell = \int_0^1 \ell(x)dx$. To ensure uniqueness of optimal taxes below, we suppose that $u'(\bar{y} - k) \cdot k$ and $v'(\ell) \cdot \ell$ are convex. In addition, we assume that the equilibrium involves a positive net rate of return on investment. Finally, we denote by $\varepsilon^k(k)$ and $\varepsilon^\ell(\ell)$ the Hicksian elasticities of capital and labor. These are given by the response of capital and labor supply to a permanent percent change in the relevant keep rates (one minus the tax rates):

$$\varepsilon^k(k) = \frac{d\ln k}{d\ln(1-\tau^k)} = -\frac{u'(\bar{y}-k)-1}{u''(\bar{y}-k)\cdot k} \geq 0 \quad \text{and}$$

$$\varepsilon^\ell(\ell) = \frac{d\ln \ell}{d\ln(1-\tau^\ell)} = \frac{v'(\ell)}{v''(\ell)\cdot \ell} \geq 0.$$

As the equation for $\varepsilon^k(k)$ makes clear, the concavity of period –1 utility, $u(\bar{y} - k)$, ensures that the marginal rate of substitution between consumption today and tomorrow is increasing in k and thus the supply of capital is not perfectly elastic, otherwise $\varepsilon^k(k)$ would be infinite.[6]

---

[5] As shown in Acemoglu, Manera, and Restrepo (in progress), this wedge can be derived from bargaining between workers and firms or from efficiency wage considerations.

[6] A complementary reason for finite $\varepsilon^k(k)$ is that the technology for investment is convex (for example, the production of k units of capital requires $\phi(k)$ units of period 0 resources, where ϕ is strictly convex). If the profits from producing capital cannot be directly taxed, our optimal tax formulae apply regardless of whether $\varepsilon^k(k)$ reflects changes in the marginal rate of substitution between consumption today and tomorrow as a function of k or a convex investment technology.

Note that our formulation assumes that τ^k is a tax on net—after depreciation—returns, not gross returns, and our formula for $\varepsilon^k(k)$ computes it as the elasticity of capital to a percent change in one minus the net tax on capital.

I.B. Equilibrium

Given taxes $\{\tau^k, \tau^\ell\}$ and the labor wedge ϱ, a market equilibrium is defined by factor prices $\{w, R\}$, a tuple of current output, consumption, capital, and labor, $\{y, c, k, \ell\}$, and an allocation of tasks to factors, such that this allocation minimizes the after-tax cost of producing each task and the markets for capital, labor, and the final good clear. The online appendix shows that the equilibrium level of output can be represented as:

$$(2) \qquad y = f(k, \ell; \theta) = \left(\left(\int_0^\theta \psi^k(x)^{\lambda-1} dx \right)^{\frac{1}{\lambda}} \cdot k^{\frac{\lambda-1}{\lambda}} + \left(\int_\theta^1 \psi^\ell(x)^{\lambda-1} dx \right)^{\frac{1}{\lambda}} \cdot \ell^{\frac{\lambda-1}{\lambda}} \right)^{\frac{\lambda}{\lambda-1}},$$

where the threshold task θ satisfies

$$(3) \qquad \theta = \theta^m(k, \ell) \equiv arg\ max_{\theta \in [0,1]} f(k, \ell; \theta).$$

Moreover, factor prices are given by the usual marginal conditions $f_k = R$ and $f_\ell = w$. Consequently, the market-clearing condition for capital is

$$(4) \qquad u'(\bar{y} - k) = 1 + (f_k - \delta) \cdot (1 - \tau^k),$$

while the market-clearing condition for labor is

$$(5) \qquad v'(\ell) = f_\ell \cdot (1 - \varrho) \cdot (1 - \tau^\ell),$$

so that the wedge ϱ and the labor tax τ^ℓ distort the labor market in similar ways.

Finally, the government budget constraint takes the form

$$(6) \qquad g \leq \tau^k \cdot (f_k - \delta) \cdot k + \tau^\ell \cdot f_\ell \cdot \ell.$$

A couple of points about this equilibrium are worth noting. As emphasized in Acemoglu and Restrepo (2018, 2019b), though the output level in the economy can be represented by a constant elasticity of substitution (CES) aggregate of capital and labor, the implications of this setup are very different from models that assume a CES production function with factor-augmenting technologies. First, there is a crucial distinction between capital intensity of production given a fixed allocation of tasks to factors and automation, represented by an increase in θ—which involves the substitution of capital for tasks previously performed by labor. This can be seen from the fact that holding the task allocation constant, the elasticity of substitution between capital and labor is λ, but when θ adjusts, the elasticity is greater. Second, further automation increases productivity but can easily reduce labor demand and the equilibrium wage because of the displacement it creates (mathematically, this works by changing the share parameters of the CES). In contrast, with a standard CES production function labor demand necessarily increases when capital becomes more productive. Third, and for the same reason, automation always reduces the labor share. Finally, our framework also clarifies that marginal increases in automation have second-order effects on aggregate output because, as shown in equation (3), the level of automation is chosen optimally.

I.C. Optimal Policy

We now characterize optimal policy by considering the choices of a benevolent social planner that sets capital and labor taxes τ^k and τ^ℓ and directly controls the extent of automation, represented by θ. We refer to the maximization problem of this planner as the Ramsey problem. As usual, this problem can be transformed so that the planner directly chooses an allocation $\{c, \ell, k, \theta\}$ that maximizes household utility subject to the resource constraint of the economy and a single implementability condition, which combines the government budget constraint in equation (6) and input market equilibrium conditions, equations (4) and (5):

$$(7) \quad \max_{c,\ell,k,\theta} \quad u(\bar{y}-k)+c-v(\ell)$$

$$\text{subject to: } c+g=f(k,\ell;\theta)+(1-\delta)\cdot k \quad \text{(resource constraint)}$$

$$g \le f(k,\ell;\theta)+(1-\delta)\cdot k-u'(\bar{y}-k)\cdot k\frac{v'(\ell)\cdot \ell}{1-\varrho}$$

$$\text{(implementability condition)}$$

Because the planner is assumed to choose the level of automation θ, we do not impose $\theta = \theta^m(k, \ell)$ as an additional constraint. We discuss issues of how the planner's choice of automation can be implemented below. Throughout, we use $\mu > 0$ to denote the multiplier on the implementability condition, which also gives the social value of public funds.

PROPOSITION 1: Optimal capital and labor taxes and automation

The unique solution to the Ramsey problem in equation (7) satisfies

$$\frac{\tau^{k,r}}{1-\tau^{k,r}}=\frac{\mu}{1+\mu}\frac{1}{\varepsilon^{k}(k)} \quad \textit{and} \quad \frac{\tau^{\ell,r}}{1-\tau^{\ell,r}}=\frac{\mu}{1+\mu}\frac{1}{\varepsilon^{\ell}(\ell)}-\frac{\varrho}{1+\mu} \tag{8}$$

and $\theta^{r}=\theta^{m}(k,\ell)$.

The proof of this proposition, like those of all other results in this paper, is provided in the online appendix. The optimal tax formulas in equation (8) follow from the first-order conditions for the maximization problem in equation (7). Uniqueness follows from the fact that the Ramsey problem is convex (the objective function is quasi-concave and the constraint set is convex).

This proposition provides simple and intuitive formulas for the optimal taxes on capital and labor related to the social value of public funds and the inverse of the elasticity of supply of these factors. The formulas show that taxes should be lower for more elastic factors, and in addition, the optimal labor tax is further lowered by the presence of labor market frictions. This latter feature is intuitive: labor market frictions reduce employment beyond the socially optimal level, and the planner corrects for this by reducing labor taxation.

An immediate corollary of this proposition provides one set of sufficient conditions for uniform (symmetric) taxation of capital and labor—$\varepsilon^{k}(k)\simeq\varepsilon^{\ell}(\ell)$ and $\varrho\simeq 0$.

Corollary 1

If $\varepsilon^{k}(k)=\varepsilon^{\ell}(\ell))$ and $\varrho=0$, uniform taxation of capital and labor is optimal.

In section III we will see that realistic values of these parameters are not too far from $\varepsilon^{k}(k)\simeq\varepsilon^{\ell}(\ell)>0$, but labor market imperfections imply $\varrho>0$, so that our framework yields lower labor taxes than capital taxes in the optimum.

Although the formulas in equation (8) apply in a two-period model, the online appendix shows that, under the key assumption that the government must run a balanced budget, these formulas extend to an infinite-horizon setting.[7] The online appendix also derives similar formulas for general preferences over consumption and leisure and clarifies the relationship

---

[7] Even if the government is allowed to incur debt or accumulate assets, the result that the optimal tax system should not simultaneously impose significant taxes on labor and zero (or small) taxes on capital extends to an infinite-horizon setting provided that the long-run elasticity of capital supply, $\varepsilon^{k}(k)$, is not infinite. Straub and Werning (2020) show that in a

between our result and Atkinson and Stiglitz's (1972) principles of optimal commodity taxation.

In line with Diamond and Mirrlees (1971), proposition 1 also shows that, once optimal taxes are imposed on capital and labor, the planner has no reason to deviate from equilibrium automation decisions, $\theta^r = \theta^m(k, \ell)$. This is because any distortions in the labor market are corrected by optimal taxes, and thus, factor prices accurately reflect the social values of capital and labor. Consequently, profit-maximizing automation decisions are optimal as well. We will see that this is no longer true when taxes are not optimal or are subject to additional constraints.

I.D. Excessive Automation with Tax Distortions

Naturally, taxes in practice need not coincide with those characterized in proposition 1 both because of additional constraints and for political economy reasons (policymakers have other objectives and face political or other unmodeled economic constraints). When that is the case, either capital or labor taxes can be (relatively) too low. The interesting case for us, both for conceptual and empirical reasons, is the one where capital taxes are too low and labor taxes are too high, and the necessary and sufficient condition for this follows from equation (8) in proposition 1 and is presented in the next corollary.

Corollary 2

If the tax system $\{\tau^k, \tau^\ell\}$ is below the peak of the Laffer curve and satisfies

$$(9) \qquad \frac{\dfrac{\tau^\ell}{1-\tau^\ell} + \varrho}{\dfrac{1}{\varepsilon^\ell(\ell)} - \dfrac{\tau^\ell}{1-\tau^\ell}} > \frac{\dfrac{\tau^k}{1-\tau^k}}{\dfrac{1}{\varepsilon^k(k)} - \dfrac{\tau^k}{1-\tau^k}},$$

then $\tau^\ell > \tau^{\ell,r}$ and $\tau^{k,r} > \tau^k$—that is, the labor tax is too high and the capital tax too low.

The inequality expressed in equation (9) is sufficient for the tax system being biased against labor and in favor of capital.[8] An important

---

representative household economy where preferences are not time-additive separable and the tax system is not constrained by other considerations, optimal taxes on both capital and labor should converge to zero. We prove in the online appendix that if in addition there are labor market distortions, then optimal long-run taxes are lower on labor than capital.

[8] The government budget constraint implies that both taxes cannot be too high or too low at the same time (provided that we are below the peak of the Laffer curve, meaning that tax revenues cannot be increased by lowering both taxes). Thus, equation (8) is sufficient for $\tau^\ell > \tau^{\ell,r}$ and $\tau^{k,r} > \tau^k$.

implication of such a biased tax structure is that there is too little employment relative to the optimal allocation in proposition 1, and thus marginal increases in employment will have first-order positive effects on welfare. We exploit this insight in the next proposition, where we take the tax system as given and consider a marginal change in automation. To do this in the simplest way, we relax the government budget constraint, equation (6), and value changes in revenue at the social value of public funds given by the multiplier μ.

PROPOSITION 2: When reducing automation improves welfare

Suppose that the tax system $\{\tau^k, \tau^\ell\}$ satisfies equation (9) (and is thus biased against labor and in favor of capital). Welfare (inclusive of fiscal costs and benefits) increases following a small reduction in θ below $\theta^m(k, \ell)$. A small reduction in θ also increases net output provided that $\varepsilon^\ell(\ell) > \varepsilon^k(k)$ and government revenue provided that $\tau^\ell \cdot (1 + \varepsilon^\ell(\ell)) > \tau^k \cdot (1 + \varepsilon^k(k))$.

This result shows that, in contrast with proposition 1, when taxes are not optimal and are biased against labor (in the sense that equation (9) holds), it is welfare improving to restrict automation below its equilibrium level. This result is intuitive in light of the observation in corollary 2 that employment is below the socially optimal level. Specifically, a small reduction in automation will create a first-order welfare gain by shifting demand from capital to labor. Distorting automation is costly, but starting from the equilibrium level of automation, $\theta^m(k, \ell)$, this cost is second-order, since $f_\theta(k, \ell; \theta^m(k, \ell)) = 0$, and hence, a small reduction in automation is welfare improving. This intuition relates proposition 2 to the notion of "so-so (automation) technologies" proposed in Acemoglu and Restrepo (2019a, 2019b): automation is not beneficial to labor when it only increases productivity by a small amount, while still creating the usual displacement of workers as tasks are reallocated from them to capital. The equilibrium condition $f_\theta(k, \ell; \theta^m(k, \ell)) = 0$ implies that automation technologies adopted at marginal tasks are, by definition, so-so. The planner is therefore happy to sacrifice these so-so technologies in order to help labor.[9]

As we will see in section III, the US tax system lies within the range that satisfies equation (9), so that there are prima facie reasons for suspecting that the level of automation may be excessively high in the US economy, as in this proposition.

---

[9] If automation decisions were constrained by available technology, that is, θ had to be less than some $\bar{\theta} < 1$ as in Acemoglu and Restrepo (2018, 2019a), we could have that $f_\theta(k, \ell; \theta^m(k, \ell)) > 0$ if $\theta^m(k, \ell) = \bar{\theta}$. In this case, productivity gains from automating marginal tasks could be positive. If they were sufficiently large, then automation would no longer be a so-so technology and proposition 2 would not apply.

One common intuition is that when confronted with a tax system with distortions, $\{\tau^k, \tau^\ell\}$, the best policy is to redress these tax distortions directly. We next show that this is not always the case. In particular, if for other reasons taxes on labor cannot be reduced below a certain threshold (which we denote by $\bar{\tau}^\ell$), then the tax system satisfies equation (9) and is biased against labor, but this does not necessarily imply that capital taxes should be increased. Rather, constrained optimal policy calls for a reduction in the equilibrium level of automation and may even involve a *lower* tax on capital. Before presenting this result, let us note that in this case we are imposing $\tau^\ell \geq \bar{\tau}^\ell$, which can be expressed as an additional constraint on the Ramsey problem in equation (7) of the form

$$v'(\ell) \leq (1-\bar{\tau}^\ell)\cdot(1-\varrho)\cdot f_\ell, \tag{10}$$

where the lower bound on labor taxes translates into an upper bound on the marginal disutility from work. In the next proposition, we denote the multiplier on this constraint by $\gamma^\ell \cdot \ell \geq 0$ (where the ℓ simply normalizes the multiplier and makes it easier to interpret).

PROPOSITION 3: Excessive automation with tax distortions

Consider the constrained Ramsey problem of maximizing equation (7) subject to the additional constraint $\tau^\ell \geq \bar{\tau}^\ell$, and suppose that in the solution to this problem equation (10) binds. Then the constrained optimal taxes and allocations are given by a labor tax of $\tau^{\ell,c} = \bar{\tau}^\ell$ and a tax subsidy on capital that satisfies

$$\frac{\tau^{k,c}}{1-\tau^{k,c}} = \frac{\mu}{1+\mu}\frac{1}{\varepsilon^k(k)} - \frac{\gamma^\ell}{1+\mu}\cdot(1-\bar{\tau}^\ell)\cdot(1-\varrho)\frac{f_{\ell k}\cdot\ell}{u'-1} \tag{11}$$

and a level of automation $\theta^c < \theta^m(k, \ell)$.

Before discussing the implications of this proposition, we explain the meaning of the constraint expressed in equation (10). The fact that this constraint is binding means that the planner would have chosen a tax rate on labor $\tau^{\ell,r}$ below $\bar{\tau}^\ell$, but the constraint forces the planner to set a higher tax on labor of $\bar{\tau}^\ell$, which results in a tax system biased against labor and in favor of capital, or in other words, the inequality in equation (9) holds. This also implies that the level of employment is below what the planner would have chosen in the unconstrained Ramsey problem.

Given this biased tax system, the planner wants automation to be less than its equilibrium level. The intuition is identical to that in proposition 2: the reduction in automation creates a second-order productivity cost but a first-order gain via its impact on increased employment. Importantly, this holds even when capital taxes can be freely adjusted.

Moreover, the optimal capital tax formula in equation (11) has an additional negative term on the right-hand side relative to equation (8). This negative term can lead not just to lower capital taxes than in the unconstrained Ramsey problem in proposition 1, but even to capital subsidies.[10] The combination of lower capital taxes and limiting the set of tasks that are automated ensures that capital gets used intensively in tasks that are (and should be) automated, while avoiding the displacement of workers from marginal tasks. Both of these margins contribute to raising employment, which increases welfare when the tax system is biased against labor. This is related to the discussion of deepening of automation in Acemoglu and Restrepo (2019a): deepening of automation, which means an increase in the use or productivity of capital in tasks that are already automated, is always beneficial for labor. What is potentially damaging to labor is the extensive margin of automation—because this displaces workers from tasks they were previously performing. Proposition 3 builds on this logic: the planner would like to reduce the range of tasks that are automated by reallocating marginal tasks back to labor and may also want to reduce capital taxes or even subsidize capital at the same time, so that automated tasks can use capital more intensively.

Proposition 3 focused on the case with a lower bound on labor taxes. An equally plausible case is one where, because of political influence of capital owners or because of concerns about capital flight, there is an upper bound on capital taxes.[11] Proposition 2 in the online appendix establishes that in this case, too, the planner would like to reduce automation below its market level, even if taxes on labor can be adjusted. The intuition is similar: the upper bound on capital taxation leads to a tax system biased in favor of capital and against labor, and this makes the displacement of labor by capital in marginal tasks socially costly.

I.E. Implementation

To ease exposition, we have so far assumed that the planner can directly control θ. We now discuss how the desired level of θ can be implemented via taxes. Recall that $k(x)$ is the capital used in task x, and so far we have

---

[10] This might at first appear surprising, especially because the program in proposition 1 is convex, so moving in the direction of the unconstrained optimum should be beneficial. However, the convexity is in the space of allocations and does not imply convexity in the space of taxes. Therefore, increasing the tax rate on capital toward $\tau^{k,r}$ is not necessarily welfare-improving.

[11] A similar constraint on capital taxation is used in the optimal taxation literature (Chamley 1986; Judd 1999; Straub and Werning 2020).

assumed that all types of capital are taxed at the same uniform rate, τ^k. In practice, taxes vary by type of capital (e.g., equipment, software, structures) and industry (because of differential depreciation allowances). In the context of our model, this can be viewed as a task-specific capital tax rate of $\tau^k(x)$. The next proposition establishes when such task-specific capital tax rates are useful and in the process further clarifies the nature of optimal policy interventions.

PROPOSITION 4: Automation tax

Suppose the planner can set task-specific capital taxes and cannot directly control automation decisions. Then, first, under the conditions of proposition 1, the planner sets a uniform capital tax rate, that is, $\tau^k(x) = \tau^k$. Second, under the conditions of proposition 3, the planner prefers to depart from uniform capital taxation. In particular, the planner can implement the level of automation $\theta^c < \theta^m(k, \ell)$ with the following tax scheme:

$$\tau^k(x) = \begin{cases} \tau^k & \text{for } x \leq \theta^c \\ \tau^k + \tau^A & \text{for } x > \theta^c \end{cases},$$

where $\tau^A > 0$ is a task-specific automation tax.

The reason (unconstrained) optimal policy has no use for task-specific taxes is intuitive: in the unconstrained Ramsey problem, there is no need to distort equilibrium automation decisions. However, in the presence of additional constraints, the planner would like to reduce automation to $\theta^c < \theta^m(k, \ell)$, and the planner can achieve this by imposing an incremental tax to capital used in tasks above θ^c. By design, these incremental taxes encourage the use of capital in tasks where capital has a comparative advantage (which helps labor via complementarities across tasks) and discourages the automation of marginal tasks (which also benefits labor by preventing its displacement). In what follows, we refer to the incremental tax on capital τ^A as an automation tax.

II. The US Tax System

In this section, we first introduce the notion of *effective taxes* on capital and labor. Effective taxes summarize the average distortion that the US tax system introduces in the use of capital and labor. We then provide formulas for effective taxes that take into account the various elements of the US tax code and their interaction with the type of financing and ownership structure of the firm making investment decisions.

II.A. *Defining Effective Taxes on Capital*

In our framework, τ^k is the effective tax on (the use of) capital. It is defined as the wedge that the tax system introduces between the internal rate of return for a firm investing in capital and the after-tax rate of return paid to investors. The US tax system includes several taxes, not just a single effective tax on the use of capital. We have personal income taxes on capital income, corporate income taxes, depreciation allowances, and many other instruments that contribute to taxes on different types of capital. Moreover, these taxes vary by form of organization (C corporation versus pass-through) and type of financing (equity versus debt).[12]

We start by providing formulas for effective taxes on the use of capital by type of asset, j, form of organization, and type of financing. To simplify the exposition, we assume the economy is in steady state—the capital-labor ratio remains constant, the tax system is not expected to change, the price of capital goods changes at a constant rate $\pi^j = q_t^j / q_{t-1}^j$, and the capital stock of type j depreciates at a constant rate $\delta^j > 0$.

The internal rate of return of investing one dollar in equipment j at time $t - 1$ is given by

$$r^{f,j} = \mathrm{mpk}^j - \bar{\delta}^j,$$

where mpk^j is the marginal product of investing one dollar in asset j and $\tilde{\delta}^j = 1 - \pi^j \cdot (1 - \delta)$ denotes the total depreciation of the asset (inclusive of investment price changes). Let us denote the after-tax steady-state rate of return to investors by r. The effective tax rate on capital of type j, $\tau^{k,j}$, can then be defined as

$$\frac{1}{1-\tau^{k,j}} = \frac{r^{f,j}}{r} = \frac{\mathrm{mpk}^j - \tilde{\delta}^j}{r}. \tag{12}$$

This formula aligns closely with the effective capital taxes in our conceptual framework in the previous section. In particular, in equation (4), $\frac{1}{1 - \tau^k}$ is equal to the wedge (ratio) between the return to the firm from

---

[12] Pass-through organizations include both S corporations and other pass-throughs, such as sole proprietor businesses and partnerships, and are subject to different tax rules, as we explain below.

using capital—$\text{mpk}^j - \tilde{\delta}^j$ here and given by $f_k - \delta$ in equation (4)—and the return demanded by investors— r here and $u'(\bar{y} - k) - 1$ in equation (4).[13]

The computation of effective tax rates requires measuring the marginal product of capital. We follow Hall and Jorgenson (1967) and back out the marginal product of capital using a representative firm's first-order condition for investment. We need to distinguish between C corporations and pass-through businesses as well as the source of financing, since each of these combinations implies a different first-order condition for investment as well as a different set of taxes on the income generated from capital.

For C corporations that finance their investment with equity, the first-order condition is

$$\text{mpk}^j = \frac{1 - \alpha^j \cdot \tau^c}{1 - \tau^c} \cdot \left(r^e + \tilde{\delta}^j\right), \tag{13}$$

where τ^c is the corporate income tax rate and $\alpha^j \in 0, 1$ are discounts from depreciation allowances, which reduce taxable income and are discussed in the next subsection. In the absence of corporate income taxes, this expression is identical to the standard user cost formula. In addition, r^e is the pretax return to equity holders. This implies that $r = r^e \cdot (1 - \tau^{e,c})$, where $\tau^{e,c}$ is the tax rate on income resulting from ownership of public equity.

Combining the formula for effective taxes in equation (12) with the first-order condition for investment in equation (13), the effective tax rate for an equity-financed C corporation is

$$\frac{1}{1 - \tau^{k,j}_{\text{c-corp,equity}}} = \frac{1}{1 - \tau^{e,c}} \cdot \left(\frac{r^e + \tilde{\delta}^j}{r^e} \frac{1 - \alpha^j \cdot \tau^c}{1 - \tau^c} - \frac{\tilde{\delta}^j}{r^e} \right). \tag{14}$$

The formula shows that the effective tax on capital depends on the taxation of capital income of equity owners, corporate income tax rates, and depreciation allowances. It reiterates that depreciation allowances can significantly offset corporate taxes. For example, with full (immediate) expensing, which corresponds to $\alpha^j = 1$, we would have $\tau^{k,j}_{\text{c-corp,equity}} = \tau^{e,c}$.

The main difference for pass-through businesses is that these organizations do not pay the corporate income tax and are only subject to

---

[13] An alternative is to use a formula for effective taxes based on gross returns: $\frac{1}{1 - \tau^{k,j}_{\text{gross}}} = \frac{\text{mpk}^j}{r + \tilde{\delta}^j}$. All of our results can be expressed in terms of gross returns, but this would require adjusting the empirical estimates of capital supply elasticities, which are in terms of net returns.

personal income taxation. Depreciation allowances in this case lower personal income tax obligations for business owners. The formula for the effective tax on the use of capital for a pass-through business that is financing its investment with (private) equity is

$$\frac{1}{1-\tau^{k,j}_{\text{passthrough,equity}}} = \left(\frac{r^e + \tilde{\delta}^j}{r^e} \frac{1-\alpha^j \cdot \tau^{o,p}}{1-\tau^{o,p}} - \frac{\tilde{\delta}^j}{r^e} \right), \tag{15}$$

where $\tau^{o,p}$ denotes the individual tax rate on the income of owners of pass-through businesses. Note again that with immediate expensing ($\alpha^j = 1$), we have $\tau^{k,j}_{\text{passthrough,equity}} = 0$.

We next turn to debt-financed investments, which allow a further tax discount by subtracting interest payments from taxable income. The presence of these additional tax discounts modifies the first-order condition for investment to

$$\text{mpk}^j = \frac{1-\alpha^j \cdot \tau^c}{1-\tau^c} \cdot \left(r^b \cdot (1-\tau^c) + \tilde{\delta}^j \right), \tag{16}$$

where r^b is the return offered to bondholders and the term $r^b \cdot (1 - \tau^c)$ captures the fact that tax liabilities are lower because of the deduction of interest rate payments. Note that the after-tax return to households that own bonds is $r = r^b \cdot (1 - \tau^{b,c})$, where $\tau^{b,c}$ is the personal income tax rate for capital income from C corporation bonds.

Combining the formula for effective taxes in equation (12) with the first-order condition for investment in equation (16), the effective tax rate for a debt-financed C corporation is

$$\frac{1}{1-\tau^{k,j}_{\text{c-corp,debt}}} = \frac{1}{1-\tau^{b,c}} \cdot \left(\frac{r^b \cdot (1-\tau^c) + \tilde{\delta}^j}{r^b} \cdot \frac{1-\alpha^j \cdot \tau^c}{1-\tau^c} - \frac{\tilde{\delta}^j}{r^b} \right). \tag{17}$$

The effective tax on capital again depends on the personal income tax rate of bondholders, corporate income tax rates, interest rate deductions, and depreciation allowances.[14] The additional tax discounts can easily lead

---

[14] Our model assumes that new and underappreciated old (already installed) capital are perfect substitutes and thus face the same tax rate. When new and old capital are imperfect substitutes, bonus depreciation allowances and other deductions will make new capital cheaper relative to already installed capital (see Auerbach and Kotlikoff 1987). To the extent that capital involved in automation tends to be new capital, this would create an additional incentive for excessive automation.

to a net subsidy to the use of capital. In particular, with full expensing ($\alpha^j = 1$), we have $\tau^{k,j}_{\text{c-corp,debt}} \approx \tau^{b,c} - \tau^c$, which is negative if bondholders face lower individual tax rates than corporations.

Owners of pass-through businesses can also subtract their interest payments on debt from their taxable income. However, if they issue bonds, payments to bondholders are subject to personal income taxation. The formula for the effective tax on the use of capital for a pass-through business that is financing its investment with debt is similar to that of a C corporation and given by

$$\frac{1}{1-\tau^{k,j}_{\text{passthrough,debt}}} = \frac{1}{1-\tau^{b,p}} \cdot \left(\frac{r^b \cdot (1-\tau^{o,p}) + \tilde{\delta}^j}{r^b} \frac{1-\alpha^j \cdot \tau^{o,p}}{1-\tau^{o,p}} - \frac{\tilde{\delta}^j}{r^b} \right), \tag{18}$$

where $\tau^{b,p}$ denotes the individual income tax rate applying to holders of pass-throughs' bonds. As before, with full expensing ($\alpha^j = 1$), we would have $\tau^{k,j}_{\text{passthrough,debt}} \approx \tau^{b,p} - \tau^{o,p}$, which is negative if bondholders face lower income taxes than owners of pass-through businesses.

II.B. Computing Effective Taxes on Capital

We compute effective taxes for equipment, software, and structures separately. For each type of capital good, we calculate effective taxes by form of organization and type of financing, and we aggregate these taxes into a single effective tax rate for the relevant type of capital using investment shares as weights. The online appendix details the sources and numbers used in our calculations. Here we outline the computations of the key objects in our formulas for effective taxes on capital: depreciation allowances, α^j; corporate income taxes and taxes on owners of equity and pass-throughs; and interest rates, economic depreciation, and investment prices.

DEPRECIATION ALLOWANCES The tax discount term, α^j, is equal to the present discounted value of depreciation allowances associated with one unit of capital purchased at time t, which can be computed as

$$\alpha^j = d^j_0 + \sum\nolimits_{s=0}^{\infty} d^j_{s+1} \cdot \prod\nolimits_{\tau=0}^{s} \frac{1-d^j_\tau}{1+r}, \tag{19}$$

where d^j_s denotes the fraction of the investment that a firm is allowed to subtract from its tax liabilities s years after the purchase.

One useful benchmark is when firms can subtract the economic depreciation of their capital goods each period. In the above formula, this means

$d_0^j = 0$ and a constant depreciation rate of δ^j from there on, which adds up to an allowance of $\tilde{\alpha}^j = \delta^j/(\delta^j + r) < 1$.

The Internal Revenue Service (IRS) and the US tax code handle depreciation allowances quite differently from this benchmark, however. The way in which depreciation allowances are determined is specified in IRS Publication 946. The current system places each type of capital under a specific class life—the number of years that a new unit of capital lasts for tax purposes—based on its characteristics and sector. The first reason why tax discounts α^j differ from the one given by constant economic depreciation, $\tilde{\alpha}^j$, is that the depreciation rate implied by a class life is different from the economic depreciation rate.

A second source of an additional tax discount is that the tax code requires taxpayers to follow specific depreciation schedules and enables front-loading of allowances. When computing their tax discount, firms may use a combination of straight line and declining balance methods that yields the highest possible discount. The straight line method allows firms to expense a constant fraction of their initial investment (or undepreciated investment in the initial year in which the method is applied) for each year of remaining tax life. The declining balance method can be used for assets with a class life below twenty years and allows firms to front-load their depreciation allowances by expensing a decreasing fraction of their initial investment each year. Assets in a class life of ten years or less can be depreciated using a 200 percent declining balance rule, which allows firms to expense their undepreciated investment at two times the rate prescribed by the straight line method (2 × 10 percent for an asset in a class life of ten years). Firms can then switch to the straight line method near the end of the asset life to maximize their allowances.[15] Assets with a class life

---

[15] As an example, consider the allowances generated by the purchase of a machine with a class life of ten years. Suppose the purchase takes place in the middle of the year. The straight line method allows a deduction of 5 percent of the cost in the first year, 10 percent for the following nine years, and 5 percent on the eleventh year. The 200 percent declining balance method gives an allowance of 10 percent in the first year (two times the straight line rate of 5 percent), 18 percent in the second year (two times the straight line rate of 10 percent times the undepreciated stock, 90 percent), 14.4 percent in the third year (two times the straight line rate of 10 percent times the undepreciated stock, 72 percent). This continues up to year seven, where the method prescribes an allowance of 5.89 percent, which is below the straight line method allowance of 6.55 percent computed on the undepreciated stock of capital and four and a half years of useful life left. Therefore, the schedule for ten-year property follows the 200 percent declining balance method until year seven and switches to a constant allowance of 6.55 percent of the undepreciated cost for the remaining four and a half years. For further discussion and examples, see the appendix in House and Shapiro (2008).

between ten and twenty years, on the other hand, can be depreciated using a 150 percent declining balance rule, while assets with a class life of more than twenty years adhere to the straight line method.

The third and final source of large discounts from depreciation allowances is recent changes in legislation, passed as part of economic stimulus plans, which introduced bonus depreciation.[16] Under current bonus depreciation provisions, most capital with a class life below twenty years enjoys a 100 percent bonus depreciation, meaning that investors can immediately expense their capital purchases as current costs. This full expensing yields the maximum discount of $\alpha^j = 1$.[17]

We compute α_t^j for 1980–2018 for each type of capital, taking into account changes in the treatment of depreciation allowances and bonus depreciation programs (excluding the further reductions in effective capital taxes generated by the 2017 tax reforms, since these did not affect the automation decisions throughout the 2010s). When computing α_t^j, we assume that firms anticipate no future changes in the tax code, so that they expect current rates to apply in the future.[18]

Figure A.4 in the online appendix plots $\tilde{\alpha}^j$ and α^j for software, equipment, and nonresidential structures. The figure shows that α^j typically exceeds $\tilde{\alpha}^j$ for software and equipment and that recent bonus depreciation provisions generated an increase in allowances, bringing α^j close to 1 for software and equipment in the 2010s.

TAX RATES ON CORPORATIONS AND CAPITAL OWNERS Effective taxes on capital also depend on taxes on corporations and the households who own

---

[16] In particular, the Job Creation and Worker Assistance Act of 2002 (JCWAA) introduced a 30 percent bonus depreciation for 2002–2003; the Jobs and Growth Tax Relief Reconciliation Act of 2003 (JGTRRA) raised the bonus to 50 percent for 2004; the Economic Stimulus Act of 2008 introduced a 50 percent bonus, extended until 2017 by successive bills; the Tax Relief, Unemployment Insurance Reauthorization, and Job Creation Act of 2010 temporarily raised the bonus to 100 percent (full expensing) between September 2010 and the end of 2011. Finally, the Tax Cuts and Jobs Act of 2017 raised bonus depreciation to 100 percent for 2018–2022.

[17] A 100 percent bonus depreciation corresponds to $d_0 = 1$ and $d_s = 0$ for all s > 0 in equation (19). As stated above, capital allowances are generally set by the schedules in IRS Publication 946, which give a specific d_s^j for all $s, j,$ such that $\sum_{s=0}^{T_j} d_s^j = 1$, for each investment type j, and where T_j is the class life for the capital type j. When bonus depreciation is $\gamma < 1$, the taxpayer obtains a first-year bonus allowance equal to γ and then follows the schedules for depreciation allowances for the undepreciated capital stock. Therefore, the bonus allowance series, d_s^j, has $\tilde{d}_s^j = (1-\gamma)d_s^j$, for all $s \geq 1$, and $\tilde{d}_s^j = \gamma + (1-\gamma)d_0^j$ in the initial period.

[18] Anticipated tax reforms create a reevaluation effect for capital that is already installed.

capital. We approximate the average marginal corporate income tax rate τ_t^c for each year as the average tax paid by C corporations:

$$\tau_t^c = \frac{\text{corporate tax revenue}}{\text{net surplus of C corporations}}.$$

The corporate tax revenues are obtained from National Income and Product Accounts (NIPA) tables. The computation of the tax base is presented in the online appendix. We start with operating surplus from corporations and subtract depreciation allowances. We then allocate a fraction of these profits to C corporations using data from the IRS on profits by type of corporation. The remaining share is accounted for by S corporations, which do not pay corporate income taxes, and this share is not included as part of the tax base in the above calculation. The share of corporate profits generated by C corporations has fallen over time from 93 percent in 1980 to 61 percent in 2018, in line with the findings of Smith and others (2019). Our calculations show that once we account for this changing share, the average tax rate on C corporations increased from 25 percent in 1981 to 35 percent in 2000 and then declined to 17.5 percent in 2018.

Note that we are computing corporate income taxes as an average of the taxes paid, rather than by using the statutory rate (46 percent in 1981, 35 percent in the intervening years, and 21 percent in 2018).[19] This is because many corporations pay less corporate income tax than implied by the statutory rate. Throughout, we interpret average taxes as averages of marginal tax rates faced by different types of firms.

Besides taxes paid by corporations, taxes paid by households on their capital income from equity and lending also contribute to the effective tax on the use of capital—the terms $\tau^{e,c}$, $\tau^{b,c}$, and $\tau^{b,p}$ in equations (14), (17), and (18). We compute $\tau^{e,c}$ as the average tax rate paid by owners of equity on their dividends and capital gains. We start by computing the share of corporate equity that is directly held by US households and is thus subject to taxation. Using data from the Board of Governors of the Federal Reserve System, we approximate this as the share of corporate equity owned by US households and nonprofit organizations serving these households, which has fallen from 58 percent in 1981 to 37 percent in 2018. We follow the

---

[19] See the IRS SOI Tax Stats—Historical Table 24, https://www.irs.gov/statistics/soi-tax-stats-historical-table-24.

CBO (2014) and assume that the remaining share is owned by funds or kept in accounts that are not subject to additional taxation.

Taxes paid by households depend on how corporate profits are realized. Qualified dividends or capital gains are taxed at a maximum capital gains tax rate specified by the IRS.[20] These include dividends on stocks held for more than sixty-one days or capital gains on stocks owned for over a year. Ordinary (nonqualified) dividends or capital gains apply to stock owned over shorter periods and are taxed at the same rate as individual income. The remaining profits are for stocks held until death, whose capital gains are never realized and thus face no taxation. We compute the share of profits realized through ordinary dividends and short-term capital gains by using data from the IRS Individual Complete Report (Publication 1304, table A) for the period 1990–2017 and the IRS Statistics of Income (SOI) Tax Stats (Sales of Capital Assets Reported on Individual Tax Returns) for the period 1990–2012. Publication 1304 reports households' ordinary dividend income from corporate stocks, while the SOI Tax Stats reports the short-term capital gains on corporate stocks. Short-term dividends and ordinary capital gains account for the bulk of realized profits from C corporations (about 60 percent in recent years). The remaining share of profits is accounted for by long-term qualified gains and dividends, or by stocks held until death whose capital gains are never realized. We assume that each of these two forms makes up an equal share of profits, which aligns with what the CBO reports for 2011.

The average tax rate on profits derived from C corporation profits (after paying corporate taxes) is thus given by

$$\tau_t^{e,c} = \begin{matrix}\text{share directly}\\ \text{owned}_t\end{matrix} \cdot \left(\begin{matrix}\text{share short-}\\ \text{term ordinary}_t\end{matrix} \cdot \tau_t^{o} + \begin{matrix}\text{share long-}\\ \text{term qualified}_t\end{matrix} \cdot \tau_t^{q} + \begin{matrix}\text{share held}\\ \text{until death}_t\end{matrix} \cdot 0\% \right).$$

Here, τ_t^o is the average tax rate on short-term ordinary capital gains and dividends, and τ_t^q is the average tax rate on long-term qualified capital

---

[20] The maximum capital gains tax rate is specified in IRS Publication 550. In 2018, taxpayers facing a marginal tax rate below 15 percent had a maximum capital gains rate of 0 percent. Taxpayers facing a marginal tax rate between 22 percent and 35 percent had a maximum capital gains tax rate of 15 percent. Finally, taxpayers facing a marginal tax rate of 35 percent faced a maximum capital gains tax rate of 20 percent.

gains and dividends. Both average taxes are computed using data from the Office for Tax Analysis for 1980–2014. In recent years, the average tax rate on ordinary short-term gains and dividends was $\tau_t^o = 24$ percent and the average tax rate on long-term qualified capital gains and dividends was $\tau_t^q = 18$ percent. Our estimates show that $\tau_t^{e,c}$ has hovered around a historical average of 15 percent and experienced a temporary reduction to 12.5 percent during the 2000s.[21]

Turning to taxation of rental income for bondholders, the CBO estimates that 52.3 percent of C corporation bonds are held directly by households, 14.9 percent generate income that is temporarily deferred for tax purposes, and the rest is held by funds or kept in accounts that are not subject to additional taxation. For pass-through entities, the share owned by households is larger, 76.3 percent, and the share deferred is 10.1 percent. Moreover, the CBO reports that the rental income owned by households is subject to personal income taxes at the average rate 27.4 percent in 2014. Supposing that temporarily deferred income is subsequently taxed at the same rate as the rest of rental income, we estimate the average tax paid by bondholders on their rental income from C corporations and pass-throughs, respectively, as $\tau^{b,c} = 16.84$ percent and $\tau^{b,p} = 23.25$ percent, and we assume that these rates have remained constant over time.

The final item required for our calculations is the tax rate paid by owners of pass-throughs, which we separate into S corporations and other pass-throughs (sole proprietor businesses and partnerships). Profits from S corporations are taxed at the individual income rate of the owners. We assume that the average tax rate paid by owners of S corporations is the same as the average tax paid by individuals earning ordinary short-term dividends and capital gains, τ_t^o.[22] In economic terms, this requires owners of S corporations to have a similar income profile as investors in public equity. In addition, part of the profits generated by S corporations accrue

---

[21] Our estimate for $\tau^{e,c}$ assumes that new investments are financed with new equity. When new investments are partly financed with retained earnings, the effective tax on capital might be even lower. In addition, accrual-equivalent taxes on capital will be generally lower than the average tax rate we use for capital gains, because gains are typically postponed relative to accrual. We thank Alan Auerbach for raising these two issues.

[22] Profits from S corporations are also taxed as corporate income by some states. To account for these taxes, we add the average state and local tax rate on businesses, which we compute by dividing state and local revenues from business taxes by the net operating surplus of corporations. State and local taxes on businesses are small in practice, with an average value near 3 percent in recent years.

only when the company is sold, and these profits are taxed at the maximum qualified rate, τ_t^q. Thus, we measure the average tax paid by owners of S corporations on their profits as

$$\tau^{o,s} = \tau_t^o - \text{share capital gains} \cdot (\tau_t^o - \tau_t^q).$$

Using data on sales of pass-through businesses reported by the IRS for 1990–2000, we estimate the average share of capital gains in S corporation profits as 25 percent and assume it has remained at this level over time. Our estimates imply that $\tau_t^{o,s}$ has been roughly constant as well, at about 27 percent, reaching 28 percent in 2018. Since self-proprietors' and partnerships' income is reported as personal income, we have no data on the tax rate faced by owners on profits, and so we assume that they face the average tax rate on income (obtained from the IRS, SOI Tax Stats), which has been approximately 14.6 percent in recent years.

Overall, our estimates imply that in 2011 the average corporate income tax was 26.4 percent (with equity holders paying an additional 11.8 percent on top of this), the average tax rate paid by S corporation owners was 23 percent, and the average tax rate paid by owners of other pass-throughs was 14.6 percent. These numbers align closely with those from the CBO and with Cooper and others (2016).[23]

INTEREST RATES, DEPRECIATION, AND INVESTMENT PRICES We assume a constant interest rate, a constant growth rate for investment prices, and a constant rate of economic depreciation for each asset that match historical averages from 1981 to 2017. We use a constant value of $r^b = 4.21$ percent per annum for bondholders, given by the average of the Moody's Seasoned AAA Corporate Bond Yield minus realized inflation between 1981 and 2017. Likewise, we use a constant value of $r^e = 4.36$ percent per annum for equity holders, which is the historical average of the real rate of return on the S&P 500 over 1957–2018. The constant growth rate for investment prices is estimated from the average change of investment price indexes by type of capital from the Bureau of Economic Analysis (BEA) fixed asset tables between 1981 and 2017. These imply an annual average growth rate of prices equal to −1.6 percent for software, −1 percent for equipment, and 2 percent for nonresidential structures. The economic depreciation rates, the δ_t^j's, are taken directly from the BEA fixed asset tables as the averages

---

[23] Using IRS data, Cooper and others (2016) estimate that in 2011 C corporations paid an average tax rate of 23 percent (plus 8.25 percent on the household side), S corporations paid an average tax rate of 25 percent, and other pass-throughs paid an average tax rate of about 14.7 percent.

for 1981–2017 (the average economic depreciation rate per annum is 23.4 percent for software, 13.9 percent for equipment, and 2.6 percent for nonresidential structures).

II.C. Effective Taxes on Labor

In our model, τ^ℓ is the effective tax on (the use of) labor. However, as with capital, there is no single tax on labor in the US tax code. Instead, labor income is subject to a number of different taxes both at the federal and local levels. Means-tested public programs may generate additional implicit taxes on labor. The effective tax on labor is given by the wedge that the tax system introduces between the marginal product of labor and the before-tax wage, mpl^f. The representative firm will demand labor until the marginal product of labor, mpl^f, equals the cost of one unit of labor given by total compensation. That is,

$$\text{mpl}^f = \text{compensation} = \text{salary} + \text{benefits}.$$

Wage income is subject to personal income tax at a rate τ^h and payroll taxes at a rate τ^p. Benefits are not taxed but might be imperfectly valued by workers, which we capture by converting them to an income-equivalent amount by multiplying them with $\varphi \in 0, 1$. Consequently, the after-tax return to work for the household is given by

$$w = \text{salary} \cdot (1 - \tau^h - \tau^p) + \text{benefits} \cdot \varphi.$$

The effective tax rate on labor is defined, analogously to the effective tax on capital, as

$$\frac{1}{1-\tau^\ell} = \frac{\text{mpl}^f}{w} \Rightarrow \tau^\ell = \frac{\text{salary} \cdot (\tau^h + \tau^p) + \text{benefits} \cdot (1-\varphi)}{\text{compensation}}.$$

We measure the terms in this expression as follows. From national accounts we obtain data on salaries and total compensation for the corporate sector. We treat employers' contributions to pensions and health insurance as part of the benefits since these are not taxed. We assume that workers outside the corporate sector receive a similar split between benefits and salaries and are therefore subject to the same effective taxes. We use a payroll tax rate of 15.3 percent, which is the statutory rate that applies to all earners with an income below \$132,900 in 2018 (a level that roughly matches the 95th percentile of income). Since the vast majority of jobs at risk of automation are performed by workers in the middle of the income distribution, the payroll tax of 15.3 percent is relevant for automation decisions and is incorporated into our effective tax rate on labor. We measure the personal income tax rate τ^h, consistent with our treatment of payroll taxes, as the average income tax paid by earners below the 95th percentile. This is computed from publicly available data from the IRS for 1986–2017. The estimate for τ^h has been stable in recent years at a level close to 10 percent.[24] Finally, we use a value of $\varphi = 0.65$ building on estimates from Gruber and Krueger (1991), Goldman, Sood, and Leibowitz (2005), and Lennon (2020), which suggest that one dollar of spending on benefits is valued on average at 65 cents by households. This increases our estimates for τ^ℓ by 3 percent.

Besides our baseline estimate for τ^ℓ described above, in the online appendix we present results using an estimate for the effective tax on labor which incorporates the implications of means-tested welfare programs. In particular, there is a range of programs, including cash transfers and credits, that are phased out as individual income increases and various social programs (such as disability insurance and unemployment insurance) in which individuals participate less when labor demand is high (see, for instance, Autor and Duggan 2003; Autor, Dorn, and Hanson 2013; Acemoglu and Restrepo 2020). As a consequence, transfers decline as labor demand rises, and this acts as an additional implicit tax on labor, τ^d. Austin, Glaeser, and Summers (2018) estimate that the public expenditures resulting from a person going into nonemployment was \$4,900 per year between 2010 and 2016 (\$6,300 for those in long-term nonemployment and \$2,300 for the short-term unemployed). This is roughly 8 percent of the average yearly worker compensation during this period, suggesting that social expenditure and disability insurance add an extra 8 percent tax to labor.

II.D. Effective Tax Rates in the United States

Figure 5 in the online appendix depicts the evolution of the average personal income tax and average capital tax rates for C corporations (including both corporate income taxes and personal income taxation) and for S corporations (whose owners only pay personal income taxes and some state-level taxes). Taxes on C corporations' profits decline significantly from 2000 onward, reflecting declines in the statutory corporate income tax rate over time. Taxes on pass-through profits have remained stable around 25 percent, and the average individual income tax has remained close to 15 percent.

---

[24] If we were to use the average payroll tax (about 10 percent in recent years) and the average income tax (about 14.6 percent in recent years), we would end up with a very similar effective tax rate on labor.

Figure 1. Effective Tax Rates on Labor, Software, Equipment, and Nonresidential Structures

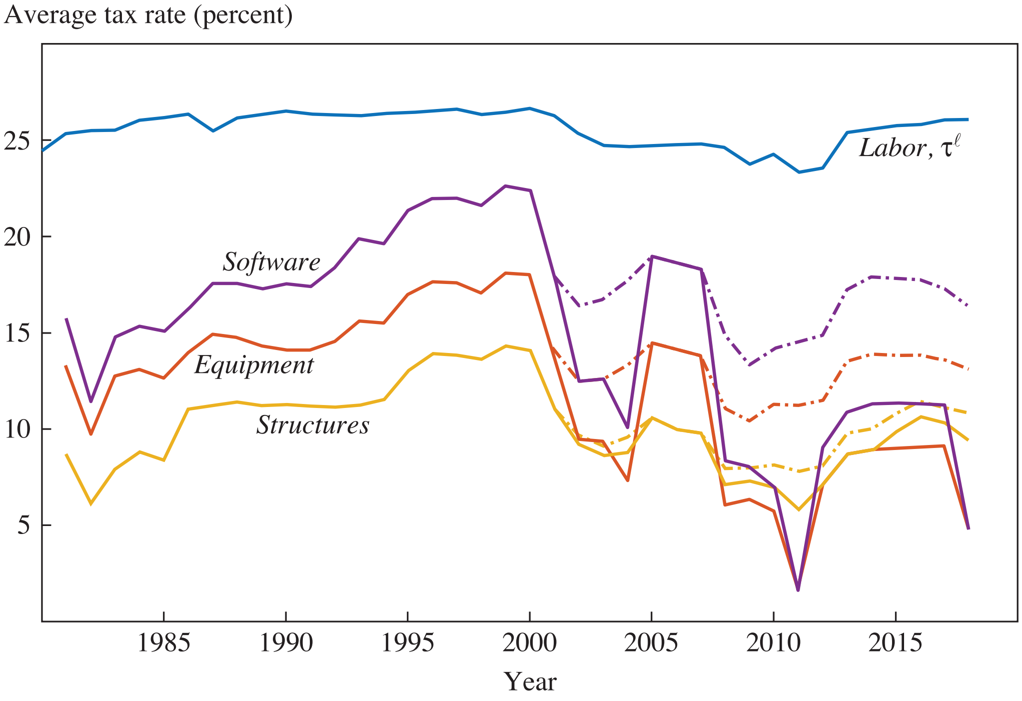

Source: Authors' calculations.

Notes: The solid lines depict the observed effective taxes. The dashed lines present the effective taxes that would result if the treatment of allowances had remained as in the year 2000.

Figure 1 presents our estimates for the effective tax rates on labor and different types of capital (in turn computed from effective tax rates on capital and depreciation allowances for C corporations, S corporations, and other pass-through businesses and the differential taxation of capital financed with debt and equity). The solid lines show the effective taxes on software, equipment, nonresidential structures, and labor.

Several points about these effective tax rates are worth noting. First, effective taxes on equipment and software are low compared to the effective taxes on labor. Our benchmark effective tax on labor (which does not include the implicit taxes implied by means-tested programs) hovers around 25.5 percent.[25] In contrast, effective taxes on both equipment capital and software in the 2010s (and before the tax reform of 2017) are around

[25] Our estimates imply that the net tax revenue collected by the government with these instruments is roughly 18.6 percent of GDP (25.5 percent × labor income in GDP +10 percent × net capital income in GDP). This figure matches closely the average share of personal income taxes, corporate taxes, and Social Security contribution in GDP for the period considered in our study (18.7 percent for 1981–2018 in NIPA table 3.1).

10 percent.[26] Second, effective taxes on equipment and software were higher in the 1990s and early 2000s and declined significantly thereafter. This decline is mostly because of the reforms summarized in footnote 16, which have increased depreciation allowances. The dashed lines in figure 1 illustrate the contribution of these reforms by plotting the (counterfactual) effective taxes on different types of capital that would have applied had the treatment of depreciation allowances remained as it was in 2000. They show that about half of the decline in the effective taxes on software and equipment capital is due to the more generous depreciation allowances introduced since 2002. Third, effective taxes on equipment and software decreased further, to about 5 percent, following the 2017 tax reform, which introduced full expensing of these capital expenditures. Finally, because depreciation allowances for structures are lower, the effective tax on nonresidential structures is higher today than tax rates on equipment and software, but in the past the ordering was reversed.

For our purposes, effective tax rates on equipment and software are more relevant, since these are the types of capital that are involved in automation. In what follows, we will summarize the US tax system as an effective tax on labor of $\tau^l = 25.5$ percent and an effective tax on capital of $\tau^k = 10$ percent (the level before the 2017 tax reforms). We will also separately discuss the implications of the reforms in the 2000s and the 2017 tax reform.

III. Does the US Tax Code Favor Automation?

In this section, we investigate whether the US tax system is biased against labor and favors excessive automation. We then explore the implications of different tax reforms.

III.A. Parameter Choices

We first review the estimates of the main parameters in our model. The parameter λ corresponds to the short-run elasticity of substitution between capital and labor. This is the elasticity of substitution between capital and labor holding the amount of automation (and, more generally, the state of technology) constant and without any compositional changes (for example,

---

[26] These effective tax rates are lower than those reported in CBO (2014). Two factors explain the differences. First, and most importantly, the CBO does not incorporate bonus depreciation allowances (based on the argument that these may not be extended in the future). Second, the CBO uses the statutory rate of corporate income tax. As noted above, we do not believe this gives an accurate estimate of the effective tax on capital, since most corporations pay less than the statutory rate.

between firms with different technologies or between industries). Under the assumption that in the short run the allocation of tasks to factors is fixed, this elasticity can be approximated by the short-run elasticity of substitution within establishments, which is estimated to be $\lambda = 0.5$ in Oberfield and Raval (2014).

The other important building block of the production side of our economy is given by the comparative advantage schedules for labor and capital, $\psi^{\ell}(x)$ and $\psi^{k}(x)$. We reduce the dimensions of these functions by assuming that they take isoelastic forms:

$$\frac{\psi^{\ell}(x)}{\psi^{k}(x)} = A \cdot x^{\zeta} \qquad \psi^{\ell}(x) = A \cdot x^{\zeta\upsilon},$$

where $\zeta \geq 0$ controls how the comparative advantage of labor changes across tasks and υ controls the relationship between the comparative and absolute advantage of labor. We take $\upsilon = 1$ as our baseline, which implies that labor is more productive at higher-index tasks (where it has a comparative advantage), while capital has a constant productivity across tasks, as in the "balanced growth" specification in Acemoglu and Restrepo (2018). The online appendix explores the opposite case in which $\upsilon = 0$ and labor is less productive in tasks where it has comparative advantage.

The parameter of comparative advantage ζ (together with λ) shapes the long-run substitution possibilities between capital and labor. In the long run, changes in factor prices will lead to endogenous development and adoption of automation technologies, and as the allocation of tasks to factors changes, there will be greater substitution between capital and labor than implied by λ. The extent of this greater substitution is shaped by the comparative advantage of labor across tasks. In particular, since $\lambda = 0.5$, a lower user cost of capital will increase the labor share of national income in the short run (because capital and labor are gross complements given θ), but as automation adjusts, the labor share could end up lower than it was before the change. Karabarbounis and Neiman (2014) estimate that a 10 percent reduction in the user cost of capital lowers the labor share by 0.83 to 1.67 percentage points in the long run. The midpoint of this range implies $\zeta = 2.12$ in the context of our model.[27]

Turning to labor market imperfections, recall that the wedge ϱ captures the difference between the wage earned by workers and workers'

---

[27] More specifically, Karabarbounis and Neiman (2014) use a constant elasticity of substitution aggregate production function without automation or reallocation of tasks and show that their estimates correspond to a long-run elasticity of substitution in the 1.2–1.5 range.

opportunity cost. This motivates measuring ϱ as the (average) permanent earning loss from job separation. The majority of the estimates of these earning losses in the labor literature are within the range of 5–25 percent with a midpoint of 15 percent.[28] Motivated by this evidence, we choose a baseline value of $\varrho = 0.15$.[29]

The remaining key parameters of our framework are the Hicksian elasticities of labor and capital supply (Hicksian elasticities are the relevant ones in our context because we are focusing on permanent tax reforms). We adopt the following functional forms for utility: $u(\bar{y} - k) = -B \cdot k^{1+1/\varepsilon^k}/(1 + 1/\varepsilon^k) - k$ and $v(\ell) = \ell^{1+1/\varepsilon^\ell}/(1 + 1/\varepsilon^\ell)$, so that the two Hicksian elasticities, $\varepsilon^k \geq 0$ and $\varepsilon^\ell \geq 0$, are constant. The parameters A and B are calibrated to match an aggregate labor share of 56 percent and a net capital share of 26 percent, with the depreciation rate fixed at 5.5 percent per year.

Because our model does not distinguish between the intensive (hours conditional on employment) and extensive (employment) margin, we use the combined elasticity for total hours of work. Chetty and others (2011) report micro elasticity estimates, obtained from differences in tax rates and wages across regions and demographic groups within a country, in the range of 0.46–0.76 (of which 0.33 comes from the intensive margin and 0.13–0.45 comes from the extensive margin). These numbers are close to macro elasticity estimates obtained from tax differences across countries, which are also around 0.7. Because there might be nonlinearities in supply elasticities (see, for example, Mui and Schoefer 2019), and because there is uncertainty about the exact supply elasticities, we explore the implications of labor supply elasticities between 0.46 and 1 in our robustness checks.[30]

---

[28] Couch and Placzek (2010) survey this literature and present their own estimates, suggesting long-run earning declines from separations of 5 percent. Jacobson, LaLonde, and Sullivan (1993) find long-run earning declines of about 25 percent. Davis and von Wachter (2011) report long-run earning losses of 10 percent in normal times and 20 percent in recessions.

[29] Some of the earning losses may be due to loss of firm-specific human capital. If productivity gains from firm-specific human capital are shared equally between firms and workers, these would also create a wedge identical in reduced form to our ϱ. We also note that there are other factors that would act like a wedge, generating additional incentives to raise employment. These include negative spillovers from nonemployment on family, friends, and communities and on political behavior (see Austin, Glaeser, and Summers 2018). Because quantifying these effects is more difficult, we are ignoring them in the current paper.

[30] In the presence of some types of labor market frictions, the extensive margin changes in employment may take place off the labor supply curve, while intensive margin changes are on the labor supply curve. In table A.3 in the online appendix we show that our main conclusions are robust if we reduce ϱ to 0.075, so that labor market frictions apply only to the extensive margin changes in employment (which make up about half of the variation given the elasticities reported in the text).

The parameter ε^k corresponds to the long-run elasticity with which the supply of capital responds to changes in net returns $d\ln k/d\ln r$ or the keep rate from net capital taxes $d\ln k/d\ln(1 - \tau^k)$ (and is thus different from the demand-side elasticities that are informative about how much investment or capital at the firm level will respond to the user cost of capital). Although there is much uncertainty about this elasticity and many theoretical analyses assume it to be infinite (for example, by imposing time-additive, discounted utility), a number of recent papers estimate it to be much smaller. These studies exploit reforms that change taxes on wealth for different groups of households and find medium-run elasticities that range from 0.2 to 0.65 over four-to-eight-year periods (see Zoutman 2018; Durán-Cabré, Esteller-Moré, and Mas-Montserrat 2019; Jakobsen and others 2020).[31] Using a calibrated life-cycle model and assuming a net after-tax return of $r = 5$ percent, Jakobsen and others (2020) show that their medium-run estimates are consistent with long-run elasticities ranging from 0.58 for the wealthy to 1.15 for the very wealthy. With a lower tax net return of 4 percent (in line with the numbers used in our computation of net effective taxes), long-run capital supply elasticities would be even lower, and conversely, with an after-tax rate of return of 7 percent, these elasticities would range between 1 for wealthy households and 1.9 for the very wealthy (see table 3 in Jakobsen and others 2020). We set our baseline capital supply elasticity to 0.65, which lies at the upper end of the medium-run elasticities reported above and is the average elasticity for the wealthy in Jakobsen and others' (2020) preferred scenario with $r = 5$ percent.[32] We explore the robustness of our results to using a higher elasticity of capital supply in the online appendix.

---

[31] These estimates are from small and fairly open economies, such as Denmark, the Netherlands, and Catalonia, and thus presumably include the response due to the international mobility of capital.

[32] We view our choice as conservative given other estimates in the literature. Brülhart and others (2016) estimate the elasticity of capital to after-tax returns using variation across Swiss cantons. They find an elasticity of 1.05 but also show that about a quarter of the effects are driven by migration across cantons and do not involve a change in savings—which is the relevant margin for optimal taxation. In their concluding remarks, they argue that once this response is accounted for, their numbers are comparable to the medium-run estimates of Jakobsen and others (2020). Kleven and Schultz (2014) estimate an elasticity of capital supply with respect to one minus the tax rate on capital income of 0.3, which would imply an even more inelastic response of capital, reinforcing our results. Finally, a related literature finds small elasticities of savings to one minus the estate tax rate, typically about 0.09–0.16 (see Joulfaian 2006; Kopczuk and Slemrod 2000), which also imply less elastic responses of the supply of capital.

III.B. Is the US Tax System Biased against Labor?

We first verify that the US tax system (with τ^ℓ = 25.5 percent and τ^k = 10 percent) is biased against labor. The estimated US taxes comfortably satisfy equation (8) when we use the elasticity estimates in the previous subsection, $\varepsilon^\ell = 0.7$ and $\varepsilon^k = 0.65$.

Equation (8) implies that current US taxes on labor are too high and US taxes on capital are too low relative to the optimum. In fact, the formulas in proposition 1 for our baseline choice of parameters imply that optimal taxes should be $\tau^{k,r}$ = 26.65 percent and $\tau^{\ell,r}$ = 18.22 percent, which contrast with the observed taxes of τ^k = 10 percent and τ^ℓ = 25.5 percent. The optimal tax on labor is lower than on capital because the supply elasticities for the two factors are similar, while there is an additional wedge for labor ($\varrho = 0.15$), which the optimal tax system corrects for.

The conclusion that the US tax system is biased against labor is robust to variations in our measurement of effective taxes and our estimates of the elasticities of the supply of capital and labor. The top panel of figure 2 documents that variations in how we compute effective taxes on capital and labor do not change this conclusion. It depicts two contour plots for τ^ℓ and τ^k that satisfy equation (8) for the baseline values of the remaining parameters ($\varepsilon^\ell = 0.7$; $\varepsilon^k = 0.65$) and for $\varrho = 0.15$ and $\varrho = 0$. All of our tax estimates lie within these sets and thus comfortably satisfy equation (9) regardless of the value of ϱ.

The bottom panel of figure 2 documents that the US tax system remains biased against labor when we vary the elasticities for the supply of capital and labor. The figure presents contour plots for combinations of elasticities ε^ℓ and ε^k that satisfy equation (9) for our baseline estimates of the US tax system (τ^ℓ =25.5 percent;τ^k = 10 percent) and again separately for $\varrho = 0.15$ and $\varrho = 0$. We find that even if the capital supply had a unitary elasticity, the US tax system would satisfy equation (9) and would continue to be biased against labor.

III.C. Implications of the US Tax System for Automation and Employment

As discussed in our theory section, the bias against labor in the US tax system will generate excessive automation and lead to lower employment than is socially optimal. We now return to our baseline parameters and investigate the implications of the pro-capital bias of the US tax system for automation, employment, the labor share, and welfare.

As a first step, we compare the implied equilibrium level of automation under the tax system in the 2010s (before the 2017 tax reform),

Figure 2. Contour Plots of Taxes and Elasticities That Verify Equation (9)

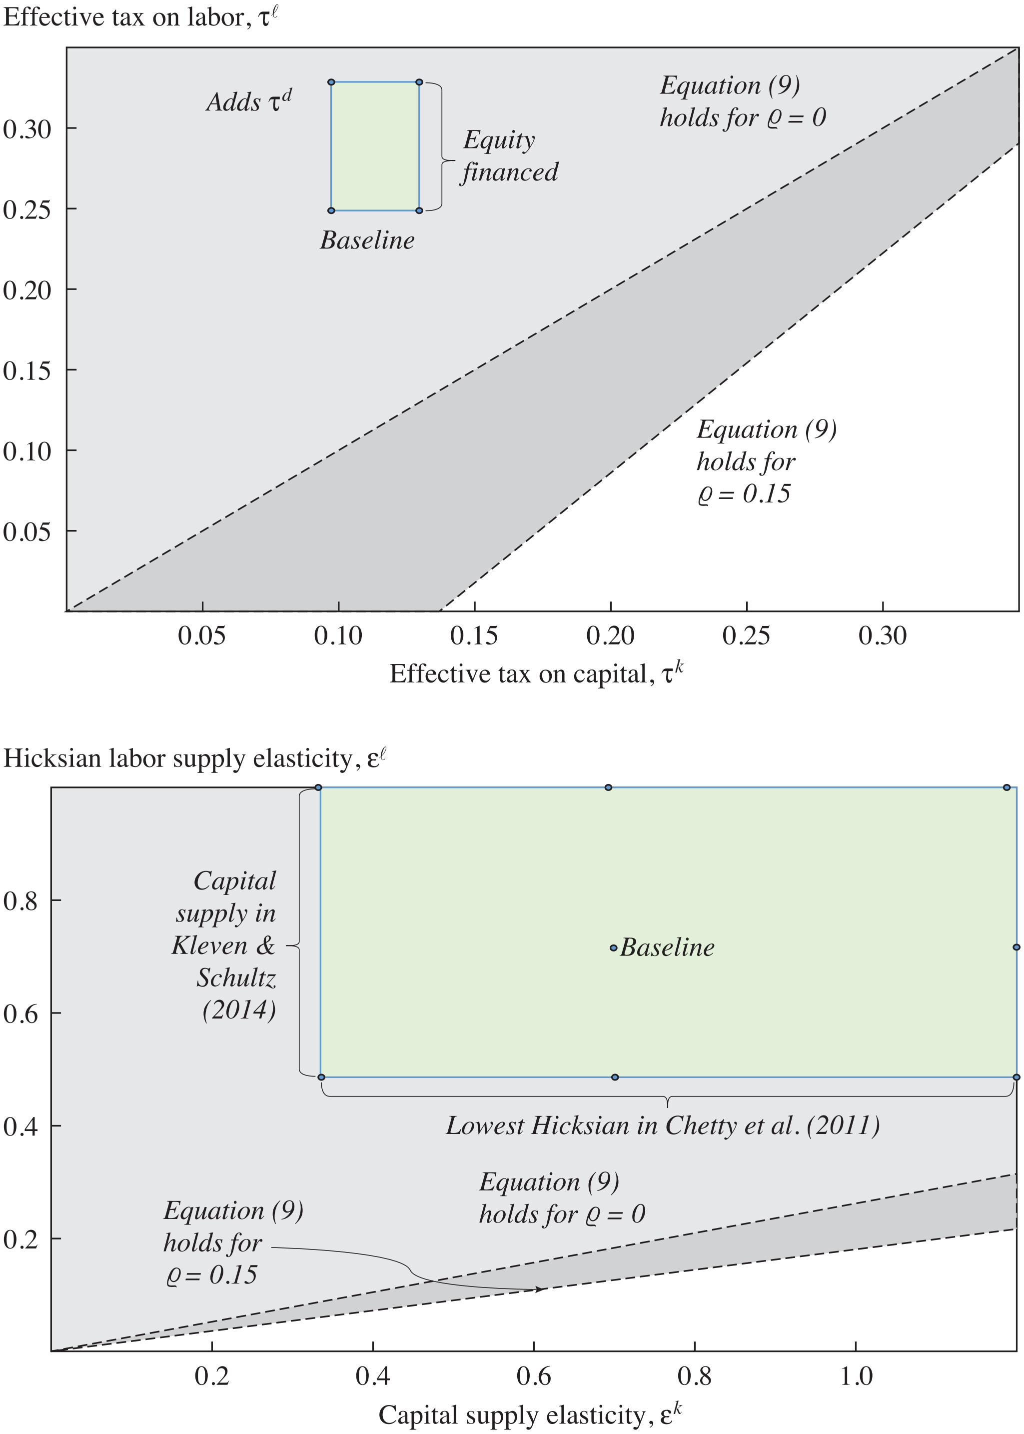

Source: Authors' calculations.

Notes: The top panel shows contour plots for estimates of the current US tax system and the bottom panel depicts contour plots for labor and capital supply elasticities to verify the robustness of the claim that the US tax system is biased against labor. Shaded boxes represent the range of estimates we consider in our robustness checks, and in each case we separately mark our baseline estimates. Equation (9) is satisfied for $\varrho = 0$ in the light gray area and for $\varrho = 0.15$ in both the light and the dark gray areas.

Table 1. Equilibrium under the Current Tax System and under Other Potential Scenarios

	Current system (1)	*Ramsey solution (2)*	*Distorting* θ *(3)*	*Distorting* θ *and changing* τ^k *(4)*	*Distorting* θ *and changing* τ^ℓ *(5)*
Tax system					
τ^k (%)	10.00	26.65	10.00	8.39	10.00
τ^ℓ (%)	25.50	18.22	25.50	25.50	24.89
θ	0.276	0.265	0.267	0.265	0.264
τ^A (%)	0.00	0.00	10.15	12.90	13.07
Aggregates (%)					
Employment	—	4.02	1.14	1.59	1.96
Labor share	56.00	56.78	57.93	58.44	58.54
Net output	—	0.44	–0.10	0.16	0.20
C. E. welfare change	—	0.38	0.09	0.14	0.18
Revenue	—	0.00	1.41	0.00	0.00

Source: Authors' calculations.

Notes: This table shows the effective capital and labor taxes, the level of automation, and the automation tax under different scenarios. It also presents the implied changes in employment, output, welfare, and government revenue, and the level of labor share in national income. The first column is for the current US tax system. The second column shows the unconstrained Ramsey solution. Column 3 considers the implications of changing the level of automation, θ via automation taxes (and no other change in policy). Column 4 additionally allows a change in the effective tax on capital, and column 5 considers a change in the effective tax on labor. Change in welfare is in terms of consumption equivalent. See the text for details.

$\tau^\ell = 25.5$ percent and $\tau^k = 10$ percent, to the equilibrium with optimal taxes, $\tau^{\ell,r} = 18.22$ percent and $\tau^{k,r} = 26.65$ percent. Columns 1 and 2 of table 1 present this comparison. Because the optimal tax system encourages the use of labor in production (relative to the US system in the 2010s), it leads to a lower level of automation than currently. Under the optimal tax system, θ declines by 4.1 percent from its equilibrium value in the 2010s.[33] This lower level of automation would also increase the labor share by 0.78 percentage point and, together with the lower labor tax, increase employment by 4.02 percent. Finally, welfare would be higher by 0.38 percent in consumption-equivalent terms (meaning that the welfare gains are equivalent to increasing consumption in period 1 by 0.38 percent). Although this increase in welfare appears small (relative to the change in employment), this is due to the usual intuition related to Harberger's

---

[33] Though the magnitude of a change in θ is not directly interpretable, we can compute the share of employment that would be displaced with the higher level of θ. Given our parameterization of λ, $\psi^\ell(x)$, and $\psi^k(x)$, reducing θ from 0.276 to 0.265 results in 3.3 percent fewer workers displaced by automation.

triangles: because changes in welfare are second order near the optimum, they tend to be smaller than changes in quantities unless we are very far away from this optimum.

In table 1, we used an effective tax rate on labor of $\tau^{\ell} = 25.5$ percent, which does not include the additional implicit tax on labor implied by means-tested programs. Table A.4 in the online appendix shows that when we incorporate this additional implicit tax on labor supply and set $\tau^{\ell} = 33.5$ percent, the employment and welfare gains from changing the current system are amplified. Moving to optimal taxes now increases employment by 6.07 percent, the labor share by 1.09 percentage points, and welfare by 0.81 percent.

The conclusion that we can achieve higher welfare through tax reforms that raise employment and reduce automation is robust in respect to variations in parameters and the measurement of taxes. Figure 3 considers the same range of taxes and parameters as in the two panels of figure 2. The contours in this figure correspond to combinations of current tax rates (top panel) and elasticities (bottom panel) that give the same employment response when we switch from the current tax system to optimal taxes. For a wide range of parameters, optimal taxes induce levels of employment that are 2–10 percent larger than in the current system.

Recall from proposition 2 that when the tax system is biased against labor, the level of automation is not only greater than in the Ramsey solution, but it is also excessively high compared to what would be socially optimal given the tax system. Column 3 in table 1 quantifies this inefficiency by computing the level of automation that would maximize welfare taking the current capital and labor taxes as given.[34] The level of automation that maximizes welfare is $\theta = 0.267$, which is 3.3 percent lower than equilibrium automation. In line with proposition 4, this lower level of automation can be implemented with an automation tax of 10.15 percent, so that a task will be automated only if replacing labor with capital reduces the cost of producing that task by more than 10.15 percent. The automation tax raises employment by 1.14 percent—partially correcting for some of the inefficiencies in the current system and raising welfare—and the

---

[34] The alternative is to follow proposition 2 and maximize the sum of the representative household's utility plus the change in revenue valued at μ (which is the social value of government funds). Here, we simply maximize welfare—given by the representative household's utility—to make the results in this column comparable to the rest of the table. Valuing additional revenues with the multiplier μ leads to higher automation taxes, since reductions in θ have the additional benefit of generating higher labor tax revenue.

Figure 3. Percent Change in Employment Moving from the Current Tax System to the Optimal Tax System

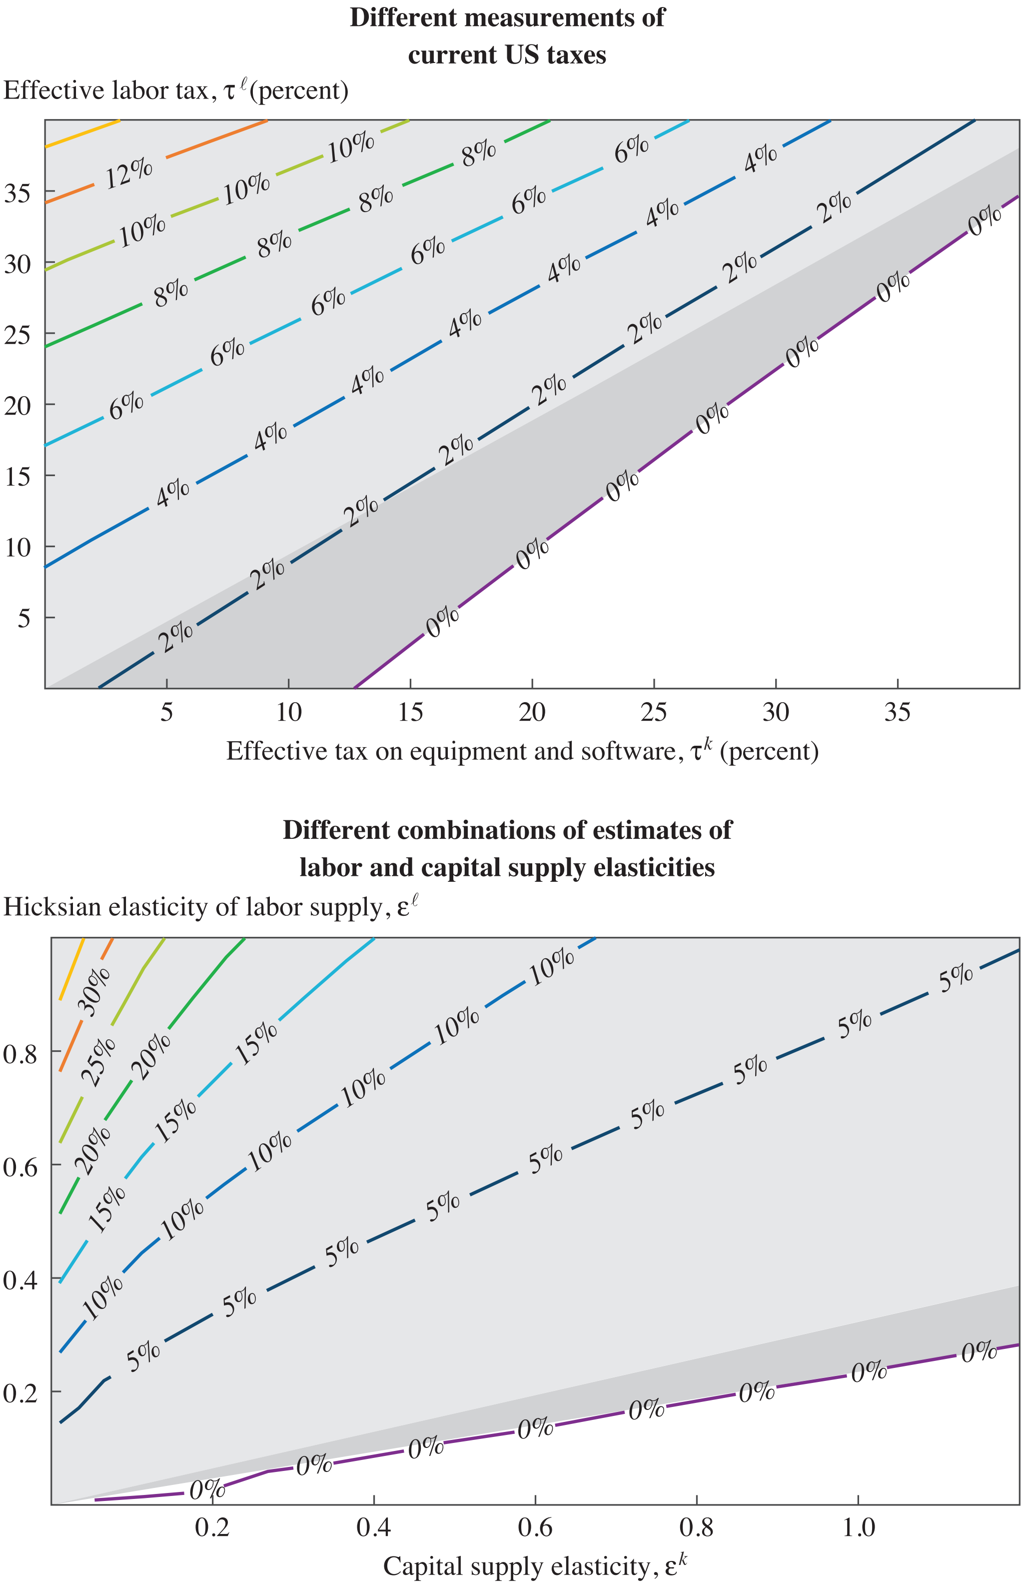

Source: Authors' calculations.

labor share by 1.93 percentage points. Even though equilibrium automation decisions are being distorted, aggregate net output remains essentially unchanged (it declines by 0.10 percent). As already noted, this is because marginal tasks automated under a biased tax system do not increase productivity much (or the automation technology being used in these tasks is so-so).

Column 3 in table 1 allows the planner to change θ, but without modifying the effective tax on capital, τ^k. We next verify that, as implied by proposition 3, if the planner could also modify τ^k (but could not reduce labor taxes), the planner would complement any reform with an automation tax to reduce automation below its market level. This is illustrated in column 4, which shows that in this case the planner achieves higher welfare through a combination of lower capital taxes (τ^k decreases to 8.39 percent) and an automation tax of 12.9 percent, which further reduces θ to 0.265. This alternative tax system would lead to a 2.44 percentage points higher labor share and 1.59 percent more employment.[35]

Finally, column 5 in table 1 turns to a setting where the planner can reduce taxes on labor and distort θ but cannot increase taxes on capital (as mentioned, this scenario may be relevant due to political constraints or fear of capital flight). In this case, the planner would combine a lower labor tax with an automation tax of 13.07 percent, reducing automation again to $\theta = 0.264$ and increasing employment by 1.96 percent and the labor share by 2.54 percentage points.[36]

In summary, our quantitative results show that the current tax system inefficiently favors automation and leads to an employment level that is below the social optimum. The best policy would be to set taxes at their optimal levels, which does not require any further distortions to automation. But if optimal taxes were infeasible, then reducing automation, with or without accompanying changes in other taxes, could reverse some of the inefficiencies in the current tax system and increase

---

[35] As mentioned above, reducing capital taxes may be optimal because the use of capital in tasks in which it has a comparative advantage benefits labor due to complementarity between tasks. In practice, capital might also complement labor in labor-intensive tasks. To capture this possibility, the task production function could be changed to $y(x) = \psi^{\ell}(x) \cdot \ell(x)^{\alpha} \cdot \tilde{k}(x)^{1-\alpha} + \psi^{k}(x) \cdot k(x)$, where $\tilde{k}(x)$ is the capital used to complement labor within tasks. Table A.5 in the online appendix shows that allowing for direct complementarities in this way (with $\alpha = 0.75$) does not change our main findings.

[36] Importantly, this can be implemented without raising *any* capital taxes. In particular, a tax on automation can also be implemented via a subsidy to labor of $\tau^A = 13.07$ percent combined with a tax of τ^A on the output of tasks above $\theta^c = 0.264$. This alternative implementation is discussed in proposition A.1 in the online appendix.

employment by 1.14–1.96 percent and the labor share in national income by 1.93–2.54 percentage points.

III.D. Recent Reforms and Effective Stimulus

As described in footnote 16, a series of reforms enacted between 2000 and the mid-2010s significantly reduced effective taxes on equipment and software (from about 20 percent in the year 2000 to about 10 percent). The Tax Cuts and Jobs Act of 2017, which came into effect in 2018, further reduced effective taxes on equipment and software to about 5 percent. These reforms aimed to raise employment by stimulating investment and overall economic activity. In this subsection, we use our calibrated model to study the effectiveness of these reforms and their implications for automation.

Our main finding is that, although all of these reforms increased employment (because they reduced effective taxes), their effects were limited and they increased employment at a large fiscal cost per job created, in large part because they encouraged additional automation. In contrast, we show that alternative reforms reducing labor taxes or combining lower capital taxes with an automation tax could have increased employment by more and at a much lower cost per job.

Column 1 of table 2 reports the market equilibrium for the capital and labor taxes in 2000—$\tau^{\ell} = 25.5$ percent and $\tau^{k} = 20$ percent. Column 2 then documents the impact of the tax cuts on capital enacted between 2000 and the mid-2010s, which reduced the effective tax on software and equipment to 10 percent and reduced government revenue by 10.49 percent. Our model implies that these tax cuts raised employment by a modest 1.01 percent, and did so at a substantial fiscal cost of $162,851 per job. As our theoretical analysis highlights, the lackluster employment response was in part because the lower taxes on capital encouraged greater automation, as shown by the increase in θ. Column 3 turns to the most recent (2017) tax cuts on capital. These are predicted to reduce government revenue by an additional 5.51 percent (or 16 percent relative to the revenue collected in 2000) and encourage further automation, with θ rising to 0.278. The resulting employment gain is again small, 1.47 percent relative to 2000 (or 0.46 percent relative to the mid-2010s), and comes at a fiscal cost per additional job of $169,857.

Columns 4–6 turn to alternative tax reforms that would have cost the same revenue as the capital tax cuts implemented between 2000 and the mid-2010s (10.49 percent of the year 2000 revenue). In column 4, we consider the implications of reducing labor taxes (for example, with a

Table 2. Comparison of Observed Tax Reforms and Reforms Costing the Same Revenue

		Observed reforms		*Alternative reforms*		
	System in 2000 with $\tau^k = 20\%$ (1)	*System in 2010s: reform to $\tau^k = 10\%$ (2)*	*System in 2018: reform to $\tau^k = 5\%$ (3)*	*Labor tax reform (4)*	*Capital tax reform with automation taxation (5)*	*Capital and labor tax reform (6)*
Tax system						
τ^k (%)	20.00	10.00	5.00	20.00	8.58	26.65
τ^ℓ (%)	25.50	25.50	25.50	21.09	25.50	18.22
θ	0.271	0.276	0.278	0.269	0.266	0.265
τ^A (%)	0.00	0.00	0.00	0.00	11.42	0.00
Aggregates						
Employment (%)	—	1.01	1.47	3.56	2.43	5.06
Capital (%)	—	5.64	8.34	1.14	3.07	–2.19
Labor share (%)	56.30	56.00	55.86	56.46	58.15	56.77
Net output (%)	—	2.38	3.49	2.83	2.55	2.83
Cost/revenue per job	—	$162,851	$169,857	$45,954	$67,316	$32,378
Revenue (%)	—	–10.49	–16.00	–10.49	–10.49	–10.49

Source: Authors' calculations.

Notes: This table shows the effective capital and labor taxes and the level of automation for different tax reforms. Column 1 presents the equilibrium under the tax system in the year 2000. Columns 2 and 3 present the resulting changes from the capital tax cuts enacted up to the mid-2010s and then the subsequent capital tax cuts enacted in 2017. Columns 4–6 show the effects of three alternative reforms that would have cost the same as the capital tax cuts enacted between 2000 and the mid-2010s. Column 4 considers cutting the effective labor tax. Column 5 considers a combination of capital tax cuts and a tax to automation. Column 6 considers a combination of lower labor taxes and higher capital taxes.

payroll tax cut) to $\tau^\ell = 21.09$ percent and keeping $\tau^k = 20$ percent as in 2000. This alternative reform would have increased employment by 3.56 percent and would achieve this at a quarter of the cost of one additional job in column 2. Part of the reason why reducing payroll taxes is much more effective in stimulating employment than cutting capital taxes is that lower payroll taxes reduce automation (θ falls to 0.269) whereas lower capital taxes further increase automation (θ increases from 0.272 to 0.276 between columns 1 and 2).

Column 5 considers another reform, this time combining lower capital taxes with an automation tax (again chosen to cost the same revenue as the tax cuts enacted between 2000 and the mid-2010s). This reform would have also raised employment by more than the reforms of the 2010s, increasing it by 2.43 percent, and would have cost $67,316 per job, which is less than half the cost per job in column 2. Notably, this policy combination involves an even larger tax cut for capital—from 20 percent to 8.58 percent. But crucially, the automation tax simultaneously rolls back any automation that the capital tax cut would have otherwise induced.[37]

Finally, column 6 considers a reform that changes both capital and labor taxes in a welfare maximizing way and costs the same revenue as the reform in column 2. By definition, this reform coincides with the Ramsey solution in column 2 of table 1, and it would have raised the effective capital tax rates to 26.65 percent and reduced the labor tax to 18.22 percent (eliminating the payroll tax almost entirely). We include it in this table to show that, in addition to the 5.06 percent additional increase in employment, such a reform would have had a much smaller cost per job—only $32,378, or about a fifth of the cost per job generated by the capital tax cuts since 2000.

Overall, this discussion shows that, because automation responds to the cost of capital, reducing capital taxes uniformly (via generous depreciation allowances or reductions in corporate taxes) is not an effective way of stimulating employment. Reforms over the last two decades that reduced capital taxes achieved only a modest increase in employment and instead encouraged further automation. Moving forward, reducing payroll taxes or accompanying tax cuts for capital with a tax on automation can more powerfully stimulate economic activity and achieve greater increases in employment at lower fiscal costs.

---

[37] A policy of reducing taxes on capital and at the same time taxing automation is equivalent to lowering the tax on capital by 11.42 percent only for tasks below $\theta = 0.266$. This exceeds the 10 percent tax cut from 2000 to the 2010s. These tax cuts for capital targeted at tasks in which it has a strong comparative advantage allow policymakers to give even larger subsidies to capital accumulation without triggering excessive automation.

III.E. Capital Distortions

Our analysis so far incorporates labor market imperfections via the wedge ϱ but ignores capital distortions. This is motivated by two considerations. First, our starting point is that because of labor market imperfections such as bargaining, search, or efficiency wages, even without any taxes, the level of employment would be too low; the wedge used in our model introduces this property in a simple way. Second, while earning losses from worker displacement provide a natural way of identifying the labor market wedge, there is no simple method for ascertaining whether there are capital wedges and how large they may be.[38] Nevertheless, we have carried out a number of exercises to verify that our conclusions are not unduly affected by this asymmetry in the treatment of capital and labor.

First, if equity finance is not subject to an additional distortion, then the deductions of interest rate payments from taxes in the case of debt finance more than undo any capital market distortions. This is because the interest rate on corporate loans is an upper bound on the capital wedge and is deducted from taxes. Therefore, we can conservatively use effective tax rates on equipment and software that would apply with full equity financing (without any of the reductions in effective capital taxes that come with debt finance). Table A.6 in the online appendix provides analogous results to table 1 in this case. The effective capital taxes are now $\tau^k = 12$ percent, but this has minimal effects on our results. Second, table A.7 in the online appendix repeats our main exercise but now assuming a capital wedge of $\varrho^k = 0.15$—the same as the labor wedge. The employment and welfare gains from moving to optimal taxes are still nontrivial even if about half as large as our baseline estimates. We conclude that our results are not driven by the assumption that there are no capital wedges or the asymmetric treatment of capital and labor.

IV. Extensions

In this section, we discuss two extensions that generalize our model and reinforce our main conclusion that the US tax code favors capital and promotes excessive automation.

---

[38] For example, large corporations that have significant cash at hand should not be using a different external rate of return than their internal rate of return, and their behavior should not be affected by a capital wedge, even if they use external funds. Smaller corporations may face a higher rate of return when borrowing funds, but if investment in these and larger corporations is highly substitutable, this may not correspond to an aggregate capital wedge.

IV.A. Human Capital Investments

The asymmetric treatment of capital and labor may further distort investments in human capital, which may interact with automation decisions. To incorporate this possibility, suppose that the efficiency unit of labor services provided by a worker is augmented by his or her human capital. Assume also that all workers have the same amount of human capital h, so that the efficiency units of labor are now $\ell_h = h \cdot \ell$.[39] The cost of investing in human capital h for ℓ workers is $\frac{\ell}{1 + 1/\varepsilon^h} \cdot h^{1+1/\varepsilon^h}$ in terms of the final good of the economy, and $\varepsilon^h > 0$. This parameter will be the elasticity of investment in human capital with respect to changes in wages. Likewise, we take the isoelastic specification of $\nu(\ell)$ used in our quantitative section, so that ε^ℓ is the constant Hicksian elasticity of labor supply.

Incorporating human capital into the labor market-clearing condition, we obtain

$$f_{\ell_h} \cdot (1 - \tau^\ell) \cdot (1 - \varrho) = \ell_h^{1/(\varepsilon^\ell + \varepsilon^h + \varepsilon^\ell \cdot \varepsilon^h)}.$$

The relevant elasticity for the supply of efficiency units of labor has now been replaced by $\varepsilon^\ell + \varepsilon^h + \varepsilon^\ell \cdot \varepsilon^h$, which incorporates the elastic response of human capital and is thus always greater than ε^ℓ. Intuitively, efficiency units of labor can be increased not just by supplying labor but by investing in human capital as well.

The next proposition characterizes optimal taxes in the presence of human capital and shows that labor taxes need to be adjusted to take into account the greater elasticity with which labor services respond to taxation. This pushes in the direction of (relatively) lower labor taxes and, conversely, higher capital taxes.

PROPOSITION 5: Optimal taxes with endogenous human capital

The solution to the Ramsey problem in an environment with human capital satisfies $\theta^r = \theta^m(k, \ell)$ and

$$\frac{\tau^{k,r}}{1 - \tau^{k,r}} = \frac{\mu}{1+\mu} \frac{1}{\varepsilon^k(k)} \quad \textit{and} \quad \frac{\tau^{\ell,r}}{1 - \tau^{\ell,r}} = \frac{\mu}{1+\mu} \frac{1}{\varepsilon^\ell + \varepsilon^h + \varepsilon^\ell \cdot \varepsilon^h} - \frac{\varrho}{1+\mu}.$$

---

[39] This formulation ignores the fact that high-human capital workers may be employed in tasks that are not automated or are complementary to automation technologies. The impact of automation on the employment and wages of different types of workers is explored in Autor, Levy, and Murnane (2003) and Acemoglu and Restrepo (in progress).

Moreover, if an economy has too low a tax on capital and excessive automation without human capital (in the sense of proposition 2), it will a fortiori have too low a tax on capital and excessive automation when there is an elastic response of human capital.

We next provide a back of the envelope quantification of the effect of human capital investments on optimal policy. To do this, we augment our analysis in the previous section with an estimate for the elasticity of human capital, ε^h. We set the elasticity of human capital supply, ε^h, to 0.092. This value is in the mid-range of estimates from the literature on high school completion (Jensen 2010; Kuka, Shenhav, and Shih 2018) and college major choice (Wiswall and Zafar 2015; Beffy, Fougère, and Maurel 2012).[40] This increases the supply elasticity of efficiency units of labor to 0.86, and as a result, the optimal labor tax is now lower, $\tau^\ell = 16.90$ percent, and the optimal capital tax is modestly higher, $\tau^k = 29.21$ percent (see table A.9 in the online appendix). Replacing the current system with optimal taxes leads to more pronounced changes: 1.06 percentage points higher labor share, 5.73 percent increase in employment, and 0.59 percent increase in welfare in consumption-equivalent terms.

IV.B. Endogenous Technology

In our baseline model, increases in θ represent both the development and the adoption of automation technologies. In principle, these two decisions are distinct, even if related. Unless automation technologies are developed, they cannot be adopted. If they are expected to be adopted, then there are greater incentives to develop them. Moreover, as emphasized in Acemoglu and Restrepo (2018), new automation technologies may come at the expense of other technological changes with different implications for capital and labor. For instance, more resources devoted to automation typically imply less effort toward the introduction of new tasks that tend to increase the labor share and demand for labor. If so, a tax structure that favors capital may distort the direction of technological change in a way that disadvantages labor. In this subsection, we provide a simple model to

[40] Jensen's (2010) experimental results imply a 0.097 high school completion elasticity in response to perceived returns. Kuka, Shenhav, and Shih (2018) estimate a high school completion elasticity of 0.019–0.086 in response to actual returns and 0.014–0.17 in response to perceived returns. Wiswall and Zafar (2015) estimate elasticities in the range of 0.036–0.062 from the response of college major choice to changes in relative wage premium. Previous estimates in Beffy, Fougère, and Maurel (2012) put the same elasticity in the range of 0.09–0.12. Taken together, these studies imply values for ε^h in the range of 0.014–0.17.

highlight these ideas and show that, with endogenous technology, optimal policy may also need to redirect the direction of technological change, and this is the case even when capital and labor taxes are set optimally.

For brevity, we borrow from the formulation of endogenous technology in Acemoglu (2007, 2010), whereby a (competitive) production sector decides how much capital and labor to use and which technology, from a menu of available technologies, to utilize, while a monopolistically competitive (or simply monopolistic) technology sector decides which menu of technologies to develop and offer to firms.

We consider a menu of technologies consisting of both automation techniques and technologies that increase the productivity and the set of tasks performed by labor, such as the introduction of new tasks considered in Acemoglu and Restrepo (2018). We summarize this menu by Θ with the convention that a higher Θ means a menu that is more biased toward automation technologies. Given menu Θ, firms choose their level of automation θ and their utilization of other technologies ω subject to the feasibility constraint $G(\theta, \omega; \Theta) \leq 0$. Therefore, the index of technologies Θ determines what combinations of automation and other technologies are feasible for final good producers. We denote the production function given θ and ω by $f(k, \ell; \theta, \omega)$, and assume that f_k/f_ℓ is increasing in θ as in our baseline model and decreasing in ω. We further assume that when Θ increases, the set $G(\theta, \omega; \Theta) \leq 0$ includes higher values of θ and lower values of ω, so that a higher Θ enables more adoption of automation technologies and less adoption of other technologies.

The profit-maximizing adoption decision solves the following problem:

$$\{\theta^m(k, \ell; \Theta), \omega^m(k, \ell; \Theta)\} = arg\ max_{G(\theta,\omega;\Theta)\leq 0} f(k, \ell; \theta, \omega).$$

The assumptions on $G(\theta, \omega; \Theta)$ imply that $\omega^m(k, \ell; \Theta)$ is decreasing in Θ and $\theta^m(k, \ell; \Theta)$ is increasing in Θ—so that a higher Θ means a menu of technologies that is more biased toward automation. Hence, as the menu of available technologies becomes more biased toward automation, it crowds out the adoption of nonautomation technologies (such as new tasks or others that increase human productivity).

Finally, we assume that the technology sector charges markups for the use of technologies by final good producers and, via this, captures a constant fraction $\kappa \in (0, 1)$ of the output of these producers (this could be microfounded by assuming that the technology sector sells machines embedding the new technology with a constant markup; see Acemoglu 2007, 2010). We denote the cost of choosing a menu of technologies

Θ by $\Gamma(\Theta)$. Thus, the maximization problem of the technology sector that determines the equilibrium bias of technology is

$$\max_{\Theta} \kappa \cdot f(k, \ell; \theta^m(k, \ell; \Theta), \omega^m(k, \ell; \Theta)) - \Gamma(\Theta). \tag{20}$$

We make the following assumptions on $\Gamma(\Theta)$:

— $\Gamma(\Theta)$ has a minimum at $\bar{\Theta} \in (0, 1)$. This assumption means that there exists a baseline bias of technology $\bar{\Theta}$, such that deviations from this baseline involve increasing costs. More specifically, deviations from $\bar{\Theta}$ can come in the direction of further automation or further effort devoted to creating new tasks (and thus less automation). Both of these will be more costly than continuing with $\bar{\Theta}$. In the dynamic framework of Acemoglu and Restrepo (2018), $\bar{\Theta}$ corresponds to the state of technology inherited from the past.

— $\Gamma(\Theta)$ is convex, which captures diminishing returns in research directed at changing the bias of technology away from the baseline level $\bar{\Theta}$.

In addition to capital and labor taxes, we allow for subsidies to the use of automation and other technologies in the final good sector to undo the effects of the markup κ and for taxes on the profits of the technology monopolist. Our results do not depend on whether such additional taxes and subsidies exist, but their presence simplifies the expressions and makes them much more closely connected to those in our baseline model in section I.

A market equilibrium satisfies the same market-clearing conditions as in our benchmark economy but is augmented to include the fact that technology adoption decisions of final good producers are given by $\theta^m(k, \ell; \Theta)$ and $\omega^m(k, \ell; \Theta)$, and the equilibrium bias of technology Θ maximizes equation (20). We assume that a market equilibrium exists and is unique and that the solution to equation (20) always involves some interior $\Theta \in (0, 1)$.

We next characterize the solution to the Ramsey problem as in proposition 1. As in our baseline model, we assume that the planner directly controls the development and adoption of technologies (these choices can be implemented with additional taxes as in section I.E).

PROPOSITION 6: Optimal taxes and automation with endogenous technology

The solution to the Ramsey problem with endogenous technology involves capital and labor taxes given as in equation (8) and undistorted adoption decisions (conditional on Θ) given by $\theta^r(k, \ell; \Theta) = \theta^m(k, \ell; \Theta)$ and $\omega^r(k, \ell; \Theta) = \omega^m(k, \ell; \Theta)$. However, if $\Theta^r \gtreqless \bar{\Theta}$, the optimal bias of technology satisfies $\Theta^r \gtreqless \Theta^m$ (i.e., the optimal and market bias of technology are the same if and only if $\Theta^r = \bar{\Theta}$).

The most important implication of this proposition is that, even with optimal taxes on capital and labor, the planner might wish to discourage the development of automation technology. This will be the case when the baseline level of technology is more geared toward automation than what the planner would like to achieve. Put differently, if the economy in question has already gone in the direction of excessively developing automation technologies (which may be a consequence of past distortions or other factors influencing the direction of past technological change), then the planner should intervene by distorting the direction of innovation. The reason for this is that the technology sector does not fully internalize the social surplus created by its technology choices, because of the presence of the term $\kappa < 1$ in equation (20), and thus will not develop the right type of technologies. This result has a close connection to one of the key insights in Acemoglu and others (2012), which established, in the context of optimal climate change policy, that if the economy starts with relatively advanced carbon-emitting, dirty technologies and relatively backward low-carbon, clean technologies, then it is not sufficient to impose Pigovian taxes; rather, optimal policy additionally calls for direct subsidies to the development of clean technologies.[41]

This result has important implications in our context. As our findings in section III suggest, past US tax policy has favored capital and automation. Because these policies have likely led to excessive development of automation technologies, it is not sufficient to simply redress the distortions in the current tax system. Instead, optimal policy may need to intervene to redirect technological change by subsidizing the creation of new tasks and temporarily discourage further effort toward automation innovations at the margin. We leave a quantitative exploration of the implications of endogenous technology development to future work.

V. Concluding Remarks

Automation is transforming labor markets and the structure of work in many economies around the world, not least in the United States. The number of robots in industrial applications and the use of specialized software, artificial intelligence, and several other automation technologies have increased rapidly in the US economy over the last few decades. There has been a concomitant decline in the labor share of national income, wages

---

[41] Note in addition that once the planner can influence the direction of automation technology and set optimal taxes on capital and labor, there is no need to distort the adoption of automation technologies.

have stagnated, and low-skill workers have seen their real wages decline. Many experts believe that these labor market trends are, at least in part, related to automation.

The general intuition among economists (and many policymakers) is that even if automation may have some adverse distributional and employment consequences, policy should not slow down (and certainly not prevent) the adoption of automation technologies because these technologies are contributing to productivity. According to this perspective, policy should instead focus on fiscal redistribution, education, and training to ensure more equally distributed gains and more opportunities for social mobility. But what if automation is excessive from a social point of view?

This paper has argued that the US tax system is likely to be encouraging excessive automation, and if so, reducing the extent of automation (or, more plausibly, slowing down the development and adoption of new automation technologies) may be welfare-improving. We have developed this argument in three steps.

First, we revisited the theory of optimal capital and labor taxation in a task-based framework where there is an explicit decision of firms to automate tasks. We also introduced, albeit in a reduced form, labor market imperfections. Consistent with the classical theory of public finance, if capital and labor taxes are set optimally, automation decisions are optimal in equilibrium. However, away from optimal capital and labor taxes or in the presence of additional constraints on tax decisions, this is no longer the case. Exploiting the structure of our task-based framework, we establish that when the tax system is already biased against labor, it is generally optimal to distort equilibrium automation. The economics of this result is simple but informative: marginal tasks that are automated bring little productivity gains—or in the terminology of Acemoglu and Restrepo (2019a, 2019b), they are "so-so automation technologies"—and as a result, the cost of reducing automation at the margin is second order. When the tax system is biased against labor, the gain from reducing automation and preventing the displacement of labor is first order because it increases employment. In fact, it may even be optimal to reduce automation while at the same time cutting capital taxes (even though the tax system is biased against labor and in favor of capital) because, in contrast to automation, the use of capital in tasks in which capital has a strong comparative advantage is complementary to workers employed in labor-intensive tasks.

Second, we delved into a detailed evaluation of the US tax system in order to map the complex tax code into effective capital and labor taxes. Our numbers suggest that the US system taxes labor heavily and favors capital significantly. While labor is taxed at an effective rate between 25.5 percent and 33.5 percent, capital faces an effective tax rate of about 5 percent (down from 10 percent in the 2010s and 20 percent in the 1990s and early 2000s).

Third, we compared the US tax system to the optimal taxes implied by our theoretical analysis. This exercise confirmed that the US tax system is biased against labor and in favor of capital. As a result, we found that moving from the current US tax system and level of automation to optimal taxation of factors and the optimal level of automation would raise employment by 4.02 percent, the labor share by 0.78 percentage point, and overall welfare by 0.38 percent in consumption-equivalent terms. If optimal taxes can be implemented, there is no need for distorting or taxing automation. If, on the other hand, optimal taxes are infeasible, more modest reforms involving a tax on automation can undo some of the inefficiencies in the current system and increase employment by 1.14–1.96 percent and the labor share by 1.93–2.54 percentage points. In this case, the constrained optimal policy always involves an automation tax in order to discourage the automation of marginal tasks which bring little productivity benefits but significant displacement of labor.

We also showed that a range of realistic generalizations (absent from our baseline framework) reinforce our conclusions and call for even more extensive changes in automation and capital taxation, and under some conditions, it may be optimal to redirect new innovations away from automation.

To simplify the analysis, we focused on an economy with a single type of labor. As noted at the beginning, automation is also associated with increases in inequality (Autor, Levy, and Murnane 2003; Acemoglu and Autor 2011; Acemoglu and Restrepo 2020, in progress). Consequently, slowing down automation may generate additional distributional benefits. These issues are discussed in Guerreiro, Rebelo, and Teles (2017), Thuemmel (2018), and Costinot and Werning (2018). A natural next step is to augment these analyses with the possibility that certain aspects of the tax system may be encouraging excessive automation.

In practice, there are many potential sources of excessive automation. Our objective in this paper has been narrow: to focus on tax reasons for excessive automation. Our companion paper, Acemoglu, Manera, and Restrepo (in progress), shows that even absent tax-related distortions, the market economy tends to generate excessive automation because bargaining power and efficiency wage considerations vary across tasks and this tends to create incentives for firms to automate beyond what is socially beneficial in order to improve their share of rents. Furthermore, as we have already noted, automation-driven job loss may generate negative spillovers on communities and political and social behavior. There may additionally be social factors and corporate strategies concerning the direction of innovation and research (the best minds in many fields being attracted to automation technologies and the most influential companies favoring automation) that further contribute to excessive automation. Quantifying the extent of these other factors is an important area for future research, especially because they have major implications for policy.

Finally, we should briefly comment on how our results relate to two popular policy proposals: wealth taxes and so-called robot taxes. Although our framework suggests that it may be beneficial to increase taxes on capital, wealth taxes on high-wealth individuals may not be the most direct way of achieving this because they would not necessarily increase the effective tax on the use of capital. Increasing corporate income taxes and eliminating or lowering depreciation allowances may be more straightforward ways of implementing higher effective taxes on capital (provided that there are no other distributional or political benefits from wealth taxes). Moreover, our framework highlights that it is not always beneficial to increase taxes on capital: when it is not feasible to implement optimal taxes, reducing automation becomes a central objective (and may even need to be combined with lower taxes on capital). Our automation tax is also different from taxes on robots for the same reasons: it is not a uniform tax on all automation technologies; rather, it is applied to technologies automating tasks above a certain threshold (which are tasks in which humans still have a significant comparative advantage). In fact, our results clarify that instead of taxing all automation technologies, optimal policy often involves subsidizing capital in tasks in which machines have a strong comparative advantage. Last, our analysis also clarifies that if the tax system is reformed so that it is no longer biased against labor and in favor of capital, then employment and welfare can be increased without an automation tax.

ACKNOWLEDGMENTS Prepared for the Brookings Institution Spring Conference of 2020. We are grateful to Alan Auerbach, Jan Eberly, Larry Katz, James Poterba, Ivan Werning, Owen Zidar, and Eric Zwick for excellent comments and suggestions. We gratefully acknowledge financial support from Brookings Institution, Google, the National Science Foundation, Schmidt Sciences, the Smith Richardson Foundation, and the Sloan Foundation.

References

Acemoglu, Daron. 2007. “Equilibrium Bias of Technology.” *Econometrica* 75, no. 5: 1371–409.

Acemoglu, Daron. 2010. “When Does Labor Scarcity Encourage Innovation?” *Journal of Political Economy* 118, no. 6: 1037–78.

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous. 2012. “The Environment and Directed Technical Change.” *American Economic Review* 102, no. 1: 131–66.

Acemoglu, Daron, and David Autor. 2011. “Skills, Tasks and Technologies: Implications for Employment and Earnings.” In *Handbook of Labor Economics*, vol. 4, edited by Orley Ashenfelter and David Card. Amsterdam: Elsevier.

Acemoglu, Daron, David Autor, Joe Hazell, and Pascual Restrepo. 2020. “AI and Jobs: Evidence from Online Vacancies.” Mimeo. Massachusetts Institute of Technology.

Acemoglu, Daron, Andrea Manera, and Pascual Restrepo. In progress. “Automation in Imperfect Labor Markets.” Massachusetts Institute of Technology.

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, no. 6: 1488–542.

Acemoglu, Daron, and Pascual Restrepo. 2019a. “Artificial Intelligence, Automation, and Work.” In *The Economics of Artificial Intelligence: An Agenda*, edited by Ajay Agrawal, Joshua Gans, and Avi Goldfarb. Chicago: University of Chicago Press.

Acemoglu, Daron, and Pascual Restrepo. 2019b. “Automation and New Tasks: How Technology Displaces and Reinstates Labor.” *Journal of Economic Perspectives* 33, no. 2: 3–30.

Acemoglu, Daron, and Pascual Restrepo. 2020. “Robots and Jobs: Evidence from US Labor Markets.” *Journal of Political Economy* 128, no. 6: 2188–244.

Acemoglu, Daron, and Pascual Restrepo. In progress. “The Labor Share, Displacement and Wage Inequality.” Mimeo. Massachusetts Institute of Technology.

Atkinson, A. B., and J. E. Stiglitz. 1972. “The Structure of Indirect Taxation and Economic Efficiency.” *Journal of Public Economics* 1, no. 1: 97–119.

Atkinson, A. B., and J. E. Stiglitz. 1976. “The Design of Tax Structure: Direct versus Indirect Taxation.” *Journal of Public Economics* 6, no. 1: 55–75.

Auerbach, Alan, and Laurence Kotlikoff. 1987. *Dynamic Fiscal Policy*. Cambridge: Cambridge University Press.

Austin, Benjamin, Edward Glaeser, and Lawrence Summers. 2018. “Jobs for the Heartland: Place-Based Policies in 21st-Century America.” *Brookings Papers on Economic Activity*, Spring, 151–232.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. “The China Syndrome: Local Labor Market Effects of Import Competition in the United States.” *American Economic Review* 103, no. 6: 2121–68.

Autor, David H., and Mark G. Duggan. 2003. "The Rise in the Disability Rolls and the Decline in Unemployment." *Quarterly Journal of Economics* 118, no. 1: 157–206.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* 118, 4: 1279–333.

Beffy, Magali, Denis Fougère, and Arnaud Maurel. 2012. "Choosing the Field of Study in Postsecondary Education Do Expected Earnings Matter?" *Review of Economics and Statistics* 94, no. 1: 334–47.

Brülhart, Marius, Jonathan Gruber, Matthias Krapf, and Kurt Schmidheiny. 2016. "Taxing Wealth: Evidence from Switzerland." Working Paper 22376. Cambridge, Mass.: National Bureau of Economic Research.

Chamley, Christophe. 1986. "Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives." *Econometrica* 54, no. 3: 607–22.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber. 2011. "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins." *American Economic Review* 101, no. 3: 471–75.

Congressional Budget Office (CBO). 2014. "Taxing Capital Income: Effective Marginal Tax Rates Under 2014 Law and Selected Policy Options." Washington. https://www.cbo.gov/sites/default/files/113th-congress-2013–2014/reports/49817–taxingcapitalincome0.pdf.

Cooper, Michael, John McClelland, James Pearce, Richard Prisinzano, Joseph Sullivan, Danny Yagan, Owen Zidar, and Eric Zwick. 2016. "Business in the United States: Who Owns It, and How Much Tax Do They Pay?" In *Tax Policy and the Economy*, vol. 30, edited by Robert A. Moffitt. Chicago: University of Chicago Press.

Costinot, Arnaud, and Iván Werning. 2018. "Robots, Trade, and Luddism: A Sufficient Statistic Approach to Optimal Technology Regulation." Working Paper 25103. Cambridge, Mass.: National Bureau of Economic Research.

Couch, Kenneth A., and Dana W. Placzek. 2010. "Earnings Losses of Displaced Workers Revisited." *American Economic Review* 100, no. 1: 572–89.

Davis, Steven J., and Till von Wachter. 2011. "Recessions and the Costs of Job Loss." *Brookings Papers on Economic Activity*, Fall, 1–72.

Diamond, Peter A., and James A. Mirrlees. 1971. "Optimal Taxation and Public Production I: Production Efficiency." *American Economic Review* 61, no. 1: 8–27.

Durán-Cabré, José M., Alejandro Esteller-Moré, and Mariona Mas-Montserrat. 2019. "Behavioural Responses to the (Re)Introduction of Wealth Taxes: Evidence from Spain." Working Paper 2019/4. Barcelona: Institut d'Economia de Barcelona.

Edgerton, Jesse. 2010. "Investment Incentives and Corporate Tax Asymmetries." *Journal of Public Economics* 94, nos. 11–12: 936–52.

Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin. 2013. "The Decline of the U.S. Labor Share." *Brookings Papers on Economic Activity*, Fall, 1–63.

Garrett, Daniel G., Eric Ohrn, and Juan Carlos Suárez Serrato. 2020. "Tax Policy and Local Labor Market Behavior." *American Economic Review: Insights* 2, no. 1: 83–100.

Goldman, Dana, Neeraj Sood, and Arleen Leibowitz. 2005. "Wage and Benefit Changes in Response to Rising Health Insurance Costs." Working Paper 11063. Cambridge, Mass.: National Bureau of Economic Research.

Goolsbee, Austan. 1998. "Investment Tax Incentives, Prices, and the Supply of Capital Goods." *Quarterly Journal of Economics* 113, no. 1: 121–48.

Gruber, Jonathan, and Alan B. Krueger. 1991. "The Incidence of Mandated Employer-Provided Insurance: Lessons from Workers' Compensation Insurance." In *Tax Policy and the Economy*, vol. 5, edited by David F. Bradford. Cambridge, Mass.: MIT Press.

Guerreiro, Joao, Sergio Rebelo, and Pedro Teles. 2017. "Should Robots Be Taxed?" Working Paper 23806. Cambridge, Mass.: National Bureau of Economic Research.

Hall, Robert E., and Dale W. Jorgenson. 1967. "Tax Policy and Investment Behavior." *American Economic Review* 57, no. 3: 391–414.

Hassett, Kevin A., and R. Glenn Hubbard. 2002. "Tax Policy and Business Investment." In *Handbook of Public Economics*, vol. 3, edited by Alan J. Auerbach and Martin Feldstein. Amsterdam: Elsevier.

House, Christopher L., and Matthew D. Shapiro. 2008. "Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation." *American Economic Review* 98, no. 3: 737–68.

Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan. 1993. "Earnings Losses of Displaced Workers." *American Economic Review* 83, no. 4: 685–709.

Jakobsen, Katrine, Kristian Jakobsen, Henrik Kleven, and Gabriel Zucman. 2020. "Wealth Taxation and Wealth Accumulation: Theory and Evidence From Denmark." *Quarterly Journal of Economics* 135, no. 1: 329–88.

Jensen, Robert. 2010. "The (Perceived) Returns to Education and the Demand for Schooling." *Quarterly Journal of Economics* 125, no. 2: 515–48.

Joulfaian, David. 2006. "The Behavioral Response of Wealth Accumulation to Estate Taxation: Time Series Evidence." *National Tax Journal* 59, no. 2: 253–68.

Judd, Kenneth L. 1985. "Redistributive Taxation in a Simple Perfect Foresight Model." *Journal of Public Economics* 28, no. 1: 59–83.

Judd, Kenneth L. 1999. "Optimal Taxation and Spending in General Competitive Growth Models." *Journal of Public Economics* 71, no. 1: 1–26.

Karabarbounis, Loukas, and Brent Neiman. 2014. "The Global Decline of the Labor Share." *Quarterly Journal of Economics* 129, no. 1: 61–103.

Kleven, Henrik Jacobsen, and Esben Anton Schultz. 2014. "Estimating Taxable Income Responses Using Danish Tax Reforms." *American Economic Journal: Economic Policy* 6, no. 4: 271–301.

Kopczuk, Wojciech, and Joel Slemrod. 2000. “The Impact of the Estate Tax on the Wealth Accumulation and Avoidance Behavior of Donors.” Working Paper 7960. Cambridge, Mass.: National Bureau of Economic Research.

Kuka, Elira, Na’ama Shenhav, and Kevin Shih. 2018. “Do Human Capital Decisions Respond to the Returns to Education? Evidence from DACA.” Working Paper 24315. Cambridge, Mass.: National Bureau of Economic Research.

Lennon, Conor. 2020. “Are the Costs of Employer-Sponsored Health Insurance Passed on to Workers at the Individual Level?” Working Paper. University of Louisville.

Mui, Preston, and Benjamin Schoefer. 2019. “The Labor Supply Curve at the Extensive Margin: A Reservation Wedge Approach.” Paper presented at the Society for Economic Dynamics Annual Meeting, St. Louis, Mo., June 28.

Oberfield, Ezra, and Devesh Raval. 2014. “Micro Data and Macro Technology.” Working Paper 20452. Cambridge, Mass.: National Bureau of Economic Research.

Smith, Matthew, Danny Yagan, Owen Zidar, and Eric Zwick. 2019. “Capitalists in the Twenty-First Century.” *Quarterly Journal of Economics* 134, no. 4: 1675–745.

Straub, Ludwig, and Iván Werning. 2020. “Positive Long-Run Capital Taxation: Chamley-Judd Revisited.” *American Economic Review* 110, no. 1: 86–119.

Suárez Serrato, Juan Carlos, and Owen Zidar. 2016. “Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms.” *American Economic Review* 106, no. 9: 2582–624.

Thuemmel, Uwe. 2018. “Optimal Taxation of Robots.” 7317. CESifo Working Paper Series. CESifo Group Munich.

Wiswall, M., and B. Zafar. 2015. “Determinants of College Major Choice: Identification Using an Information Experiment.” *Review of Economic Studies* 82, no. 2: 791–824.

Yagan, Danny. 2015. “Capital Tax Reform and the Real Economy: The Effects of the 2003 Dividend Tax Cut.” *American Economic Review* 105, no. 12: 3531–63.

Zeira, Joseph. 1998. “Workers, Machines, and Economic Growth.” *Quarterly Journal of Economics* 113, no. 4: 1091–117.

Zoutman, Floris T. 2018. “The Elasticity of Taxable Wealth: Evidence from the Netherlands.” Working Paper.

Zwick, Eric, and James Mahon. 2017. “Tax Policy and Heterogeneous Investment Behavior.” *American Economic Review* 107, no. 1: 217–48.

Comments and Discussion

COMMENT BY

LAWRENCE F. KATZ Daron Acemoglu, Andrea Manera, and Pascual Restrepo have reassessed the theory of optimal capital and labor taxation in an elegant task-based framework in which firms make explicit decisions about whether or not to automate tasks. The framework also can account for labor market and capital market imperfections. The authors find that the optimal labor tax rate should be lower than the optimal capital tax rate under (what they argue are) realistic conditions for the US economy of similar effective capital and labor supply elasticities and greater labor market imperfections than capital market imperfections (e.g., a positive labor wedge and no capital wedge). In contrast, the authors carefully document that the actual US effective labor tax rate in the 2010s of 25.5–33.5 percent is much higher than the effective capital tax rate of 5–10 percent for equipment and software. Thus, they conclude that the US tax system is distorted against labor. Furthermore, their framework implies that a tax system biased against labor generates excessive automation with marginal automation being “so-so automation” with only second-order productivity gains but first-order welfare losses from labor displacement. The US tax system’s bias against labor thereby leads to lower employment and a lower labor share of national income than optimal in their framework. Finally, the authors argue that if political constraints or international capital flight concerns prevent raising capital taxes to the optimal level, then the second-best policy response is to try to limit automation directly or impose an automation tax to reduce excessive so-so automation and raise employment and welfare.

The authors have produced a provocative, creative, and impressive analysis that certainly makes one rethink how technological advances can potentially harm workers when they are task-replacing automation as opposed to labor-augmenting technological change of the type assumed in standard Solow growth models. And the authors show how a tax system designed (or lobbied) to encourage capital investment can distort automation decisions and lead to too much automation and too little employment. Nevertheless, I do have some concerns about the practical feasibility of their second-best policy of automation restrictions or automation taxes only for capital investment to replace tasks beyond some automation threshold where labor still has a comparative advantage. And I have questions concerning the implications of spillovers across firms or network effects of automation advances as well as how open economy considerations have an impact on the analysis.

In table 1, the authors simulate the potential benefits of task-specific automation taxes that increase capital taxation on automation investments in tasks where labor still has a comparative advantage. But it is not clear how policymakers and tax authorities can identify the task-specific impact of different capital investments. The authors argue that their recommended automation taxes differ from (uniform) robot taxes in only taxing automation above a threshold where it is strongly labor displacing and not taxing capital investments in tasks where machines have a strong comparative advantage. In practice, the US tax system can differentiate taxes on capital assets by asset life and asset categories (structures versus equipment versus software), but it is not obvious how these distinctions translate into different impacts on different types of tasks. More research is likely needed on how investments in different types of capital assets have an impact on different types of worker tasks, such as looking at impacts by occupation and using Bureau of Labor Statistics Occupational Information Network (BLS O*NET) task measures by occupation to determine how differing taxes by capital asset classes might be fine-tuned to more closely approximate an automation tax to discourage marginal (so-so) automation along the lines suggested by the authors. Still, one worries that this will be a daunting task and that even within detailed asset classes (e.g., software) some investments may be labor augmenting and some strongly displacing. And differences in tax rates by asset classes that may be similar to each other based on expected labor displacement effects could generate a whole new range of tax shenanigans through the relabeling of different capital assets to avoid the automation tax or to get favored tax treatment.

Another way to target the automation displacement threshold might be to link capital tax rates to whether a firm keeps employment stable or increases employment as opposed to displacing workers after new investments. But so many difficult-to-measure factors have an impact on employment decisions that again one worries that such a policy could generate unintended distortions. Thus, a more direct and feasible policy approach might be to focus on labor subsidies for workers more at risk of displacement from automation. Such policies could end up looking more like traditional policy responses of wage subsidies for disadvantaged and less-skilled workers or education and training policies to help at-risk workers become more complementary to new technologies.

Open economy considerations also need to be integrated into the authors' framework and policy analysis. Automation that generates only modest productivity gains or cost savings could lead to substantial job displacement in a closed economy setting but could be essential to maintain or expand domestic employment in internationally competitive industries. In fact, Aghion and others (2020) find that for France automation investments significantly increase industry employment (even for less-skilled workers) in industries facing international competition but not in other industries. The implication is that policies that try to limit marginal automation to increase domestic employment could have the opposite effect in the face of foreign competition.

A further difficulty relates to how to identify marginal (or so-so) automation that displaces workers and has only second-order productivity benefits from breakthrough automation investments that could have spillovers to other firms or help create new markets, new tasks, and employment opportunities as emphasized by Acemoglu and Restrepo (2019). For example, one might conclude that autonomous long-haul trucks will end up having only so-so automation, displacing many long-haul truck drivers with only small productivity benefits, and use the tax system to discourage adoption of autonomous vehicles. On the other hand, autonomous long-haul trucks could initially spur new infrastructure and complementary public and private investments that change the nature of interstate highways and linkages of long-haul and short-haul trucking to massively reduce longer-distance domestic transport costs or increase transport speed for products now supplied only locally (e.g., perishable products or artisanal goods), creating new markets, more short-haul trucking and logistics jobs, and overall improved employment activities. Public infrastructure investments, R & D policies, education and training policies, and technology extension policies may alter the impacts of automation such that so-so automation for individual firms that appears labor displacing might be beneficial for workers in the broader economy. The more positive impacts of automation on productivity growth and worker outcomes in the mid-twentieth century of shared prosperity and stable labor share as opposed to the twenty-first century of declining labor share and weak wage growth might have more to do with such complementary policies than changes in tax treatment of capital versus labor.

Finally, the authors' conclusion that optimal labor taxes should be lower than optimal capital taxes and that the US tax system is distorted against labor, in the sense of their equation (9), depends on the capital and labor supply elasticities being of similar magnitude or of labor market distortions being larger than capital market distortions. The authors admirably present a lot of robustness checks to different values for the labor and capital supply elasticities and the size of labor and capital market distortions. But I do remain concerned that the large international mobility of capital means that the effective capital supply elasticity to US taxes could be in the high range of their estimates (even greater than one) and that effective labor supply elasticities could be below the range they consider. Further disaggregation across education, age, and gender groups in labor markets with different labor supply elasticities could help sharpen the analysis. And the large decline in the worker power in the United States documented by Stansbury and Summers (2020) over recent decades implies a substantial reduction in the labor wedge that could be of similar magnitude to (and thereby offset) the decline in capital taxes versus labor taxes and should have operated to ameliorate the problem of excessive automation in the authors' framework from tax distortions.

REFERENCES FOR THE KATZ COMMENT

Acemoglu, Daron, and Pascual Restrepo. 2019. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33, no. 2: 3–30.

Aghion, Philippe, Céline Antonin, Simon Bunel, and Xavier Jaravel. 2020. "What Are the Labor Market and Product Market Effects of Automation? New Evidence from France." Working Paper. https://www.xavierjaravel.com/papers.

Stansbury, Anna, and Lawrence H. Summers. 2020. "The Declining Worker Power Hypothesis: An Explanation for the Recent Evolution of the American Economy." In the present volume of *Brookings Papers on Economic Activity*.

COMMENT BY

ERIC ZWICK Acemoglu, Manera, and Restrepo provide a thought-provoking perspective on capital taxes and automation. The paper is clearly written and the model's logic is intuitive. I especially appreciate their careful attention to detail in mapping the model onto its empirical analogues in the US tax system. My comments focus on capital tax questions and quibbles, both in mapping the model to the data and in applying the results to inform capital tax policy. To me, the question—Where does automation rank relative to other considerations in evaluating tax policy?—remains open. But I suspect we have not heard the last from this team. In the words of a famous robot from popular culture, "[They'll] be back!"

EVALUATING KEY CAPITAL MARKETS ASSUMPTIONS The paper's primary contribution is a set of theoretical results on the potential for welfare-improving taxation of automated tasks or complementary capital. These results are then calibrated under certain assumptions to allow quantitative statements (1) on whether the tax code suboptimally favors automation and (2) on what optimal tax rates on labor and capital should be. Under the offered set of assumptions, the model's results hold. My first comment concerns which assumptions are important for the main results and whether these assumptions are suitably chosen. I focus on three parameters of interest: the elasticity of capital supply, the capital markets wedge, and the effective capital tax rate.

First, the model's results accord with a Ramsey rule intuition that optimal taxes are inversely proportional to the elasticities of capital and labor supply. Thus, the relative elasticity of capital versus labor is crucial for the quantitative exercise. The authors' baseline parameterization sets the capital supply elasticity to 0.65, slightly below the labor supply elasticity of 0.7. The source for this assumption is recent quasi-experimental research on wealth and savings taxes, for example, the recent paper by Jakobsen and others (2020) which studies a recent wealth tax reform in Sweden.

I believe the conventional wisdom on labor versus capital remains that capital supply is more elastic than labor supply, so I was surprised to see them calibrated to be about the same. Taking estimates from the wealth tax literature to this setting requires a nontrivial extrapolation, namely, that the local effects estimated for a small subpopulation can inform the aggregate capital supply elasticity.

Moreover, the wealth and savings tax literature are not the only useful sources for such estimates. For example, House and Shapiro (2008) use the first round of bonus depreciation incentives to estimate the elasticity of capital supply, subject to the assumption that demand elasticities for long-lived goods in response to temporary subsidies are infinite. Their capital supply elasticities range from 6 to 14, an order of magnitude larger than the assumed elasticity here. Of course, the tradition in theoretical corporate tax incidence going back to Harberger (1962) and Feldstein (1974) has been to permit the capital supply elasticity to approach infinity in the medium and long runs. While I believe finite capital supply elasticities are well justified and supported in the data, the bottom line is that the working assumption here appears to me to be nonstandard, potentially controversial, and quantitatively relevant.

Second, the core theoretical results follow from the interaction between automation and the assumption of an unrelated labor market wedge that implies equilibrium employment is suboptimally low. Key here is that there is not a similar capital market wedge. Of course, there is a large body of research on capital market wedges, deriving from information asymmetries, agency problems, imperfect contractibility, and so on. An influential line of *Brookings Papers* going back to Fazzari, Hubbard, and Petersen (1988), along with the entire field of corporate finance, seems to invite us to consider relaxing the perfect capital markets assumption. It might be worth generalizing the results to a setting with a capital market wedge and making statements a function of the relative wedges. I suspect some of the results on the interaction between capital taxes and automation policies depend on the assumption of no capital market frictions.

Third, the authors make a tremendous and laudable attempt to calibrate the model to match the recent history and current level of effective labor and capital taxes. In modeling capital taxes, they account for the multitudinous idiosyncrasies of the American tax system, including depreciation incentives, business entity taxes across all corporate forms, payout taxes, effective taxes on debt versus equity finance, and differences across structures, equipment, and software investment. I recommend that any interested reader spend time in the paper's tax appendix, which details these calculations and will be helpful to others working in this area.

In the baseline calibration, the paper calculates an effective tax rate on capital in the 2010s of approximately 10 percent. I have one concern about this calculation, which pertains to the temporary nature of bonus depreciation. The main empirical inputs into these effective rates are National Income and Product Accounts (NIPA) aggregates for corporate tax revenues relative to gross operating surplus. These measures are poorly suited to capture the dynamic effects of bonus depreciation on effective tax rates. Tax revenues fall temporarily in stimulus years, but at the mechanical expense of higher taxes in the future. I worry that the current calculation of effective rates does not fully account for this dynamic. Nevertheless, it is relatively uncontroversial to claim that the tax burden on capital has fallen over time, which is the more important message of the paper.

My overall takeaway on these capital market assumptions is that defensible changes to them would quantitatively alter the authors' results. At the same time, the qualitative results are on firmer ground. Accounting for automation in considering tax policy changes could well be an important consideration going forward.

THE ROLE OF BONUS DEPRECIATION When I was on the "job market" presenting the results from Zwick and Mahon (2017), I had to work pretty hard to convince people that bonus depreciation had any effect on anything. Times have changed. The logic in the current paper implies that, not only does bonus depreciation matter for capital accumulation, it also materially biases the factor mix of production away from labor and toward automation.

How much does bonus depreciation distort investment and disfavor labor? Bonus depreciation accelerates the timing of deductions but does not change their amount. Thus, in contrast to changes in the corporate tax rate or an investment tax credit, its value to firms is driven only by discounting.[1] When interest rates are low, as they have been for the last decade, the effective subsidy is relatively small.[2] Accordingly, the aggregate effective tax rate shouldn't be very sensitive to bonus depreciation incentives. As a result, we likely need a very large substitution elasticity between labor and automation for bonus depreciation to be quantitatively relevant for aggregate automation trends.

This logic is also why the Joint Committee on Taxation (JCT) and the Treasury Department do not score bonus depreciation incentives as being very expensive in the ten-year budget window.[3] As noted above, this logic is also somewhat at odds with the authors' current approach to measuring the impact of bonus depreciation on effective tax rates without accounting for future tax payments.

---

[1] My reading of the 1980s tax history is that the investment tax credit in the 1981–1986 period, when combined with the accelerated cost recovery system (ACRS) and a more generous treatment of passive losses, was more generous to capital than recent changes to accelerated depreciation for equipment expenditures. Relative to this period, I am not sure that current capital taxes are significantly more favorable.

[2] Zwick and Mahon (2017) argue this may not be the case for firms facing financial frictions.

[3] For example, in the JCT score for the Tax Cuts and Jobs Act, the ten-year revenue estimate for expanded depreciation incentives shows an increase in revenue in later years, even before the policy sunsets. Over the ten-year window, the cost of depreciation changes is less than 10 percent of the cost of the corporate tax rate cut. See, for example, JCX-67-17, "Estimated Budget Effects of the Conference Agreement for H.R.1."

It would be terrific to see additional empirical work evaluating the effects of bonus depreciation on labor markets. To date, we have seen work documenting that investment increases overall, and more so for firms valuing liquidity and immediate benefits (House and Shapiro 2008; Zwick and Mahon 2017). One labor market effect of greater investment demand is through output effects among capital suppliers, where we would expect to see *higher* employment.

Zwick and Mahon (2017) also present evidence that the wage bill increases at the firm level among firms buying more equipment due to bonus. This result has been confirmed at the local labor market level by Garrett, Ohrn, and Suárez Serrato (2020) and Ohrn (2019). These two papers also show that employment either remains unchanged or increases. In a fascinating recent study, Tuzel and Zhang (2019) find that, among firms buying more equipment in response to depreciation incentives, skilled labor increases while unskilled labor falls.

Taken together, the existing evidence points toward potential complementarity between labor and capital demand induced by bonus depreciation. In the terminology of Acemoglu and Restrepo (2018), it is possible the productivity effect dominates the displacement effect. More work is needed to establish the robustness of these results and to investigate the extent to which such investment incentives promote automation.

A TASK-SPECIFIC TAX IN PRACTICE? The authors highlight the theoretical appeal of an *automation tax*, described as “an additional tax on the use of capital in tasks where labor has a comparative advantage.” The intuition for this result is clear: such a tax has a first-order benefit in increasing labor demand away from its inefficiently low level and only a second-order cost because so-so automation is only so-so. But what are the practical policy implications?

Perhaps I have read too many papers by Joel Slemrod, but I couldn’t help imagining the various strategies that firms and their consultants could devise to avoid such a tax. The literature on automation has attempted to identify those occupations that face automation risk, but we are very far from being able to codify such a system into policy. Were we to take on such a challenge, we would surely face the implications of what I call the *Slemrod conjecture* (Slemrod and Kopczuk 2002): tax avoidance is weakly increasing in the number of tax instruments.

This risk has been realized in the case of bonus depreciation, arguably a much easier policy to implement and enforce, as it builds upon preexisting rules. Subsequent to bonus depreciation’s enactment, a new consulting service called “cost segregation” has grown in popularity. These consultants

Figure 1. Cumulative Labor Share Decline with and without Manufacturing

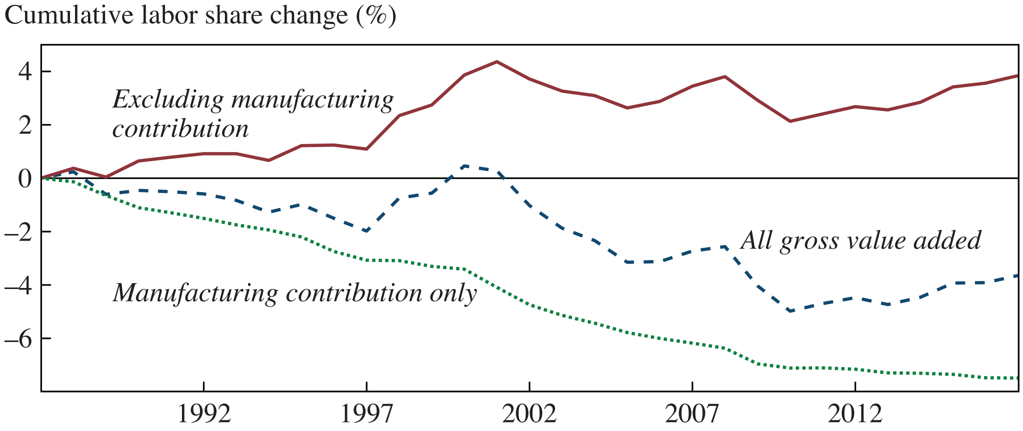

Source: Smith and others (2019, fig. 5); data from US Bureau of Economic Analysis (BEA).

advise companies on how to adjust their accounting to relabel ineligible investment expenditures (for example, fixed internal features of new buildings) as shorter-lived expenditures to be depreciated under bonus. This industry has even spawned an organization of cost segregation experts, the American Society of Cost Segregation Professionals, who have developed standards, educational material, and even a code of ethics. While we can debate the likely employment effects of an automation tax for workers, I am more confident (and also concerned) that such a tax will help promote the full employment of accountants.

REVISITING THE LABOR SHARE'S DECLINE My final comment concerns the authors' broader motivating question, which has animated research in macroeconomics over the last five to ten years: What is driving the decline in the labor share? A more specific version of this question concerns the role of tax policy in the labor share's decline.

As a starting point, let us remind ourselves that manufacturing is the most quantitatively important sector of the economy for understanding the labor share's trend since the 1980s. This fact is reasonably well known but sometimes underemphasized (figure 1).[4] Less well appreciated is that

---

[4] For example, both Karabarbounis and Neiman (2014) and Autor and others (2020) emphasize the broad-based nature of declines in the labor share. See Charles, Hurst, and Schwartz (2019) for a recent survey with new facts.

this secular decline in the manufacturing sector's contribution to the labor share was offset—fully until 2000 and partly since then—by a rise in the contribution from services. To see this, it is important to recognize evolving tax incentives to characterize owner-manager payments as labor versus profits in the skilled service sector (Smith and others 2019).

These facts are useful for the automation story, because manufacturing is one of the sectors most exposed to the rise of robots and other process automation (Acemoglu and Restrepo 2020). Perhaps automation induced by recent changes in tax policy is an important driver of labor share declines, especially in manufacturing?

Figure 2 takes a closer look at the trends in employment within manufacturing over the two time periods of study in Acemoglu and Restrepo (2020), where we have sorted the industries in each time period from lowest to highest in their exposure to automation.[5]

In the 1993–2000 period, only textile manufacturing and aerospace manufacturing contribute to employment declines, with textiles accounting for most of the decline. In the 2000–2007 period, the decline in manufacturing employment is broad-based and especially large in electronics but also continues for textiles. Notably, the within-manufacturing correlation between employment declines and automation exposure is present in the later period, consistent with Acemoglu and Restrepo (2020), but not especially pronounced.

In contrast, the evidence on the aggregate role of trade exposure and offshoring for the decline in employment and the labor share is well established and quite strong (Elsby, Hobijn, and Şahin 2013; Autor, Dorn, and Hanson 2013; Pierce and Schott 2016). An open question concerns whether tax incentives amplify these forces. Federal tax policy up to and including the most recent round of tax reforms features strong incentives to locate both profits and real activity offshore. One could argue that bonus depreciation, by prioritizing capital expenditures within the United States, leans against these incentives. I would love to see more research in this area.

---

[5] While it would be ideal to look at contributions to the aggregate labor share directly, mapping industries to the broader US Bureau of Economic Analysis (BEA) sectors is nontrivial. Employment declines likely offer a useful and policy-relevant view into understanding labor share trends.

Figure 2. Employment Declines Sorted by Automation Exposure

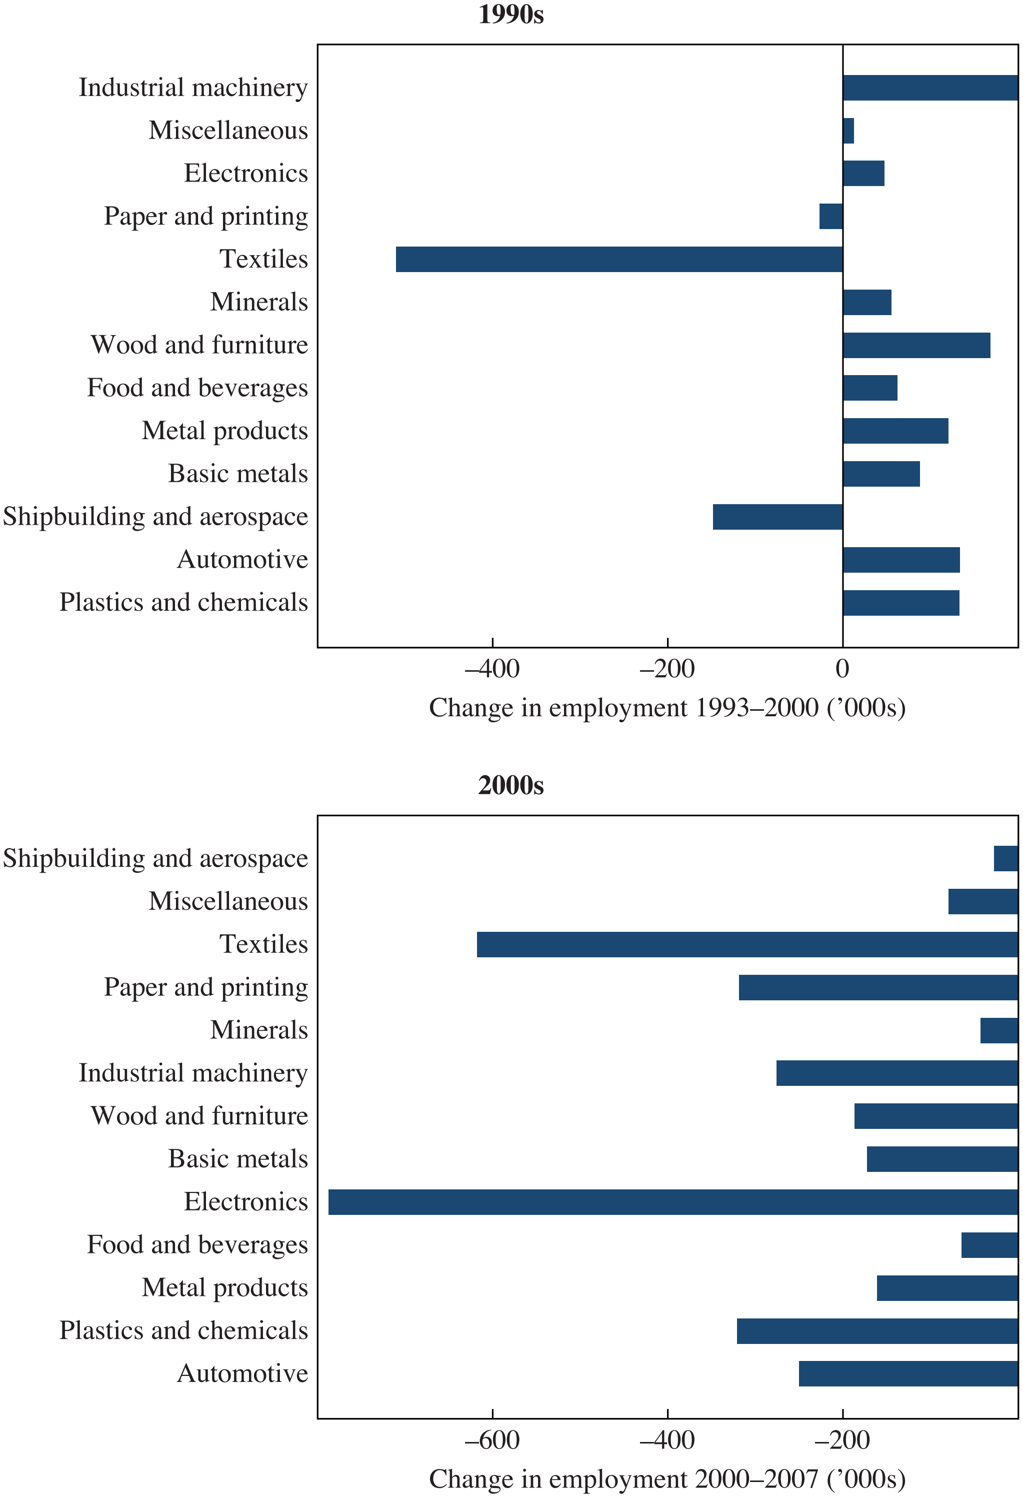

Source: Author’s calculations combining manufacturing employment statistics from the Census with automation exposure from Acemoglu and Restrepo (2020).

Note: The industries are sorted from lowest to highest in terms of automation exposure. This measure differs across the early and later periods, hence the difference in industry ordering.

REFERENCES FOR THE ZWICK COMMENT

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, no. 6: 1488–542.

Acemoglu, Daron, and Pascual Restrepo. 2020. “Robots and Jobs: Evidence from US Labor Markets.” *Journal of Political Economy* 128, no. 6: 2188–244.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. “The China Syndrome: Local Labor Market Effects of Import Competition in the United States.” *American Economic Review* 103, no. 6: 2121–68.

Autor, David H., David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2020. “The Fall of the Labor Share and the Rise of Superstar Firms.” *Quarterly Journal of Economics* 135, no. 2: 645–709.

Charles, Kerwin Kofi, Erik Hurst, and Mariel Schwartz. 2019. “The Transformation of Manufacturing and the Decline in US Employment.” *NBER Macroeconomics Annual* 33:307–72.

Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin. 2013. “The Decline of the U.S. Labor Share.” *Brookings Papers on Economic Activity*, Fall, 1–63.

Fazzari, Steven M., R. Glenn Hubbard, and Bruce C. Petersen. 1988. “Financing Constraints and Corporate Investment.” *Brookings Papers on Economic Activity*, no. 1: 141–206.

Feldstein, Martin. 1974. “Incidence of a Capital Income Tax in a Growing Economy with Variable Savings Rates.” *Review of Economic Studies* 41, no. 4: 505–13.

Garrett, Daniel G., Eric Ohrn, and Juan Carlos Suárez Serrato. 2020. “Tax Policy and Local Labor Market Behavior.” *American Economic Review: Insights* 2, no. 1: 83–100.

Harberger, Arnold C. 1962. “The Incidence of the Corporation Income Tax.” *Journal of Political Economy* 70, no. 3: 215–40.

House, Christopher L., and Matthew D. Shapiro. 2008. “Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation.” *American Economic Review* 98, no. 3: 737–68.

Jakobsen, Katrine, Kristian Jakobsen, Henrik Kleven, and Gabriel Zucman. 2020. “Wealth Taxation and Wealth Accumulation: Theory and Evidence from Denmark.” *Quarterly Journal of Economics* 135, no. 1: 329–88.

Karabarbounis, Loukas, and Brent Neiman. 2014. “The Global Decline of the Labor Share.” *Quarterly Journal of Economics* 129, no. 1: 61–103.

Ohrn, Eric. 2019. “The Effect of Tax Incentives on US Manufacturing: Evidence from State Accelerated Depreciation Policies.” *Journal of Public Economics* 180:1–14.

Pierce, Justin R., and Peter K. Schott. 2016. “The Surprisingly Swift Decline of US Manufacturing Employment.” *American Economic Review* 106, no. 7: 1632–62.

Slemrod, Joel, and Wojciech Kopczuk. 2002. “The Optimal Elasticity of Taxable Income.” *Journal of Public Economics* 84, no. 1: 91–112.

Smith, Matthew, Danny Yagan, Owen Zidar, and Eric Zwick. 2019. "The Rise of Pass-Throughs and the Decline of the Labor Share." Working Paper. http://ericzwick.com/labshare/labshare.pdf.

Tuzel, Selale, and Miao Ben Zhang. 2019. "Economic Stimulus at the Expense of Routine-Task Jobs." Working Paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3072733.

Zwick, Eric, and James Mahon. 2017. "Tax Policy and Heterogeneous Investment Behavior." *American Economic Review* 107, no. 1: 217–48.

GENERAL DISCUSSION Olivier Blanchard was disappointed that there was not more direct evidence on the effect of relative taxes on automation. He expected that there would be differences in the relative tax rates across countries, sectors, and distinct types of investments, which would lead countries or sectors to choose different technologies. Blanchard also wondered whether the reduced form labor supply specification used in the paper accounted for the fact that some low-skilled workers are going to be unemployable, because of the minimum wage effect. In addition to low-skilled workers, some skilled workers are going to be replaced by machines, and if their skills are not transferrable, there will be a large loss in income.

Robert Hall found the observations made in the paper striking and carefully thought out. He noted that the conclusions made by the authors stemmed from the parameters chosen for the capital supply elasticity. In the early 1970s, Robert Barro had a convincing argument that the capital supply elasticity was infinite.[1] In that case, Hall suggested that the optimal capital tax rate would be around 10 percent compared to the 27 percent optimal capital tax suggested by the authors. Hall also expressed related concerns about the authors' decision to use 0.65 as the capital supply elasticity. He thought that a capital supply elasticity that is less than one is extremely low and should not be taken as a fact. This is because capital in the United States and other wealthy economies is supplied by wealthy individuals, thus the behavior of the typical consumer is not relevant to the question of the capital supply because they are typically not marginal participants in the capital market. He recognized that there are obstacles to measuring capital supply elasticity but suggested that accounting for extreme inequality and uneven distribution of wealth will have a big effect on correctly measuring the capital supply elasticity.

---

1. Robert J. Barro, "Are Government Bonds Net Wealth?," *Journal of Political Economy* 82, no. 6 (1974), 1095–117.

James Poterba observed that it would be helpful to describe how far the United States tax system deviated from one that would deliver an undistorted “production efficiency” outcome, a benchmark raised in work by Diamond and Mirrlees.[2] He pointed out that there are other potential externalities that the authors might want to consider. For example, some argue that equipment should receive more favorable tax treatment than structures because there are some positive externalities associated with equipment investments, such as the acceleration of technology deployment. Also, low tax rates on software investments might in part serve as a response to imperfections in the intellectual property market, such as limitations on patenting software. Poterba did not have a way to attach numbers to such considerations or to other capital market imperfections, but he believed that quantitative analysis of these factors would result in a much richer discussion of the tax burdens on capital and labor.

Daron Acemoglu agreed that the elasticity of capital supply is an important parameter and that there was a great deal of uncertainty around it. However, he stated that capital may be very elastic when we focus on firms, but the aggregate supply coming from the consumer side may not be as elastic. Regarding Robert Hall’s comment, Acemoglu noted that infinite elasticity of capital is a natural benchmark that economists gravitate toward, but models that assume infinite elasticity of capital supply are extremely special and do not reflect the reality. Further, evidence does not support an infinite elasticity of capital. Acemoglu stated that the best papers, from his point of view, found that the elasticity of capital supply was around 0.7, but this is a noisy area in the literature. He agreed with James Poterba that this is an area where further discussion is needed.

Acemoglu further argued that a high or infinite supply elasticity would not have an impact on employment. Thus, the large effects on employment shown in the paper were a result of the supply elasticity parameters that Acemoglu and his coauthors chose. He noted that evidence suggests that capital is not as elastic as economists normally presume. Citing the work of Diamond and Mirrlees mentioned earlier, Acemoglu argued that although the theorem suggested was correct, it only worked under extremely specific assumptions. The theorems do not apply in the absence of the extremely specific assumptions. Acemoglu agreed that the practicality of an automation tax needs to be rethought. However, he suggested that implementing the Ramsey taxes would eliminate the need

---

[2] Peter A. Diamond and James A. Mirrlees, “Optimal Taxation and Public Production I: Production Efficiency,” *American Economic Review* 61, no. 1 (1971), 8–27.

for automation taxes. If capital is excessively subsidized as shown in the paper, and if implementing an automation tax is impractical, then reducing the excessive subsidy to capital and reducing payroll taxes would be welfare enhancing.

Pascual Restrepo clarified that they had conducted several robustness checks with different values of capital elasticity. The results on welfare and employment gain are robust, even for higher values of capital elasticity.