

Loebbing, Jonas

Working Paper

Redistributive Income Taxation with Directed Technical Change

CESifo Working Paper, No. 8743

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Loebbing, Jonas (2020) : Redistributive Income Taxation with Directed Technical Change, CESifo Working Paper, No. 8743, Center for Economic Studies and Ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/229561>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Redistributive Income Taxation with Directed Technical Change

Jonas Loebbing

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Redistributive Income Taxation with Directed Technical Change

Abstract

What are the implications of (endogenous) directed technical change for the design of redistributive income taxes? I study this question in a Mirrleesian economy augmented to include endogenous technology development and adoption choices by firms. Under certain conditions, any progressive tax reform induces technical change that compresses the pre-tax wage distribution. The key intuition is that progressive tax reforms tend to increase labor supply of less skilled relative to more skilled workers, which induces firms to develop and use technologies that are more complementary to the less skilled. These directed technical change effects make the optimal tax scheme more progressive, raising marginal tax rates at the right tail of the income distribution and lowering them at the left tail. For reasonable calibrations, the impact of directed technical change on the optimal tax is quantitatively important: optimal marginal tax rates are reduced substantially for incomes below the median and increase monotonically over the bulk of the income distribution instead of being U-shaped (as in most of the previous literature).

JEL-Codes: H210, H230, H240, J310, O330.

Keywords: optimal taxation, directed technical change, endogenous technical change, wage inequality.

Jonas Loebbing
Center for Economic Studies
LMU Munich / Germany
jonas.loebbing@econ.lmu.de
Website: <https://sites.google.com/view/jonasloebbing>

November 2020

I am indebted to Felix Bierbrauer, Peter Funk, and Dominik Sachs for detailed comments and discussions. Moreover, I thank Daron Acemoglu, Spencer Bastani, Christian Bredemeier, Pavel Brendler, Sebastian Findeisen, Raphael Flore, Clemens Fuest, Tobias Foell, Emanuel Hansen, Christian Holzner, Bas Jacobs, Eckhard Janeba, Sebastian Koehne, Keith Kuester, Claus Thustrup Kreiner, Etienne Lehmann, Benjamin Lockwood, Pascal Michaillat, Jakob Miethe, Andreas Peichl, Johannes Pfeifer, Paul Schempp, Florian Scheuer, Stefanie Stantcheva, Michèle Tertilt, Uwe Thue mmel, Christian Traxler, Hitoshi Tsujiyama, Daniel Waldenström, Nicolas Werquin, Han Ye, and participants at various presentations for helpful comments. This research was partly funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy–EXC 2126/1–390838866 and by the Joachim Herz Stiftung.

1 Introduction

Technical change is widely considered an important cause of the rise in wage inequality observed in many developed economies over the last decades (e.g. [Goldin and Katz, 2008](#); [Acemoglu and Autor, 2011](#)). Technologies, however, are developed and adopted by firms pursuing economic objectives. Thus, the extent to which technical change is skill biased and therefore raises wage inequality is endogenous and potentially responds to changes in economic policy ([Acemoglu, 1998, 2007](#)). This raises the question: how should redistributive policies account for the endogeneity of technology?

In this paper, I consider the design of non-linear labor income taxes as one of the major tools of redistributive policy in practice. I ask: how should we design non-linear labor income taxes when technology is determined endogenously through the profit-maximizing decisions of firms? Merging [Mirrlees \(1971\)](#) with [Acemoglu \(2007\)](#), I develop a model with directed technical change and endogenous labor supply choices that is sufficiently tractable for the analysis of arbitrarily non-linear income taxes. Within this model, I show that under certain conditions, a progressive tax reform – a reform that raises marginal tax rates on high incomes and lowers them on low incomes – can steer technology in a more equitable direction. In particular, a progressive tax reform stimulates the supply of low-skilled relative to high-skilled labor. This incentivizes firms to direct their investment towards technologies that are relatively complementary to low-skilled workers, which in turn raises the relative wages of the low-skilled. In consequence, the progressive tax reform induces technical change that reduces pre-tax wage inequality.

I show that this effect makes the optimal non-linear income tax more progressive, raising optimal marginal tax rates at the top and lowering them at the bottom of the income distribution. Intuitively, a more progressive tax not only reduces inequality in post-tax wages but also exerts a negative effect on pre-tax wage inequality through the induced technical change. The optimal tax capitalizes on this (p)redistributive effect.

Calibrations based on estimates from the empirical literature on directed technical change imply that these effects can be large. When including directed technical change, optimal marginal tax rates fall by up to 17 percentage points for incomes below the median and rise by up to 8 percentage points for high incomes.

My results advance the understanding of the implications of technical change for the design of redistributive income taxes. Previous work on optimal income taxation and technical change has mostly treated technical change as exogenous (e.g. [Ales, Kurnaz and Sleet, 2015](#); [Jacobs and Thuemmel, 2018b](#)).¹ My work adds the directed technical change perspective, whereby technical change is endogenous and its direction shaped by the structure of labor supply. Previous work on directed technical change, in turn, has not considered its implications for the design of redistributive policy. I add this normative perspective here.

Moreover, my analysis speaks to the question whether and how to regulate skill-biased technologies – a main driver of growing wage inequality. Previous work has focused on direct taxes on specific

¹Exceptions are [Jones \(2019\)](#) and [Jagadeesan \(2019\)](#), who focus on the speed, not the direction of technical change; see Section 2.

technologies, such as industrial robots (e.g. [Thuemmel, 2018](#); [Guerreiro, Rebelo and Teles, 2018](#); [Costinot and Werning, 2020](#)). Such direct taxes are optimal in theory ([Naito, 1999](#)), but it is unlikely that they can be targeted perfectly to all kinds of skill-biased technologies in practice (see Section 2 for further discussion). Hence, I focus on the (presumably large) set of technologies for which a separate direct tax is not available and show how these technologies can be regulated indirectly via the income tax.²

My first contribution is to develop a model that provides a detailed micro foundation of directed technical change while remaining sufficiently tractable for the analysis of arbitrarily non-linear taxes. The advantage of having a full-fledged micro foundation of technology choices is that it allows me to state precisely the restrictions on government policy that enter my tax analysis. In particular, besides the standard assumption that income taxes cannot be conditioned on worker skills, I assume that direct taxes on technologies cannot be conditioned on the skill bias of the respective technology, as discussed above. This makes it optimal to use the income tax to re-direct technical change.

Despite its detailed micro foundation, the model's equilibrium can be characterized by a parsimonious set of equations. Thereby, the non-linear tax analysis becomes tractable. Importantly, the equations characterizing wages and technology are identical to those studied extensively by the theory of directed technical change ([Acemoglu, 2007, 2010](#); [Loebbing, 2020](#)). Thus, existing results on directed technical change are immediately applicable to my framework. Moreover, work on directed technical change has shown that these equilibrium equations arise from a multitude of different micro foundations of directed technical change ([Acemoglu, 2007](#)). This makes my tax analysis generic within the theory of directed technical change.

In the first step, I use the model to study the effect of tax reforms on the wage distribution. Based on the techniques developed by [Sachs, Tsyvinski and Werquin \(2020\)](#), I derive a closed-form solution for the impact of arbitrarily non-linear tax reforms on wages. Using this closed-form solution, I show that progressive tax reforms tend to induce technical change that reduces skill premia and thereby lowers inequality in pre-tax wages. This effect is based on the simple logic of complementarity: a progressive reform increases the relative supply of low-skilled workers, which incentivizes investment into technologies complementary to the low-skilled.

While directed technical change effects tend to lower wage inequality in response to a progressive tax reform, they are confounded by the effects from imperfect substitutability of different worker types, as analyzed by [Sachs et al. \(2020\)](#). In particular, when holding technology constant, an increase in the labor supply of a given worker type reduces this type's wage but raises the wages of other workers. I find that the effects of these within-technology complementarities tend to counteract directed technical change: with additional structure on the aggregate production function, they raise skill premia in response to any progressive tax reform, whereas directed technical change lowers skill premia. Yet, under conditions that are supported by existing empirical evidence (e.g. [Carneiro, Liu and Salvanes, 2019](#), see Section 7.1 below), directed technical change effects dominate and skill premia decline in response to progressive tax reforms.

Next, I analyze the implications of these effects for the shape of the optimal tax schedule. I provide a

²[Jones \(2019\)](#) and [Jagadeesan \(2019\)](#) follow a similar path, analyzing optimal income taxation when the speed, but not the direction, of technical change is endogenous and cannot be controlled directly via R&D subsidies.

transparent formula for optimal marginal tax rates, extending the formula of [Sachs et al. \(2020\)](#) by the effects from directed technical change. The formula shows that directed technical change provides a force for higher marginal tax rates at the top and lower marginal tax rates at the bottom of the income distribution. To obtain precise analytical results, I compare the optimal tax with directed technical change to an appropriately specified benchmark with exogenous technology.³ I show that, indeed, optimal marginal tax rates are higher at the upper end of the income distribution and lower at the lower end when accounting for directed technical change. The key force behind this result is that a more progressive tax – besides redistributing post-tax incomes – can achieve predistribution, that is, a redistribution of pre-tax incomes, via directed technical change.

The benchmark case without directed technical change still accounts for the effects of within-technology complementarity between worker types (as in [Sachs et al., 2020](#), [Rothschild and Scheuer, 2013](#), and [Stiglitz, 1982](#)). While this is the natural benchmark to cleanly identify the impact of directed technical change, another important reference point is the case where wages are completely exogenous (as, e.g., in [Mirrlees, 1971](#), [Diamond, 1998](#), and [Saez, 2001](#)). Comparing the optimal tax with directed technical change to the benchmark with exogenous wages, the results depend on whether directed technical change effects dominate the effects from within-technology complementarity or not. If directed technical change dominates, the optimal tax is even more progressive than with exogenous wages. Hence, directed technical change not only mitigates the effects from within-technology complementarity on the optimal tax, but potentially even reverses them.

Are these effects quantitatively relevant? To answer this question, I calibrate the model based on estimates from the empirical literature on directed technical change. I first use the calibration to assess the explanatory power of my theory for the evolution of tax progressivity and wage inequality in the US. In particular, the US tax and transfer system underwent a series of regressive reforms between 1970 and 2005 (e.g. [Piketty and Saez, 2007](#)). At the same time, US wage inequality increased strongly, an evolution often attributed to skill-biased technical change. My theory suggests that skill-biased technical change may have been, to some extent, caused by the regressive tax reforms. I thus ask whether, through lens of my model, the regressive reforms of the US tax and transfer system played a significant role for the concurrent rise in US wage inequality. Using an approximation to the cumulative reforms of the US tax and transfer system between 1970 and 2005, I find that the reforms can account for an increase of up to 2.8% in the 90-10-percentile ratio of the wage distribution. Importantly, in the absence of directed technical change, the reforms would have caused a decline in the 90-10 ratio of about 2%, due to the effects from imperfect worker substitutability. Putting these results into perspective, the increase of 2.8%, potentially caused by tax reforms and the resulting technical change, amounts to a share of 9% of the total increase in the 90-10-percentile ratio of US wages between 1970 and 2005.

These effect sizes translate into a large impact of directed technical change on the optimal tax scheme. Relative to the benchmark with exogenous technology – the case studied by [Sachs et al. \(2020\)](#) – optimal marginal tax rates are reduced by up to 17 percentage points on incomes below the median and increased by up to 8 percentage points on high incomes. Relative to the benchmark with exogenous wages

³The benchmark includes the self-confirming policy equilibrium ([Rothschild and Scheuer, 2013](#)) for a planner who ignores directed technical change.

– the case studied by [Mirrlees \(1971\)](#), [Saez \(2001\)](#), and many others – the results are more nuanced. The empirical literature provides a range of estimates of directed technical change effects. For values at the lower end of this range, directed technical change and the effects from within-technology complementarity offset each other approximately. In this case, the optimal tax with directed technical change is close to the one obtained with exogenous wages. For values at the higher end of the range, however, directed technical change strongly dominates; consequently, the optimal tax with directed technical change is substantially more progressive than with exogenous wages. Remarkably, the U-shape of optimal marginal tax rates – a prominent result with exogenous wages ([Diamond, 1998](#)) – largely vanishes in this case. Instead of being U-shaped, optimal marginal tax rates increase monotonically over the bulk of the income distribution.

The paper is structured as follows. The next section, Section 2, discusses the relation to previous work. Section 3 introduces the model. Section 4 presents the results from the theory of directed technical change that form a building block of the subsequent tax analysis. The analysis of tax reforms and their effects on the wage distribution is provided in Section 5, while Section 6 contains the optimal tax analysis. Finally, Section 7 presents the quantitative results and Section 8 concludes.

2 Related Literature

My analysis connects the literature on directed technical change with the literature on the optimal design of non-linear labor income taxes.

In the literature on optimal taxation, [Ales et al. \(2015\)](#) and [Jacobs and Thuemmel \(2018b\)](#) analyze how labor income taxes respond optimally to skill-biased (or routine-biased) technical change (see also [Jacobs and Thuemmel, 2018a](#), and [Heathcote, Storesletten and Violante, 2020](#)). They treat technical change as an exogenous change in the production technology. I contribute to the research program on technical change and tax design by treating the skill bias of technical change as endogenous, such that it responds to the tax system.

[Jones \(2019\)](#) and [Jagadeesan \(2019\)](#) also study the design of optimal income taxes when technical change is endogenous, finding that optimal labor income taxes are reduced relative to a setting with exogenous technology. In their models, the speed of technical change responds to tax policy, but technical change is always unbiased, leaving relative wages between worker types unaffected. In my model, in contrast, the skill bias of technical change is endogenous and responds to tax policy while the speed of technical change is kept constant. My contribution is thus complementary to theirs.⁴

My analysis builds on [Sachs et al. \(2020\)](#) and [Stiglitz \(1982\)](#), who analyze optimal taxation (and tax incidence) in models with an aggregate production function that features complementarity between different types of workers. They find that accounting for such complementarity reduces optimal marginal tax rates at the top and increases them at the bottom of the income distribution. I extend the analysis to incorporate directed technical change. Directed technical change counteracts the effects from within-

⁴[Acemoglu, Manera and Restrepo \(2020\)](#) study the optimal taxation of labor and capital in a setting with endogenous automation technology building on [Acemoglu and Restrepo \(2018\)](#). Since they consider a Ramsey setting with a representative agent, they do not analyze inequality and redistribution between different types of labor.

technology complementarity between workers, raising marginal tax rates at the top and lowering them at the bottom. Depending on the strength of the effects, directed technical change either mitigates or even reverses the impact of within-technology complementarity on the optimal tax.

In a conceptually similar contribution, [Rothschild and Scheuer \(2013\)](#) extend [Stiglitz \(1982\)](#) to incorporate endogenous occupation choices of workers. They find that occupational switching mitigates, but never overcompensates complementarity between occupations. Hence, with endogenous occupation choices, optimal marginal tax rates are bounded between those with exogenous wages (on the progressive end) and the [Stiglitz \(1982\)](#) case with complementarities but without occupational switching (on the regressive end). With directed technical change, in contrast, overcompensation of within-technology complementarity is possible and the optimal tax can be more progressive than with exogenous wages.

Finally, my analysis is related to studies of optimal technology taxes (e.g. [Thuemmel, 2018](#); [Guerreiro et al., 2018](#); [Costinot and Werning, 2020](#)). These studies assume that direct taxes can be targeted perfectly to specific technologies. In contrast, I consider the case where direct taxes on specific technologies are infeasible (as do [Jones, 2019](#), and [Jagadeesan, 2019](#)). While direct taxes on skill-biased technologies are optimal in theory by the results of [Naito \(1999\)](#), they are challenging to implement in practice. In particular, they require precise information about the labor market impact of a well-defined technology and its applications. Moreover, the government must be able to monitor which technology is used in which ways by any individual firm. While these requirements may be met for industrial robots – which are studied by [Thuemmel \(2018\)](#), [Guerreiro et al. \(2018\)](#), and [Costinot and Werning \(2020\)](#) – they are unlikely to be satisfied in more than a few well-studied cases.⁵ Whenever a relevant share of skill-biased technologies cannot be targeted perfectly by direct taxes, however, there is scope for exploiting directed technical change effects via the income tax. Thus, my approach is complementary to work on direct technology taxes in that it applies to technologies for which a specific direct tax is not available.⁶

Starting from the theory of directed technical change, I build on the seminal ideas of [Acemoglu \(1998\)](#) and [Kiley \(1999\)](#) and explore their normative implications, in particular for the design of redistributive labor income taxes. In doing so, I use the theoretical advances by [Acemoglu \(2007\)](#) and [Loebbing \(2020\)](#) as a building block in my analysis. Specifically, their results lend structure to the relationship between labor supply and production technology, which I exploit to analyze the relationship between taxes and technology.

I use empirical work on directed technical change to quantify my results. In particular, I use estimates from [Lewis \(2011\)](#), [Dustmann and Glitz \(2015\)](#), [Morrow and Treffer \(2017\)](#), and [Carneiro et al. \(2019\)](#) to calibrate the strength of directed technical change effects in my model. The empirical literature on directed technical change is discussed in more detail in Section 7.⁷

⁵[Slavik and Yazici \(2014\)](#) consider the differential taxation of structure and equipment capital, based on the assumption that equipment capital is skill-biased. In that case, the targeted technology is not narrow but the targeting is imprecise: there are likely many different forms of equipment capital, which are skill biased to different degrees. This again gives a role for indirect targeting through the income tax, even if the optimal tax differential between equipment and structures is in place.

⁶Another important contribution to the analysis of taxes and technology is provided by [Akcigit, Hanley and Stantcheva \(2019b\)](#), who design optimal R&D subsidies for heterogeneous firms. Their approach is not directly related to wage inequality, however.

⁷[Akcigit, Grigsby, Nicholas and Stantcheva \(2019a\)](#) show empirically that innovation responds to taxes (see also [Akcigit and](#)

3 A Model of Directed Technical Change with Non-Linear Income Taxes

I merge a standard [Mirrlees \(1971\)](#) model of optimal income taxation with a directed technical change model in the spirit of [Acemoglu \(2007\)](#).⁸ The model features a continuum of differentially skilled workers, as in [Mirrlees \(1971\)](#), and a production sector that extends the monopolistic competition setup of [Romer \(1990\)](#) to include multiple types of technologies. Crucially, technologies differ in their complementarity relationships with different types of workers; some technologies are biased towards more skilled, others towards less skilled workers. Investment into these technologies is endogenous and responds, via the structure of labor supply, to the income tax. These technology responses and their implications for the optimal income tax are the central object of interest in the paper.

3.1 Workers

Workers are modeled in a standard way. There is a continuum of workers with different types $\theta \in \Theta = [\underline{\theta}, \bar{\theta}] \subset \mathbb{R}$. Types are distributed according to the strictly positive density function $h : \theta \mapsto h_\theta$, with cumulative distribution function H .

Workers' utilities depend on consumption c_θ and labor supply l_θ according to

$$u_\theta = c_\theta - v(l_\theta) ,$$

where v represents disutility from labor. The disutility from labor is twice continuously differentiable with $v' > 0$ and $v'' > 0$ everywhere.

Wages are denoted by w_θ , such that workers' pre-tax incomes are $y_\theta = w_\theta l_\theta$. Income taxes are given by the tax function $T : y_\theta \mapsto T(y_\theta)$ and the retention function corresponding to tax T is denoted by R_T . Hence, workers' budget constraints are

$$c_\theta = R_T(w_\theta l_\theta) + S , \tag{1}$$

where S is a lump-sum transfer the government uses to meet its budget constraint.

Workers choose their labor supply to maximize utility, taking wages as given. The first-order condition is (assuming that the tax function is differentiable):

$$v'(l_\theta) = R'_T(w_\theta l_\theta) w_\theta . \tag{2}$$

3.2 Final Good Firms

There is a continuum of mass one of identical final good firms indexed by i . They produce a final consumption good (the numéraire) according to the continuously differentiable production function

[Stantcheva, 2020](#)), but they do not study the directed technical change channel considered here.

⁸The model is set up in static terms but it is straightforward to construct a dynamic model with a balanced growth path that, once detrended, is equivalent to the static model's equilibrium (see Appendix B in [Loebbing, 2016](#)). The tax analysis in the static model is thus equivalent to a tax analysis on the balanced growth path of a corresponding dynamic model, where the brackets of the tax function grow at the rate of total output.

$G(L_i, Q_i)$. The first input $L_i = \{L_{i,\theta}\}_{\theta \in \Theta}$ collects the amounts of all different types of labor used by firm i . The second input $Q_i = \{Q_{i,j}\}_{j \in \{1,2,\dots,J\}}$ collects the variables $Q_{i,j}$, each of which is an aggregate of a continuum of technology-emboding intermediate goods:⁹

$$Q_{i,j} = \int_0^1 \phi_{j,k} q_{i,j,k}^\alpha dk .$$

The variables $q_{i,j,k}$ denote the amount of intermediate good (j,k) used by firm i , while the parameter $\alpha \in (0,1)$ governs the substitutability of intermediates of the same type j . The variables $\phi_{j,k}$ represent the quality of the corresponding intermediate goods. These quality, or productivity, levels of the different types of intermediate goods are the endogenous component of technology in the model. Their determination is described in detail below.

Along the k -dimension, intermediate inputs are completely symmetric. This dimension of heterogeneity is only used to introduce monopolistic competition and will disappear in equilibrium (due to symmetry). The j -dimension, in contrast, captures substantive differences between intermediate inputs (or, between the technologies embodied therein): inputs of different types j may have different effects on relative wages and thus different implications for wage inequality.

For $J = 1$ and $G(L_i, Q_i) = Q_{i,1} f(L_i)$ with f homogeneous of degree $1 - \alpha$, we would obtain the production sector of Romer (1990). There would be a single type of technology, which is factor-neutral and never affects relative wages. Instead, we allow for multiple types of technology, $J > 1$, and do not restrict them to be factor-neutral; that is, their embodiments $Q_{i,j}$ may change relative wages between worker types.

With this structure of final good production, we can write the output of firm i as $\tilde{G}(L_i, \phi, q_i)$ where $\phi = \{\phi_{j,k}\}_{(j,k) \in \{1,2,\dots,J\} \times [0,1]}$ and $q_i = \{q_{i,j,k}\}_{(j,k) \in \{1,2,\dots,J\} \times [0,1]}$ collect qualities and quantities of all different intermediate inputs. I assume that the function $\tilde{G}$ is homogeneous in q , such that a proportional increase in all intermediate inputs does not change relative wages. This ensures that the optimal uniform subsidy on intermediate inputs is purely Pigouvian, which allows for a clean separation of efficiency and redistributive concerns (see the description of the government below).

Moreover, let the function $\tilde{G}$ be linear homogeneous and concave in the rival inputs (L, q) , satisfying the standard microeconomic replication argument (e.g. Romer, 1994). Then, the final good sector admits a representative firm and we can drop the index i in what follows.

Final good firms' profit maximization leads to the following demand for all labor types θ :

$$w_\theta = D_{L_\theta} \tilde{G}(L, \phi, q) . \quad (3)$$

The operator D_{L_θ} denotes Gateaux differentiation with respect to L in direction of L_θ .¹⁰ Labor market clearing requires that the aggregate labor demand L_θ equals the sum of individual workers' labor supply,

$$L_\theta = l_\theta h_\theta \quad \text{for all } \theta . \quad (4)$$

⁹The case with a continuum of different intermediate good types $j, j \in [0, J]$, can be treated analogously.

¹⁰See Appendix A.1 for rigorous definitions of the infinite-dimensional derivatives used throughout the paper.

Demand for intermediate good $q_{j,k}$ is given by

$$p_{j,k} = \alpha \phi_{j,k} q_{j,k}^{\alpha-1} \frac{\partial G(L, Q)}{\partial Q_j}, \quad (5)$$

where $p_{j,k}$ is the price of the intermediate good.

3.3 Technology Firms

The technology-emboding intermediate goods are produced under monopolistic competition by technology firms. Each good (j, k) is produced by a single technology firm, which I label by the index (j, k) of its output. Technology firm (j, k) produces its output at constant marginal cost η_j from final good and receives an ad valorem sales subsidy of ξ (see the description of the government for details). It sets the post-subsidy price $p_{j,k}$ to maximize profits, given by

$$((1 + \xi)p_{j,k} - \eta_j) q_{j,k},$$

subject to the demand from final good firms (5). Since the demand from final good firms is isoelastic, the profit-maximizing price is given by a constant markup over marginal cost net of the subsidy:

$$p_{j,k} = \frac{\eta_j}{(1 + \xi)\alpha}. \quad (6)$$

Technology firms can invest R&D resources to improve the quality of their output. The resulting distribution of quality levels over the different types of intermediate inputs represents the endogenous state of technology in the model. Analyzing its implications for the design of income taxes is the central purpose of the paper.

For a given firm (j, k) , a quality level of $\phi_{j,k}$ costs $C_j(\phi_{j,k})$ units of R&D resources, where the cost function C_j is smooth, convex, and strictly increasing for every j . Firm (j, k) 's profits as a function of its quality level $\phi_{j,k}$ are

$$\pi_{j,k}(\phi_{j,k}) = \max_q \left\{ \alpha \phi_{j,k} \frac{\partial G(L, Q)}{\partial Q_j} q^\alpha - \eta_j q - p^r C_j(\phi_{j,k}) \right\},$$

where p^r denotes the competitive market price of R&D resources. Via an envelope argument, the first-order necessary condition for the choice of quality is given by

$$\alpha \frac{\partial G(L, Q)}{\partial Q_j} q_{j,k}^\alpha = p^r \frac{dC_j(\phi_{j,k})}{d\phi_{j,k}}, \quad (7)$$

where $q_{j,k}$ is assumed to take its profit-maximizing value implied by equation (6). Regarding sufficient conditions, one can verify that the optimal $q_{j,k}$ grows at the rate $1/(1 - \alpha)$ in $\phi_{j,k}$, such that the left-hand side of equation (7) grows at rate $\alpha/(1 - \alpha)$ in $\phi_{j,k}$. I assume henceforth that $dC_j/d\phi_{j,k}$ grows at a rate greater than $\alpha/(1 - \alpha)$ in $\phi_{j,k}$. This ensures that the first-order condition has a unique solution, which

identifies a maximum. Since profits are symmetric across all firms (j, k) of the same type j , uniqueness of the profit maximum implies that the choices of all firms of type j are the same and we can drop the k -index henceforth.

The supply of R&D resources is exogenous and given by $\bar{C}$. Their price adjusts to guarantee market clearing,

$$\sum_{j=1}^J C_j(\phi_j) = \bar{C}.$$

The assumption of a fixed amount of R&D resources allows to focus on the effects of labor income taxes on the direction instead of the speed of technical change.¹¹

3.4 Government

The government levies an income tax, a profit tax, and a uniform subsidy on intermediate goods that does not differentiate between intermediate good types.¹²

Since final good firms' production function $\tilde{G}$ is homogeneous in intermediate goods, a uniform tax on intermediates can only lead to proportional changes in their quantities, which in turn leave relative wages unaffected. Therefore, the intermediate good tax cannot be used to alleviate the distortions from redistributive labor taxes as in [Naito \(1999\)](#). It follows that the optimal intermediate good subsidy is purely Pigouvian and set at $\xi = (1 - \alpha)/\alpha$. This ensures that the price of intermediate goods equals marginal cost, $p_j = \eta_j$ for all j . I assume that this optimal subsidy is in place throughout the analysis.

The profit tax is levied on technology firms and the owners of R&D resources. As is standard in the literature on optimal labor income taxation, I assume that the profit tax is confiscatory to avoid a role for the distribution of firm ownership without a meaningful theory of wealth formation in the model.¹³

The income tax T is the central object of interest in the paper. Note that, without an income tax, the equilibrium allocation is efficient due to the Pigouvian intermediate good subsidy. Hence, the only motive to tax income is redistribution.¹⁴

Taken together, taxes and subsidies generate the following government revenue,

$$S(y) = \int_{\Theta} T(y_{\theta}) h_{\theta} d\theta + p^r \bar{C} + \sum_{j=1}^J \pi_j - \sum_{j=1}^J \xi p_j q_j, \quad (8)$$

which is redistributed lump-sum across workers.

¹¹For an analysis of optimal income taxes when the speed, but not the direction, of technical change is endogenous, see [Jagadeesan \(2019\)](#) and [Jones \(2019\)](#).

¹²From an informational perspective, I assume that the government neither observes individual workers' types nor the types of intermediate goods produced by individual technology firms. The former gives rise to the standard restriction that income taxes cannot be conditioned on worker types while the latter implies that the government cannot tax different intermediate goods at different rates.

¹³Note that confiscatory profit taxes are part of the optimal tax policy whenever ownership shares of firms increase and marginal welfare weights decrease in workers' income levels at the optimum.

¹⁴For an analysis of optimal income taxes with Pigouvian elements, see, for example, [Rothschild and Scheuer \(2016\)](#) and [Lockwood, Nathanson and Weyl \(2017\)](#).

3.5 Equilibrium

An equilibrium of the model, given a tax function T , is a collection of quantities and prices such that all firms maximize profits, workers maximize utility, and all markets clear.

Despite the detailed micro structure of the model, the equilibrium variables of interest for the tax analysis can be characterized by a parsimonious set of equations. To derive these equations, note first that aggregate net production at labor input l and a given set of quality levels ϕ can be written as (because intermediate good prices equal marginal cost):

$$F(l, \phi) := \max_q \left\{ \tilde{G}(\{h_\theta l_\theta\}_{\theta \in \Theta}, \phi, q) - \sum_{j=1}^J \eta_j q_j \right\}. \quad (9)$$

Note that I used labor market clearing (4) to replace the aggregate labor input L by the individual labor input l to save on notation in the following. Via an envelope argument, the labor demand equation (3) then implies that, in equilibrium, wages are given by

$$w_\theta(l, \phi) = \frac{1}{h_\theta} D_{l_\theta} F(l, \phi), \quad (10)$$

where the adjustment factor $1/h_\theta$ is necessitated by the switch from aggregate to individual labor inputs in the aggregate production function.

The condition for profit-maximizing quality choices of technology firms (7) coincides with the first-order condition for a maximum of aggregate production with respect to quality ϕ (simply called technology, henceforth) when ϕ is restricted to the set of feasible technologies

$$\Phi = \left\{ \phi \in \mathbb{R}_+^J \mid \sum_{j=1}^J C_j(\phi_j) \leq \bar{C} \right\}.$$

Thus,

$$\phi^*(l) \in \operatorname{argmax}_{\phi \in \Phi} F(l, \phi) \quad (11)$$

is an equilibrium technology. In the following I focus on equilibria in which technology satisfies (11). Existence of other equilibria can be ruled out by imposing assumptions that guarantee strict quasiconcavity of F in ϕ under the constraint $\phi \in \Phi$.

Finally, we can simplify the expression for the government's budget surplus. To this end, note that marginal cost pricing of intermediate goods implies that technology firms' profits are equal to the total amount of subsidies minus the cost for R&D resources:

$$\sum_{j=1}^J \pi_j = \sum_{j=1}^J ((1 + \xi)p_j - \eta_j) q_j - p^r \bar{C} = \sum_{j=1}^J \xi p_j q_j - p^r \bar{C}.$$

It follows that the revenue from corporate taxes and the expenses on technology good subsidies offset

each other exactly in equation (8), such that the expression for government revenue shrinks to

$$S(y) = \int_{\Theta} T(y_{\theta}) h_{\theta} d\theta . \quad (12)$$

We can now fully characterize the equilibrium values for consumption levels c , labor inputs l , wages w , technology ϕ , and government revenue S by workers' budget constraint (1), their first-order condition (2), the wage equation (10), the technology condition (11), and the equation for government revenue (12). These equilibrium conditions provide the starting point for the tax analysis in the following sections.

The wage and technology equations (10) and (11) are identical to the conditions characterizing equilibrium in the directed technical change models presented in Acemoglu (2007) and Loebbing (2020, Appendix C.2). Introducing endogenous labor supply and a government then gives rise to the remaining three equations. In this sense, my tax analysis does not depend on the details of the presented model but holds more generally within a large class of directed technical change models.

Throughout the analysis, I impose the following assumptions on the aggregate production function F and the income tax T .¹⁵

Assumption 1.

1. *The aggregate production function F is twice continuously differentiable.*
2. *The Gateaux derivative $D_{l_{\theta}} F$ is strictly positive everywhere for all θ .*
3. *The maximizer $\arg\max_{\phi \in \Phi} F(l, \phi)$ is unique for all l and Gateaux differentiable in l everywhere.*
4. *Whenever an exogenous tax T is considered, it is twice continuously differentiable and satisfies $T'(y_{\theta}) < 1$ and $P_T(y_{\theta}) \frac{v'(l_{\theta})}{v''(l_{\theta}) l_{\theta}} > -1$ for all θ in equilibrium under T .*

The first three parts of the assumption ensure that wages and technology respond smoothly to changes in labor supply, while the last part ensures that labor supply responds smoothly to changes in wages and taxes. In particular, the last part ensures that workers' second-order conditions are satisfied strictly under a given tax T , such that the elasticities of labor supply with respect to the wage and the marginal retention rate are well defined for all workers (see Appendix A.2).

3.6 Special Case: CES Production Function

An important special case of the model is obtained when final good firms' production function features a constant elasticity of substitution between worker types while technology firms' R&D cost functions

¹⁵The assumptions on aggregate production F can, of course, be derived from assumptions on final good firms' production function G . I impose them on F directly because this seems more transparent and, as argued above, the equations characterizing the relevant equilibrium variables are much more general than the specific model from which they are derived here.

are isoelastic. In particular, let final good firms' production function be given by

$$\tilde{G}(L, \tilde{\phi}, q) = \left[\int_{\underline{\theta}}^{\bar{\theta}} \left(\tilde{\kappa}_{\theta} L_{\theta}^{1-\alpha} \int_0^1 \tilde{\phi}_{\theta,k} q_{\theta,k}^{\alpha} dk \right)^{\frac{\tilde{\sigma}-1}{\sigma}} d\theta \right]^{\frac{\sigma}{\tilde{\sigma}-1}},$$

where $\tilde{\kappa}$ is a continuously differentiable function that assigns an exogenous productivity level to each type of worker; $\tilde{\sigma}$ is the elasticity of substitution between worker types in the production of an individual firm; and $\tilde{\phi}$ denotes the endogenous quality levels of the different types of intermediate inputs. Note that for each worker type θ there exists a corresponding intermediate good type, such that it is convenient to index intermediate good types directly by the worker type index θ instead of using a separate index j .

In addition, let technology firms' R&D cost functions be given by

$$C_{\theta}(\tilde{\phi}_{\theta,k}) = \tilde{\phi}_{\theta,k}^{\tilde{\delta}},$$

where $\tilde{\delta}$ controls the convexity of R&D costs.

From these functional forms, the aggregate production function F can be derived as (see Appendix A.3)

$$F(l, \phi) = \left[\int_{\underline{\theta}}^{\bar{\theta}} (\kappa_{\theta} \phi_{\theta} l_{\theta} h_{\theta})^{\frac{\sigma-1}{\sigma}} d\theta \right]^{\frac{\sigma}{\sigma-1}}, \quad (13)$$

where κ and σ are composed of the parameters $\tilde{\kappa}$, $\tilde{\sigma}$, α , and η (the unit cost of intermediate goods).¹⁶ The new technology variable ϕ is a transformation of the endogenous quality levels $\tilde{\phi}$, defined by $\phi_{\theta} := \tilde{\phi}_{\theta}^{1/(1-\alpha)}$ for all θ . Accordingly, the R&D cost for a technology level ϕ_{θ} is given by $\phi_{\theta}^{\tilde{\delta}(1-\alpha)}$. Using the substitution $\delta := \tilde{\delta}(1-\alpha)$, the set of feasible technologies Φ can thus be expressed as

$$\Phi = \left\{ \phi : \theta \mapsto \phi_{\theta} \in \mathbb{R}_+ \mid \int_{\underline{\theta}}^{\bar{\theta}} \phi_{\theta}^{\delta} d\theta \leq \bar{C} \right\}, \quad (14)$$

where δ governs the ease with which productivity of one intermediate good type can be transformed into productivity of another intermediate good type by the R&D sector.

From equation (13), wages in the CES case are obtained as

$$w_{\theta}(l, \phi) = (\kappa_{\theta} \phi_{\theta})^{\frac{\sigma-1}{\sigma}} (l_{\theta} h_{\theta})^{-\frac{1}{\sigma}} F(l, \phi)^{\frac{1}{\sigma}}. \quad (15)$$

I use the CES case, as summarized by equations (13) to (15), for the quantitative analysis in Section 7. It will also allow for some additional analytical insights into the effects of tax reform and the shape of optimal taxes in Sections 5 and 6.

¹⁶Appendix A.3 provides the exact expressions for κ and σ in terms of $\tilde{\kappa}$, $\tilde{\sigma}$, α , and η . These expressions also show that continuous differentiability of $\tilde{\kappa}$ translates into continuous differentiability of κ .

4 Directed Technical Change

Income taxes affect technical change via differential changes in labor supply across worker types. An important building block of my tax analysis is, therefore, the relationship between labor supply and technical change. This relationship is studied by the theory of directed technical change.

In the following, I briefly review the central results of the theory of directed technical change for the purpose of the tax analysis below. For that, take labor supply l as exogenous for the moment and consider wages and technology as functions of l , as given by the equilibrium conditions (10) and (11). Starting from these conditions, the theory of directed technical change asks how changes in labor supply affect relative wages via the endogenous adjustment of technology.

4.1 Weak Relative Bias

The first main result of directed technical change, called weak relative bias of technology (Acemoglu, 2007), identifies conditions under which any increase in relative skill supply induces skill-biased technical change. I have shown in previous work that quasisupermodularity of aggregate production in technology is sufficient and close to necessary for weak relative bias (Loebbing, 2020).

In particular, let a technology be skill biased relative to another if it leads to pervasively higher skill premia.

Definition 1. A technology ϕ is more skill biased than another technology $\tilde{\phi}$ if, for any labor input l , all skill premia are greater under ϕ than under $\tilde{\phi}$, that is,

$$\frac{w_\theta(l, \phi)}{w_{\tilde{\theta}}(l, \phi)} \geq \frac{w_\theta(l, \tilde{\phi})}{w_{\tilde{\theta}}(l, \tilde{\phi})}$$

for all $\theta \geq \tilde{\theta}$.¹⁷ We write $\phi \succeq^{sb} \tilde{\phi}$.

Given the skill-bias order $\succeq^{sb}$, quasisupermodularity of the aggregate production function is imposed as follows.

Assumption 2. Aggregate production F is quasisupermodular in ϕ under the skill-bias order $\succeq^{sb}$. In particular, for any labor input l and any two technologies ϕ and $\tilde{\phi}$, the following holds: if $F(l, \phi) \geq F(l, \underline{\phi})$ for all $\underline{\phi}$ with $\phi, \tilde{\phi} \succeq^{sb} \underline{\phi}$, then there must exist $\bar{\phi} \succeq^{sb} \phi, \tilde{\phi}$ with $F(l, \bar{\phi}) \geq F(l, \tilde{\phi})$.¹⁸

In words, if ϕ raises output relative to all technologies that are less skill biased than both ϕ and $\tilde{\phi}$, then there must be a technology that is more skill biased than both ϕ and $\tilde{\phi}$ and raises output relative

¹⁷Below, I will assume that types are ordered such that a higher θ implies a higher wage, either at the initial tax system (in the tax reform analysis) or at the optimal tax (in the optimal tax analysis). In this sense, a higher θ -type is a more skilled type of worker and the ratio $w_\theta/w_{\tilde{\theta}}$, for $\theta \geq \tilde{\theta}$, is indeed a skill premium.

¹⁸Note that this slightly deviates from the original definition of quasisupermodularity by Milgrom and Shannon (1994). For their definition, we would first have to assume that the set $(\Phi, \succeq^{sb})$ has a lattice structure, that is, for any two technologies ϕ and $\tilde{\phi}$ there exist supremum and infimum in Φ . Then, quasisupermodularity would be defined using infimum and supremum instead of arbitrary technologies below and above ϕ and $\tilde{\phi}$. In particular, for any l and any $\phi, \tilde{\phi}$, if $F(l, \phi) \leq F(l, \tilde{\phi})$, then $F(l, \bar{\phi}) \geq F(l, \underline{\phi})$, where $\bar{\phi}$ and $\underline{\phi}$ denote infimum and supremum of ϕ and $\tilde{\phi}$. My definition is slightly less restrictive (and sufficiently restrictive for the present purpose).

to $\tilde{\phi}$. Intuitively, this requires that changes in technology that raise skill premia in different parts of the wage distribution must not be substitutes. To see this, suppose there are two innovations, one that raises skill premia in the top half of the wage distribution and one that raises skill premia in the bottom half. Let $\underline{\phi}$ be the state of technology without either of the two innovations, ϕ be the state with the innovation that raises skill premia at the top, $\tilde{\phi}$ the state with the innovation that raises skill premia at the bottom, and $\bar{\phi}$ the state with both innovations. Then, if the innovation that raises skill premia at the top leads to an output gain absent the other innovation (i.e. the move from $\underline{\theta}$ to θ raises output), quasisupermodularity requires that it also leads to an output gain when the other innovation is already implemented (i.e. the move from $\tilde{\theta}$ to $\bar{\theta}$ raises output).

Due to the additive structure in the CES aggregators, the CES setup in (13) and (14) is exactly the case where innovations that affect different parts of the wage distribution are independent of each other: rearranging the productivity levels ϕ_θ in one part of the type space has no bearing on whether a certain rearrangement of productivity levels in another part of the type space raises output or not. Hence, the CES function satisfies quasisupermodularity.

In [Loebbing \(2020, Section 6 and Appendix C.3\)](#), I present further important production structures that satisfy quasisupermodularity. For example, in an assignment model à la [Costinot and Vogel \(2010\)](#) with capital as an additional production factor (as, e.g., in [Acemoglu and Restrepo, 2018](#)) and endogenous capital productivity, aggregate production is quasisupermodular in capital productivity. Note also that, while I am not aware of empirical tests of quasisupermodularity in the aggregate production process, the implications of quasisupermodularity presented in Lemma 1 below receive support in the empirical literature. I discuss this literature in Section 7, when quantifying the results of the tax analysis.

Taking quasisupermodularity as given, the following result applies.

Lemma 1 (cf. Theorem 2, [Loebbing, 2020](#)). *Take any labor input l and let dl be a change in the labor input such that dl_θ/l_θ increases in θ (an increase in relative skill supply). Then, the technical change induced by dl raises all skill premia, that is,*

$$\left. \frac{d}{d\mu} \frac{w_\theta(l, \phi^*(l + \mu dl))}{w_{\tilde{\theta}}(l, \phi^*(l + \mu dl))} \right|_{\mu=0} \geq 0 \quad (16)$$

for all $\theta \geq \tilde{\theta}$.

Proof. See Appendix [A.5](#). □

The intuition behind weak relative bias is that an increase in relative skill supply raises the profitability of technologies that are relatively complementary to high-skilled workers; increased investment into these technologies in turn raises the relative productivity of the high-skilled.

4.2 Strong Relative Bias

Weak relative bias is about the effects of labor supply changes on relative wages that are transmitted via directed technical change. These effects are important to identify precisely what is added by directed

technical change. Yet, in general, a labor supply change also has a direct effect on relative wages that is independent of the adjustment of technology. This within-technology substitution effect is formally given by

$$\left. \frac{d}{d\mu} \frac{w_\theta(l + \mu dl, \phi^*(l))}{w_{\tilde{\theta}}(l + \mu dl, \phi^*(l))} \right|_{\mu=0}.$$

Due to concavity of the aggregate production function in l , the within-technology substitution effect of an increase in relative skill supply on skill premia is typically negative. It therefore counteracts the directed technical change effect.

If the directed technical change effect dominates the within-technology substitution effect, we say that there is strong relative bias of technology (Acemoglu, 2007). In particular, there is strong relative bias if and only if the total effect of an increase in relative skill supply on skill premia is positive; that is, for any dl such that dl_θ/l_θ increases in θ :

$$\left. \frac{d}{d\mu} \frac{w_\theta(l + \mu dl, \phi^*(l))}{w_{\tilde{\theta}}(l + \mu dl, \phi^*(l))} \right|_{\mu=0} + \left. \frac{d}{d\mu} \frac{w_\theta(l, \phi^*(l + \mu dl))}{w_{\tilde{\theta}}(l, \phi^*(l + \mu dl))} \right|_{\mu=0} \geq 0 \quad \text{for all } \theta \geq \tilde{\theta}. \quad (17)$$

In Loebbing (2020), I show that strong relative bias can occur only if the aggregate production function fails to be quasiconcave in labor when accounting for the endogenous adjustment of technology. Such a failure of quasiconcavity can arise from the monopolistic competition setup presented above (Acemoglu, 2007; Loebbing, 2020). Thus, strong relative bias is possible in the present framework.

In the CES case, there is a simple parametric condition for strong relative bias. In particular, one can show that, in the CES case, the total effect of a labor supply change dl on the (log) skill premium between types $\theta \geq \tilde{\theta}$ is given by (see Appendix A.3)

$$\underbrace{-\frac{1}{\sigma} \left(\frac{dl_\theta}{l_\theta} - \frac{dl_{\tilde{\theta}}}{l_{\tilde{\theta}}} \right)}_{=:\gamma^{CES}} + \underbrace{\frac{(\sigma-1)^2}{(\delta-1)\sigma^2 + \sigma} \left(\frac{dl_\theta}{l_\theta} - \frac{dl_{\tilde{\theta}}}{l_{\tilde{\theta}}} \right)}_{=:\rho^{CES}}.$$

The first term is the within-technology substitution effect while the second term captures directed technical change. For an increase in relative skill supply, we have that $dl_\theta/l_\theta \geq dl_{\tilde{\theta}}/l_{\tilde{\theta}}$; hence, the within-technology substitution effect is negative whereas the directed technical change effect is positive. The total effect is positive if and only if

$$\gamma^{CES} + \rho^{CES} \geq 0, \quad (18)$$

which provides a simple condition for strong relative bias in the CES case.

5 Tax Reforms with Directed Technical Change

Starting from a given tax T , a tax reform is represented by the change from T to $T + \mu\tau$, where $\mu \in \mathbb{R}_+$ is a scalar and $\tau : y \mapsto \tau(y) \in \mathbb{R}$ is a twice continuously differentiable, real-valued function. In this notation, μ is the scaling factor of the tax reform while τ indicates its direction: If $\tau(y)$ is positive

(negative) at some income level y , the reform raises (lowers) the tax burden for workers who earn y .

The curvature of τ , relative to the curvature of T , governs the progressivity of the reform. In particular, define the rate of progressivity of a tax schedule T as minus the elasticity of the marginal retention rate R'_T with respect to income (e.g. [Sachs et al., 2020](#)):

$$P_T(y) := -\frac{R''_T(y)y}{R'_T(y)}.$$

The rate of progressivity measures how quickly the marginal retention rate falls in income. It constitutes an important measure of how progressive a tax schedule is (e.g. [Heathcote, Storesletten and Violante, 2017](#)). Consequently, I call a tax reform progressive (regressive) if the post-reform tax has a higher (lower) rate of progressivity than the pre-reform tax everywhere.

Definition 2. Starting from tax T the reform (τ, μ) is progressive (regressive) if and only if

$$P_{\tilde{T}}(y) \geq (\leq) P_T(y) \quad \text{for all } y,$$

where $\tilde{T} := T + \mu\tau$ denotes the post-reform tax function.

This definition of progressivity is equivalent to the following relation between the direction of the reform τ and the initial tax schedule T .

Lemma 2. Starting from any tax function T , the reform (τ, μ) is progressive if and only if

$$\frac{\tau'(y)}{1 - T'(y)} \geq \frac{\tau'(\tilde{y})}{1 - T'(\tilde{y})} \quad \forall y \geq \tilde{y}.$$

Proof. See Appendix [A.6](#). □

This equivalence will turn out useful in the analysis below.

In the following, I focus on the local effects of a reform in the direction of τ , that is, the effects on economic outcomes of changing T to $T + \mu\tau$ as $\mu \rightarrow 0$. Note that this does not lead to confusion with the definition of progressivity, because, as indicated by Lemma 2, the definition of progressivity only depends on the direction τ of a reform but not on the scaling factor μ .

Moreover, I assume without loss of generality that worker types are ordered according to their wage under the initial tax system, that is, $w_\theta(T) \leq w_{\tilde{\theta}}(T)$ if $\theta \leq \tilde{\theta}$ under the initial tax T .

5.1 Effects on the Wage Distribution

I analyze the effects of progressive and regressive tax reforms on relative wages between worker types. The focus is on how these effects are shaped by the presence of directed technical change.

The effects of taxes on wages are transmitted via labor supply. To express them compactly, I introduce the following notation. First, consider the within-technology substitution effects of changes in labor

supply, that is, the effects on wages when holding technology constant. The elasticity of the wage of type θ with respect to the type's own labor supply at constant technology is denoted by

$$\gamma_{\theta}^{own} := \frac{l_{\theta}}{w_{\theta}} \frac{\partial w_{\theta}(l_{\theta}, l, \phi^*(l))}{\partial l_{\theta}} .$$

Here, the notation $w_{\theta}(l_{\theta}, l, \phi^*(l))$ shall indicate that the type's own labor supply l_{θ} typically has an effect on the wage that is distinct from the effect of other changes in the labor input function l . Consider for example equation (15) for wages in the CES case: labor supply l_{θ} enters the equation for wage w_{θ} separately from the overall labor input l . The elasticity of w_{θ} with respect to a different type's labor input $l_{\tilde{\theta}}$ is then denoted by

$$\gamma_{\theta, \tilde{\theta}} := \frac{l_{\tilde{\theta}}}{w_{\theta}} D_{l_{\tilde{\theta}}} w_{\theta}(l_{\theta}, l, \phi^*(l)) ,$$

where $D_{l_{\tilde{\theta}}}$, as before, stands for the Gateaux derivative in direction of $l_{\tilde{\theta}}$.

Next, consider the directed technical change effects of labor supply changes, that is, the effects on wages transmitted via induced technical change. These effects can be written conveniently as the difference between the total effect and the within-technology substitution effect. In particular, let $w_{\theta}^*(l_{\theta}, l) := w_{\theta}(l_{\theta}, l, \phi^*(l))$ denote wages as a function of labor inputs, accounting for the endogenous adjustment of technology. Then, the directed technical change elasticity of wage w_{θ} with respect to type θ 's own labor supply is

$$\rho_{\theta}^{own} := \frac{l_{\theta}}{w_{\theta}} \frac{\partial w_{\theta}^*(l_{\theta}, l)}{\partial l_{\theta}} - \gamma_{\theta}^{own} .$$

Analogously, the cross-type directed technical change elasticity of w_{θ} with respect to $l_{\tilde{\theta}}$ is

$$\rho_{\theta, \tilde{\theta}} := \frac{l_{\tilde{\theta}}}{w_{\theta}} D_{l_{\tilde{\theta}}} w_{\theta}^*(l_{\theta}, l) - \gamma_{\theta, \tilde{\theta}} .$$

Finally, to describe the effects of tax reforms on economic outcomes, I write equilibrium variables as a function of the tax, that is, the equilibrium value of a variable x (e.g. wages or labor inputs) under tax T is denoted by $x(T)$.¹⁹

The relative change in wage w_{θ} in response to the tax reform τ is then denoted by

$$\hat{w}_{\theta, \tau}(T) := \frac{1}{w_{\theta}(T)} \left. \frac{dw_{\theta}(T + \mu\tau)}{d\mu} \right|_{\mu=0}$$

and, analogously, the relative change in labor input l_{θ} is

$$\hat{l}_{\theta, \tau}(T) := \frac{1}{l_{\theta}(T)} \left. \frac{dl_{\theta}(T + \mu\tau)}{d\mu} \right|_{\mu=0} .$$

¹⁹Note that in some cases this involves an abuse of notation. Above, for example, I write $w_{\theta}(l, \phi)$ to denote wages as a function of labor inputs and technology; now I use $w_{\theta}(T)$ to denote wages as a function of the tax. The latter is meant as a short cut for $w_{\theta}(l(T), \phi^*(l(T)))$, where $l(T)$ denotes labor inputs under tax T .

With this notation, the effects of tax reform τ on wages can be written as

$$\widehat{w}_{\theta,\tau}(T) = \gamma_{\theta}^{own} \widehat{l}_{\theta,\tau}(T) + \int_{\underline{\theta}}^{\bar{\theta}} \gamma_{\theta,\tilde{\theta}} \widehat{l}_{\tilde{\theta},\tau}(T) d\tilde{\theta} + \rho_{\theta}^{own} \widehat{l}_{\theta,\tau}(T) + \int_{\underline{\theta}}^{\bar{\theta}} \rho_{\theta,\tilde{\theta}} \widehat{l}_{\tilde{\theta},\tau}(T) d\tilde{\theta} \quad (19)$$

for all θ . The first two terms capture the within-technology substitution effects of labor supply changes on wages while the last two terms capture the directed technical change effects.

Labor supply, in turn, responds directly to the tax reform but also to the induced change in wages. In particular, workers' first-order condition (2) determines labor supply as a function of the marginal retention rate and the wage, $l_{\theta}(R'_T, w_{\theta})$. Let the elasticity of $l_{\theta}(R'_T, w_{\theta})$ with respect to the marginal retention rate be denoted by²⁰

$$\epsilon_{\theta}^R := \frac{R'_T(l_{\theta} w_{\theta})}{l_{\theta}} \frac{\partial l_{\theta}(R'_T, w_{\theta})}{\partial R'_T} ;$$

and the elasticity with respect to the wage by

$$\epsilon_{\theta}^w := \frac{w_{\theta}}{l_{\theta}} \frac{\partial l_{\theta}(R'_T, w_{\theta})}{\partial w_{\theta}} .^{21}$$

Then, the response of labor inputs to the tax reform τ can be written as

$$\widehat{l}_{\theta,\tau}(T) = -\epsilon_{\theta}^R \frac{\tau'(y_{\theta}(T))}{1 - T'(y_{\theta}(T))} + \epsilon_{\theta}^w \widehat{w}_{\theta,\tau}(T) , \quad (20)$$

where the first term captures the direct effect of the reform and the second term accounts for the general-equilibrium feedback via wages.

Equations (19) and (20) form a fixed-point problem that captures the equilibrium nature of the problem: labor supply responds to the tax reform and alters wages; the wage adjustment, in turn, feeds back to labor supply. It is instructive to first solve equations (19) and (20) for the labor input response. The result can then be used in equation (19) to obtain a solution for the wage response. To solve for the labor input response, plug (19) into (20):

$$\widehat{l}_{\theta,\tau}(T) = -\epsilon_{\theta}^R \frac{\tau'(y_{\theta}(T))}{1 - T'(y_{\theta}(T))} + \epsilon_{\theta}^w (\gamma_{\theta}^{own} + \rho_{\theta}^{own}) \widehat{l}_{\theta,\tau}(T) + \epsilon_{\theta}^w \int_{\underline{\theta}}^{\bar{\theta}} (\gamma_{\theta,\tilde{\theta}} + \rho_{\theta,\tilde{\theta}}) \widehat{l}_{\tilde{\theta},\tau}(T) d\tilde{\theta} . \quad (21)$$

This fixed-point equation expresses the labor input responses $\widehat{l}_{\theta,\tau}$, via wage responses, as a function of themselves. I solve it by an iteration procedure.²² Within the iteration steps, I disentangle the feedback

²⁰Note that the marginal retention rate R'_T is a real-valued function of income, so the derivative with respect to the marginal retention rate is again a functional derivative. I provide a rigorous definition of this derivative in Appendix A.2.

²¹Note that these labor supply elasticities apply to the implicit solution of workers' first order condition $l_{\theta}(R'_T, w_{\theta})$ and thus account for the potential non-linearity of the retention function R_T . See Appendix A.2 for how they relate to the (more standard) labor supply elasticity with a linearized tax function.

²²This procedure uses the techniques developed by Sachs et al. (2020). It slightly deviates from their procedure, even when ignoring the addition of directed technical change effects, because I do not use what Sachs et al. (2020) call "elasticities of equilibrium labor" (see their Lemma 1). In the present setting, the advantage of my approach is that it allows me to derive analytical results for the slope of the effects over the type space and thus for their impact on relative labor inputs and relative wages.

effects that are purely transmitted via directed technical change from those transmitted via within-technology substitution. Thereby, I obtain a decomposition of the total labor input response into a substitution and a directed technical change component. The slope of the directed technical change component over the type space can then be signed for the case of a progressive tax reform, using the predictions of directed technical change theory. If there is strong relative bias, the sign extends to the slope of the total response of labor inputs.

Lemma 3. *Fix an initial tax T and suppose that*

$$\sup_{\theta \in \Theta} [(\epsilon_{\theta}^w \rho_{\theta}^{own})^2] + \int_{\underline{\theta}}^{\bar{\theta}} \int_{\underline{\theta}}^{\bar{\theta}} (\epsilon_{\theta}^w \rho_{\theta, \tilde{\theta}})^2 d\tilde{\theta} d\theta + 2\sqrt{\int_{\underline{\theta}}^{\bar{\theta}} \int_{\underline{\theta}}^{\bar{\theta}} (\epsilon_{\theta}^w \rho_{\theta}^{own} \epsilon_{\theta}^w \rho_{\theta, \tilde{\theta}})^2 d\tilde{\theta} d\theta} < 1 \quad (22)$$

$$\sup_{\theta \in \Theta} [(\epsilon_{\theta}^w \zeta_{\theta}^{own})^2] + \int_{\underline{\theta}}^{\bar{\theta}} \int_{\underline{\theta}}^{\bar{\theta}} (\epsilon_{\theta}^w \zeta_{\theta, \tilde{\theta}})^2 d\tilde{\theta} d\theta + 2\sqrt{\int_{\underline{\theta}}^{\bar{\theta}} \int_{\underline{\theta}}^{\bar{\theta}} (\epsilon_{\theta}^w \zeta_{\theta}^{own} \epsilon_{\theta}^w \zeta_{\theta, \tilde{\theta}})^2 d\tilde{\theta} d\theta} < 1, \quad (23)$$

where $\zeta_{\theta, \tilde{\theta}} := \gamma_{\theta, \tilde{\theta}} + \rho_{\theta, \tilde{\theta}}$ and $\zeta_{\theta}^{own} := \gamma_{\theta}^{own} + \rho_{\theta}^{own}$.²³

Then, the relative effect of tax reform τ on labor inputs can be written as

$$\hat{l}_{\theta, \tau}(T) = \sum_{n=0}^{\infty} \hat{l}_{\theta, \tau}^{(n)}(T) \quad (24)$$

for all $\theta \in \Theta$, where

$$\begin{aligned} \hat{l}_{\theta, \tau}^{(0)}(T) &= -\epsilon_{\theta}^R \frac{\tau'(y_{\theta}(T))}{1 - T'(y_{\theta}(T))} \\ \hat{l}_{\theta, \tau}^{(n)}(T) &= \epsilon_{\theta}^w \zeta_{\theta}^{own} \hat{l}_{\theta, \tau}^{(n-1)}(T) + \epsilon_{\theta}^w \int_{\underline{\theta}}^{\bar{\theta}} \zeta_{\theta, \tilde{\theta}} \hat{l}_{\tilde{\theta}, \tau}^{(n-1)}(T) d\tilde{\theta} \quad \forall n > 0. \end{aligned}$$

The total effect on labor inputs can be decomposed as follows,

$$\hat{l}_{\theta, \tau}(T) = \underbrace{-\epsilon_{\theta}^R \frac{\tau'(y_{\theta}(T))}{1 - T'(y_{\theta}(T))}}_{=: \widetilde{DE}_{\theta, \tau}(T)} + \underbrace{\sum_{n=1}^{\infty} \widetilde{TE}_{\theta, \tau}^{(n)}(T)}_{=: \widetilde{TE}_{\theta, \tau}(T)} + \underbrace{\sum_{n=1}^{\infty} \widetilde{SE}_{\theta, \tau}^{(n)}(T)}_{=: \widetilde{SE}_{\theta, \tau}(T)}, \quad (25)$$

²³Conditions (22) and (23) ensure that the series in equations (24) and (25) converge. They are sufficient but generally not necessary for convergence. If the conditions are not satisfied, the equilibrium may be unstable in the sense that an increase in some types' labor inputs may trigger a wage adjustment that is more than sufficient to justify the initial increase in labor inputs. I check that the conditions are satisfied in the quantitative analysis.

where (omitting the argument T)

$$\begin{aligned}
\widetilde{TE}_{\theta,\tau}^{(1)} &= \epsilon_{\theta}^w \rho_{\theta}^{own} (-\epsilon_{\theta}^R) \frac{\tau'(y_{\theta}(T))}{1 - T'(y_{\theta}(T))} + \epsilon_{\theta}^w \int_{\underline{\theta}}^{\bar{\theta}} \rho_{\theta,\bar{\theta}} (-\epsilon_{\bar{\theta}}^R) \frac{\tau'(y_{\bar{\theta}}(T))}{1 - T'(y_{\bar{\theta}}(T))} d\bar{\theta} \\
\widetilde{TE}_{\theta,\tau}^{(n)} &= \epsilon_{\theta}^w \rho_{\theta}^{own} \widetilde{TE}_{\theta,\tau}^{(n-1)} + \epsilon_{\theta}^w \int_{\underline{\theta}}^{\bar{\theta}} \rho_{\theta,\bar{\theta}} \widetilde{TE}_{\bar{\theta},\tau}^{(n-1)} d\bar{\theta} \quad \forall n > 1 \\
\widetilde{SE}_{\theta,\tau}^{(1)} &= \epsilon_{\theta}^w \gamma_{\theta}^{own} (-\epsilon_{\theta}^R) \frac{\tau'(y_{\theta}(T))}{1 - T'(y_{\theta}(T))} + \epsilon_{\theta}^w \int_{\underline{\theta}}^{\bar{\theta}} \gamma_{\theta,\bar{\theta}} (-\epsilon_{\bar{\theta}}^R) \frac{\tau'(y_{\bar{\theta}}(T))}{1 - T'(y_{\bar{\theta}}(T))} d\bar{\theta} \\
\widetilde{SE}_{\theta,\tau}^{(n)} &= \epsilon_{\theta}^w \gamma_{\theta}^{own} (\widetilde{TE}_{\theta,\tau}^{(n-1)} + \widetilde{SE}_{\theta,\tau}^{(n-1)}) + \rho_{\theta}^{own} \widetilde{SE}_{\theta,\tau}^{(n-1)} \\
&\quad + \epsilon_{\theta}^w \int_{\underline{\theta}}^{\bar{\theta}} \left[\gamma_{\theta}^{own} (\widetilde{TE}_{\theta,\tau}^{(n-1)} + \widetilde{SE}_{\theta,\tau}^{(n-1)}) + \rho_{\theta}^{own} \widetilde{SE}_{\theta,\tau}^{(n-1)} \right] d\bar{\theta} \quad \forall n > 1.
\end{aligned}$$

If the labor supply elasticities ϵ_{θ}^w and ϵ_{θ}^R are constant in θ (e.g., because the disutility of labor is isoelastic and T has a constant rate of progressivity), then the components $\widetilde{DE}_{\theta,\tau}$ and $\widetilde{TE}_{\theta,\tau}$ are decreasing in θ for any progressive reform τ .

Moreover, if ϵ_{θ}^w and ϵ_{θ}^R are constant in θ and there is strong relative bias (i.e., (17) is satisfied), then the total labor input response $\widehat{l}_{\theta,\tau}$ is decreasing in θ for any progressive reform τ .

Proof. See Appendix A.7. □

Equation (24) expresses the labor input change induced by reform τ as the sum over successive rounds of general-equilibrium adjustments, capturing feedback loops from labor supply to wages and back to labor supply. The first summand $\widehat{l}_{\theta,\tau}^{(0)}(T)$ is the direct effect of the reform on labor supply, holding wages constant. The direct adjustment of labor supply in turn affects wages, which then feeds back into labor supply. This first-round feedback effect is captured by $\widehat{l}_{\theta,\tau}^{(1)}(T)$. The labor supply change $\widehat{l}_{\theta,\tau}^{(1)}(T)$ then induces another adjustment of wages, which again affects labor supply, and so on.²⁴

Equation (25) decomposes the total labor input change into three components. The first is the direct effect of reform τ , holding wages constant. The second term isolates the part of the general-equilibrium feedback in which the effect of labor supply on wages is purely transmitted via directed technical change. The third term collects the remaining parts of the feedback, containing within-technology substitution effects from labor supply on wages.

If labor supply elasticities are constant across workers, the effect of the first two components on relative labor inputs can be signed. In particular, the labor supply elasticities ϵ_{θ}^w and ϵ_{θ}^R depend on the elasticity of workers' disutility of labor $v(l)$ and on the rate of progressivity of the initial tax schedule (see Appendix A.2). If the disutility of labor is isoelastic and the initial tax schedule has a constant rate of progressivity, the labor supply elasticities are constant as well.²⁵ In this case, both the direct effect

²⁴Mathematically, the series representation in equation (24) is the Neumann series expansion of the solution to the fixed point equation (21). In particular, the fixed point equation can be written abstractly as $(I - X)\widehat{l}_{\tau} = Z$, where I denotes the identity function, X is a linear operator on the space of real-valued functions on Θ , and Z is the direct effect of τ on labor supply. Inverting $I - X$ yields $\widehat{l}_{\tau} = (I - X)^{-1}Z$. By Neumann series expansion, this is equivalent to $\widehat{l}_{\tau} = \sum_{n=0}^{\infty} X^n Z$.

²⁵Heathcote et al. (2017) show that a tax schedule with a constant rate of progressivity is a good approximation to the actual tax and transfer system in the US. Thus, the restriction to an initial tax with constant rate of progressivity is empirically

$\widetilde{DE}_{\theta,\tau}$ and the technical change effect $\widetilde{TE}_{\theta,\tau}$ of a progressive reform on labor inputs are decreasing in θ , which implies that they reduce the labor input of more skilled relative to less skilled workers.

This result follows from weak relative bias. Intuitively, a progressive tax reform raises marginal tax rates by more for more skilled workers. With constant labor supply elasticities and holding wages constant, this leads to a reduction in relative skill supply (i.e., the direct effect $\widetilde{DE}_{\theta,\tau}$ decreases in θ). By weak bias, this reduction in relative skill supply induces technical change that reduces skill premia (equalizing technical change, henceforth). Again under constant labor supply elasticities, such equalizing technical change feeds back into a further reduction in relative skill supply, which, in turn, induces further equalizing technical change. Summing over the thus induced rounds of reductions in relative skill supply eventually gives rise to the term $\widetilde{TE}_{\theta,\tau}$, which must therefore also reduce relative skill supply (i.e., it decreases in θ).

Finally, if there is strong relative bias, directed technical change dominates within-technology substitution effects and the slope of $\widetilde{TE}_{\theta,\tau}$ extends to the slope of the total general-equilibrium adjustment term $\widetilde{TE}_{\theta,\tau} + \widetilde{SE}_{\theta,\tau}$. In this case, the total impact of a progressive reform on relative skill supply is unambiguously negative: the negative direct effect is amplified by the general-equilibrium feedback.

Next, we can plug the results from Lemma 3 into equation (19) to obtain the following expression for the relative wage effects of a tax reform τ .

Proposition 1. *Fix an initial tax T and let conditions (22) and (23) be satisfied.*

Then, the relative effect of tax reform τ on wages can be written as

$$\begin{aligned} \widehat{w}_{\theta,\tau} = & \underbrace{\rho_{\theta}^{own} \left(\widetilde{DE}_{\theta,\tau}(T) + \widetilde{TE}_{\theta,\tau}(T) \right) + \int_{\underline{\theta}}^{\bar{\theta}} \rho_{\theta,\tilde{\theta}} \left(\widetilde{DE}_{\tilde{\theta},\tau}(T) + \widetilde{TE}_{\tilde{\theta},\tau}(T) \right) d\tilde{\theta}}_{=:TE_{\theta,\tau}(T)} \\ & + \underbrace{\gamma_{\theta}^{own} \widehat{l}_{\theta,\tau}(T) + \int_{\underline{\theta}}^{\bar{\theta}} \gamma_{\theta,\tilde{\theta}} \widehat{l}_{\tilde{\theta},\tau}(T) d\tilde{\theta} + \rho_{\theta}^{own} \widetilde{SE}_{\theta,\tau}(T) + \int_{\underline{\theta}}^{\bar{\theta}} \rho_{\theta,\tilde{\theta}} \widetilde{SE}_{\tilde{\theta},\tau}(T) d\tilde{\theta}}_{=:SE_{\theta,\tau}(T)} \end{aligned} \quad (26)$$

for all $\theta \in \Theta$, where $\widetilde{DE}_{\theta,\tau}$, $\widetilde{TE}_{\theta,\tau}(T)$, $\widetilde{SE}_{\theta,\tau}(T)$ and $\widehat{l}_{\theta,\tau}(T)$ are given by Lemma 3.

If the labor supply elasticities ϵ_{θ}^w and ϵ_{θ}^R are constant in θ (e.g., because the disutility of labor is isoelastic and T has a constant rate of progressivity), then the component $TE_{\theta,\tau}(T)$ is decreasing in θ for any progressive reform τ .

Moreover, if ϵ_{θ}^w and ϵ_{θ}^R are constant in θ and there is strong relative bias (i.e., (17) is satisfied), then the total wage response $\widehat{w}_{\theta,\tau}$ is decreasing in θ for any progressive reform τ .

Proof. See Appendix A.8. □

The terms in equation (26) follow directly from Lemma 3. The term $TE_{\theta,\tau}$ captures the directed technical change effect on wages induced by the direct response of labor supply to the tax reform (as given by $\widetilde{DE}_{\theta,\tau}$) and by the general-equilibrium feedback on labor supply that itself is purely transmitted via

relevant.

directed technical change (as given by $\widetilde{TE}_{\theta,\tau}$). The term $SE_{\theta,\tau}$ captures all within-technology substitution effects on wages plus the directed technical change effect that is induced by the general-equilibrium feedback effects transmitted via within-technology substitution effects.

With constant labor supply elasticities, the technical change component $TE_{\theta,\tau}$ causes a fall in skill premia in response to any progressive tax reform. Intuitively, this term collects all successive rounds of general-equilibrium feedback from directed technical change to labor supply and back to technical change. The direct effect of a progressive reform on relative skill supply is negative, which induces equalizing technical change. This equalizing technical change further reduces relative skill supply, which then again induces equalizing technical change, and so on. We thus obtain a sum of equalizing technical changes, which must be equalizing itself (i.e., decreasing in θ).

With strong relative bias, the same reasoning holds for the entire wage response to the tax reform. The direct effect of a progressive reform on relative skill supply is negative. By strong relative bias, the ensuing fall in relative skill supply leads to a fall in skill premia, even when accounting for both directed technical change and within-technology substitution effects. This fall in skill premia leads to a further reduction in relative skill supply, which further depresses skill premia, and so on. Hence, with strong relative bias, a progressive tax reform leads to a reduction in skill premia.

Without strong relative bias, within-technology substitution effects can be decisive for the total effect on relative wages. Within-technology substitution effects are exclusively determined by the aggregate production function F . Since, so far, we have imposed too little structure on F to sign the within-technology substitution effects of a reduction in relative skill supply unambiguously, no clear result for the total relative wage effect of a progressive tax reform emerges without strong relative bias.

5.2 CES Case

To gain more insights into the relative wage effects of tax reforms in the absence of strong bias, I specialize the previous results to the case with a CES production function. The resulting expression also allows for a straightforward quantitative analysis of the effects of tax reforms on wage inequality, which I pursue in Section 7.

Under the CES production structure, the wage elasticities γ_{θ}^{own} , $\gamma_{\theta,\tilde{\theta}}$, ρ_{θ}^{own} , and $\rho_{\theta,\tilde{\theta}}$ take especially simple forms, being independent of type θ . In particular, the elasticity of wage w_{θ} with respect to type θ 's own labor supply when holding technology constant is given by²⁶

$$\gamma_{\theta}^{own} = \gamma^{CES} \left(:= -\frac{1}{\sigma} \right) \text{ for all } \theta, \quad (27)$$

while the cross-type elasticity at constant technology is

$$\gamma_{\theta,\tilde{\theta}} = -\gamma^{CES} \frac{l_{\tilde{\theta}} w_{\tilde{\theta}} h_{\tilde{\theta}}}{F(l, \phi)} \text{ for all } \theta, \tilde{\theta}.$$

²⁶See Appendix A.3 for detailed derivations of the wage elasticities in the CES case.

For the own-type directed technical change elasticity, we obtain

$$\rho_{\theta}^{own} = \rho^{CES} \left(:= \frac{(\sigma - 1)^2}{(\delta - 1)\sigma^2 + \sigma} \right) \text{ for all } \theta, \quad (28)$$

and the cross-type directed technical change elasticity becomes

$$\rho_{\theta, \tilde{\theta}} = -\rho^{CES} \frac{l_{\tilde{\theta}} w_{\tilde{\theta}} h_{\tilde{\theta}}}{F(l, \phi)} \text{ for all } \theta, \tilde{\theta}. \quad (29)$$

Using these elasticities in Proposition 1, we obtain the following expression for the relative wage effects of a tax reform τ in the CES case.

Corollary 1. *Fix an initial tax T and assume that F and Φ are CES as introduced in Section 3.6. Moreover, let the elasticities ϵ_{θ}^w and ϵ_{θ}^R be constant in θ , that is, $\epsilon_{\theta}^w = \epsilon^w$ and $\epsilon_{\theta}^R = \epsilon^R$ for all $\theta \in \Theta$. Then, the relative wage effect of tax reform τ satisfies*

$$\begin{aligned} \widehat{w}_{\theta, \tau}(T) = & \rho^{CES}(-\bar{\epsilon}^R) \frac{\tau'(y_{\theta}(T))}{1 - T'(y_{\theta}(T))} - \rho^{CES} \int_{\underline{\theta}}^{\bar{\theta}} \frac{w_{\tilde{\theta}} l_{\tilde{\theta}} h_{\tilde{\theta}}}{F(l(T), \phi^*(T))} (-\bar{\epsilon}^R) \frac{\tau'(y_{\tilde{\theta}}(T))}{1 - T'(y_{\tilde{\theta}}(T))} d\tilde{\theta} \\ & + \gamma^{CES}(-\bar{\epsilon}^R) \frac{\tau'(y_{\theta}(T))}{1 - T'(y_{\theta}(T))} - \gamma^{CES} \int_{\underline{\theta}}^{\bar{\theta}} \frac{w_{\tilde{\theta}} l_{\tilde{\theta}} h_{\tilde{\theta}}}{F(l(T), \phi^*(T))} (-\bar{\epsilon}^R) \frac{\tau'(y_{\tilde{\theta}}(T))}{1 - T'(y_{\tilde{\theta}}(T))} d\tilde{\theta} \end{aligned} \quad (30)$$

for all $\theta \in \Theta$, where

$$\bar{\epsilon}^R := \frac{\epsilon^R}{1 - (\gamma^{CES} + \rho^{CES})\epsilon^w}.$$

For any progressive tax reform τ , the directed technical change effect on wages (i.e., the first line of equation (30)) decreases in θ , while the within-technology substitution effect (i.e., the second line) increases in θ .

If there is strong relative bias (i.e., (18) is satisfied), the total wage response $\widehat{w}_{\theta, \tau}$ is decreasing in θ for any progressive reform τ . Otherwise, the total wage response is increasing in θ for any progressive reform.

Proof. See Appendix A.8. □

Equation (30) separates the directed technical change effect of a tax reform (the first line) from the within-technology substitution effect (the second line). It shows that the directed technical change effect of a progressive tax reform always reduces skill premia. Hence, any progressive reform induces equalizing technical change. Inversely, a regressive reform induces skill-biased technical change.

Crucially, the CES case allows to sign the within-technology substitution effect of a tax reform. In particular, the within-technology substitution effect always counteracts the directed technical change effect: in case of a progressive reform, it raises skill premia, in case of regressive reform, it lowers them. Whether directed technical change or within-technology substitution effects dominate, depends on whether there is strong relative bias. With strong relative bias, directed technical change dominates and, in accordance with Proposition 1, any progressive tax reform lowers skill premia. Without strong relative bias, the effects are reversed.

The finding that progressive (regressive) tax reforms induce equalizing (skill-biased technical change) has several consequences. First, it is empirically relevant. The US tax and transfer system, for example, underwent regressive reforms since the 1980s (Piketty and Saez, 2007) while US wage inequality surged, a development often contributed in parts to skill-biased technical change. My results suggest that skill-biased technical change may be a consequence of the contemporaneous regressive tax reforms, which I examine quantitatively in Section 7.

Moreover, the insight that progressive tax reforms reduce the skill-bias of technology suggests that, when accounting for directed technical change, the optimal tax system may become more progressive. This is what I turn to next.²⁷

6 Optimal Taxes with Directed Technical Change

In this section, I characterize the optimal tax system for a given welfare function. The focus is on how the optimal tax is affected by directed technical change.

Welfare is measured by a Bergson-Samuelson welfare function $V : \{u_\theta\}_{\theta \in \Theta} \mapsto V(\{u_\theta\}_{\theta \in \Theta})$ that is strictly increasing in all arguments. The marginal welfare weight of an individual worker of type θ is obtained as

$$g_\theta(\{u_\theta\}_{\theta \in \Theta}) = \frac{1}{h_\theta} D_{u_\theta} V(\{u_\theta\}_{\theta \in \Theta}) ,$$

where D_{u_θ} , analogously to the derivative D_{l_θ} used before, is the Gateaux derivative in direction of the individual utility level u_θ .²⁸

I assume that V is scaled such that the average welfare weight across all workers equals one at the optimal tax and impose that g is continuous in θ whenever u is continuous. In addition and more substantially, the welfare function is supposed to value equity across workers in the following sense.

Assumption 3. *For any utility profile $\{u_\theta\}_{\theta \in \Theta}$ such that u_θ increases in θ , the marginal welfare weights $g_\theta(\{u_\theta\}_{\theta \in \Theta})$ decrease in θ .*

This assumption ensures that redistributing consumption from workers with high utility to workers with low utility improves welfare.

For the characterization of optimal tax rates it is convenient to denote the average marginal welfare weight of all workers above a given type θ by

$$\tilde{g}_\theta := \frac{1}{1 - H_\theta} \int_\theta^{\bar{\theta}} g_{\tilde{\theta}} h_{\tilde{\theta}} d\tilde{\theta} .$$

Moreover, for a function $x : (\theta, z) \mapsto x_\theta(z)$ (e.g., wages or labor inputs) that depends on θ and poten-

²⁷Besides its impact on the optimal tax scheme, directed technical change also modifies the welfare effects of given tax reforms. I discuss this in Appendix C.1.

²⁸See Appendix A.1 for the rigorous definition of D_{l_θ} , which transfers to D_{u_θ} .

tially further variables z , I denote the derivative of x with respect to θ by

$$x'_\theta(z) := \frac{dx_\theta(z)}{d\theta}$$

and the corresponding semi-elasticity by

$$\hat{x}_\theta(z) := \frac{x'_\theta(z)}{x_\theta(z)}.$$

Finally, without loss of generality, let worker types be ordered according to their wages under the optimum tax schedule, that is, under the optimal tax $w_\theta \leq w_{\tilde{\theta}}$ if $\theta \leq \tilde{\theta}$.

6.1 Optimal Tax Formula

To derive optimal tax rates, I follow the mechanism design approach to optimal taxation. For that, write welfare as a function of consumption and labor allocations instead of utility levels:

$$W(c, l) := V(\{u_\theta(c_\theta, l_\theta)\}_{\theta \in \Theta}).$$

The goal is to find the consumption-labor allocation that maximizes welfare $W(c, l)$ subject to the aggregate resource constraint and to incentive compatibility constraints across worker types. The optimal tax schedule is then obtained as the tax that implements the welfare-maximizing allocation.

The aggregate resource constraint is given by

$$\int_{\underline{\theta}}^{\bar{\theta}} c_\theta h_\theta d\theta = F(l, \phi^*(l)). \quad (31)$$

Incentive compatibility requires

$$u_\theta = \max_{\tilde{\theta} \in \Theta} \left\{ c_{\tilde{\theta}} - v\left(\frac{w_{\tilde{\theta}} l_{\tilde{\theta}}}{w_\theta}\right) \right\} \quad \text{for all } \theta.$$

I restrict attention to instances of the model where the labor input under the optimal tax is continuously differentiable in θ . Moreover, I assume that this property extends from labor inputs to wages as follows.

Assumption 4. *If l is continuously differentiable in θ , then $D_{l_\theta} F(l, \phi)$ is continuously differentiable in θ for all $\phi \in \Phi$.*

*Moreover, the worker density h is continuously differentiable in θ .*²⁹

Under this restriction and with the wage function w_θ increasing in θ at the optimum, the incentive

²⁹The worker density being C^1 ensures that the CES case satisfies Assumption 4.

compatibility constraint is equivalent to the following conditions:

$$c'_\theta = v'(l_\theta)(w'_\theta l_\theta + w_\theta l'_\theta) \frac{1}{w_\theta} \quad \text{for almost every } \theta, \quad (32)$$

$$y'_\theta \geq 0 \quad \text{for almost every } \theta. \quad (33)$$

As is usual in the literature, I drop the monotonicity requirement (33) and study the relaxed problem of maximizing welfare subject to (31) and (32).³⁰

From the incentive compatibility and resource constraints, consumption levels can be derived as a function of labor inputs. I substitute this function into the welfare function W and compute first-order conditions with respect to labor inputs. Using workers' first-order conditions to reintroduce marginal tax rates then yields an insightful expression for optimal marginal tax rates. As usual in the literature (e.g. [Diamond, 1998](#)), this expression features the elasticity of labor supply with respect to the marginal retention rate that would be obtained if, hypothetically, the tax function were linear in a neighborhood of the respective income level. This elasticity, denoted by e_θ , is formally given by

$$e_\theta = \frac{v'(l_\theta)}{v''(l_\theta)l_\theta}. \quad (34)$$

Proposition 2. *Suppose the labor input l under the optimal tax is continuously differentiable in θ . Then, at every type θ , optimal marginal tax rates satisfy the following conditions.*

$$\frac{T'(y_\theta)}{1 - T'(y_\theta)} = PE_\theta^* + SE_\theta^* + TE_\theta^*,$$

where

$$\begin{aligned} PE_\theta^* &= \left(1 + \frac{1}{e_\theta}\right) \frac{1 - N_{w_\theta}(1 - \tilde{g}_\theta)}{n_{w_\theta} w_\theta} \\ SE_\theta^* &= \int_{\underline{\theta}}^{\bar{\theta}} \frac{1 - T'(y_{\tilde{\theta}})}{1 - T'(y_\theta)} \frac{1 - H_{\tilde{\theta}}}{h_\theta w_\theta} (1 - \tilde{g}_{\tilde{\theta}}) y_{\tilde{\theta}} D_{l_\theta} \hat{w}_{\tilde{\theta}}(l_{\tilde{\theta}}, l, \phi^*(l)) d\tilde{\theta} \\ TE_\theta^* &= \int_{\underline{\theta}}^{\bar{\theta}} \frac{1 - T'(y_{\tilde{\theta}})}{1 - T'(y_\theta)} \frac{1 - H_{\tilde{\theta}}}{h_\theta w_\theta} (1 - \tilde{g}_{\tilde{\theta}}) y_{\tilde{\theta}} \left(D_{l_\theta} \hat{w}_{\tilde{\theta}}^*(l_{\tilde{\theta}}, l) - D_{l_\theta} \hat{w}_{\tilde{\theta}}(l_{\tilde{\theta}}, l, \phi^*(l)) \right) d\tilde{\theta}, \end{aligned}$$

all variables are evaluated at equilibrium under the optimal tax T , and N and n denote the cumulative distribution and the density function of wages at the optimum.³²

Moreover, if $\limsup_{\theta \rightarrow \bar{\theta}} l'_\theta < \infty$ and $\liminf_{\theta \rightarrow \underline{\theta}} l'_\theta > -\infty$ under the optimal tax,³³ then the following holds:

³⁰In all simulations of optimal taxes, I verify that the monotonicity condition (33) holds at the optimum.

³¹See Appendix A.2 for the relation of this elasticity to the labor supply elasticities introduced in the tax reform analysis of Section 5, which account for the potential non-linearity of the tax function.

³²Recall from Section 5 that $w_\theta^*(l_{\tilde{\theta}}, l) := w_{\tilde{\theta}}(l_{\tilde{\theta}}, l, \phi^*(l))$ denotes the wage of type $\tilde{\theta}$ as a function of labor inputs when accounting for the endogenous adjustment of technology. Hence, the difference $D_{l_\theta} \hat{w}_{\tilde{\theta}}^*(l_{\tilde{\theta}}, l) - D_{l_\theta} \hat{w}_{\tilde{\theta}}(l_{\tilde{\theta}}, l, \phi^*(l))$ in the term TE_θ^* captures the directed technical change effect of labor input l_θ on the return to skill $\hat{w}_{\tilde{\theta}}$.

³³These conditions guarantee that the distribution of labor inputs is well behaved at the top and at the bottom, in the sense that its density is continuous and strictly positive on some neighborhood of the top or the bottom type, respectively. This

1. $TE_{\underline{\theta}}^* \leq 0$ and $TE_{\bar{\theta}}^* \geq 0$.

2. If there is strong relative bias (i.e., (17) holds), then $TE_{\underline{\theta}}^* + SE_{\underline{\theta}}^* \leq 0$ and $SE_{\bar{\theta}}^* + TE_{\bar{\theta}}^* \geq 0$.

Proof. See Appendix A.10. □

Proposition 2 provides an expression that decomposes the optimal marginal tax rates into three terms. The first term PE_{θ}^* is the standard expression from a setting with exogenous wages. It is zero at the bottom and the top income level, reflecting the well-known result that the optimal marginal tax rate is zero for the highest and the lowest income earner when wages are exogenous.

The last term, TE_{θ}^* , captures the impact of directed technical change effects on the optimal tax. It is negative at the bottom and positive at the top income. The intuition behind this result is closely related to the insight that a progressive tax reform induces equalizing technical change. By reducing marginal tax rates at the bottom and increasing them at the top, the optimal tax schedule stimulates the relative labor supply of less skilled workers, thus inducing firms to use technologies with a higher relative productivity for low-skilled workers. This raises low-skilled workers' wages relative to those of high-skilled workers. In the mechanism design problem, the ensuing compression in the pre-tax wage distribution slackens high-skilled workers' incentive compatibility constraints and widens the scope for redistribution.

The middle term, SE_{θ}^* , stems from within-technology substitution effects. To sign this term individually, further structure on aggregate production is required.

If there is strong relative bias, however, the total general-equilibrium adjustment term $SE_{\theta}^* + TE_{\theta}^*$ inherits its sign from the directed technical change term TE_{θ}^* , being negative at the bottom and positive at the top. Since the partial equilibrium term PE_{θ}^* is zero at the bottom and top, these signs extend directly to optimal tax rates: the optimal marginal tax at the bottom is negative, while the optimal marginal tax at the top is positive.

This finding relates directly to the well-known result of Stiglitz (1982) that the optimal marginal tax at the top is negative with a general production structure featuring complementarity between worker types. If there is strong relative bias, directed technical change overturns this result, leading to a positive top tax rate instead.

It is well known in the literature, however, that results for the optimal marginal tax rate at the highest income do not extend approximately to incomes even slightly below the top. In particular, if the bulk of the right tail of the income distribution is well approximated by a Pareto distribution, the optimal marginal tax rate is approximately constant over this range of incomes and differs significantly from the rate at the very top (e.g. Diamond, 1998). The following Corollary thus provides an expressions for the optimal marginal tax rate in the upper Pareto tail of the income distribution when accounting for directed technical change.

in turn allows to evaluate the effects of perturbing the labor input function at the top type $\bar{\theta}$ and the bottom type $\underline{\theta}$ on relative labor inputs and invoke the directed technical change results of Section 4.

Corollary 2. *Let the conditions of Proposition 2 be satisfied and suppose that the disutility of labor is isoelastic with $e_\theta = e$ for all θ . Moreover, suppose that the wage distribution satisfies*

$$\lim_{\theta \rightarrow \bar{\theta}} \frac{1 - N_{w_\theta}}{n_{w_\theta} w_\theta} = \frac{1}{a},$$

the terms SE_θ^ and TE_θ^* satisfy*

$$\lim_{\theta \rightarrow \bar{\theta}} SE_\theta^* = \overline{SE} \quad \text{and} \quad \lim_{\theta \rightarrow \bar{\theta}} TE_\theta^* = \overline{TE},$$

and welfare weights satisfy

$$\lim_{\theta \rightarrow \bar{\theta}} g_\theta = g^{top}$$

at the optimal tax.

Then, the optimal tax T satisfies

$$\lim_{\theta \rightarrow \bar{\theta}} \frac{T'(y_\theta)}{1 - T'(y_\theta)} = \left(1 + \frac{1}{e}\right) \frac{1}{a} (1 - g^{top}) + \overline{SE} + \overline{TE}.$$

Moreover, the following holds:

1. $\overline{TE} \geq 0$.
2. *If there is strong relative bias (i.e., (17) holds), $\overline{SE} + \overline{TE} \geq 0$.*

Proof. See Appendix A.11. □

The corollary shows that, while the level of the optimal marginal tax rate may differ greatly in the upper Pareto tail and at the very highest income, the impact of directed technical change is qualitatively the same. In both cases, the directed technical change adjustment term calls for higher marginal tax rates. Moreover, if there is strong relative bias, the directed technical change term dominates the within-technology substitution term, such that the total general-equilibrium adjustment $\overline{SE} + \overline{TE}$ is positive. In this case, directed technical change effects lead to an upwards adjustment of marginal tax rates in the Pareto tail even relative to the formula obtained in a setting with exogenous wages.

A drawback of Corollary 2 is that it requires the Pareto coefficient and the general-equilibrium adjustment terms to converge in the right tail of the income distribution under the optimal tax. It would be desirable to obtain the results under empirically observable conditions instead. This becomes possible in the CES case, which I consider in the following.

6.2 CES Case

As in the tax reform analysis of Section 5, the CES case allows for additional analytical insights, especially in the absence of strong relative bias.

Lemma 8 in Appendix A.12 specializes the adjustment terms TE_θ^* and SE_θ^* to the CES case. It shows

that, at the top income, the directed technical change term is given by

$$\rho^{CES}(1 - g_{\bar{\theta}}) ,$$

which is positive because $g_{\bar{\theta}} \leq 1$. At the bottom income, the term takes the form

$$\rho^{CES}(1 - g_{\underline{\theta}}) ,$$

which is negative. While this verifies the general predictions from Proposition 2, the CES case also allows to sign the within-technology substitution term SE_{θ}^* . At the top income, it is given by

$$\gamma^{CES}(1 - g_{\bar{\theta}})$$

and hence negative (recall that $\gamma^{CES} < 0$). At the bottom income, it takes the form

$$\gamma^{CES}(1 - g_{\underline{\theta}}) ,$$

which is positive. Thus, in the CES case, within-technology substitution effects counteract the impact of directed technical change on optimal marginal tax rates, calling for lower marginal tax rates at the top and higher rates at the bottom.³⁴

Using the expressions for SE_{θ}^* and TE_{θ}^* from Lemma 8 in the formula provided by Proposition 2, the following transparent formula for optimal marginal tax rates in the CES case is obtained.

Proposition 3. *Suppose the conditions of Proposition 2 are satisfied and F and Φ take the CES form introduced in Section 3.6. Then, at every type θ , optimal marginal tax rates satisfy the following conditions:*

$$\frac{T'(y_{\theta})}{1 - T'(y_{\theta})} = \left(1 + \frac{1}{e_{\theta}}\right) \frac{1 - B_{\beta_{\theta}}}{b_{\beta_{\theta}}\beta_{\theta}} (1 - \tilde{g}_{\theta}) + \gamma^{CES}(1 - g_{\theta}) + \rho^{CES}(1 - g_{\theta}) , \quad (35)$$

where all variables are evaluated at equilibrium under the optimal tax T , the function $\beta : \theta \mapsto \beta_{\theta}$ is given by

$$\beta_{\theta} := \kappa_{\theta}^{1+\gamma^{CES}+\rho^{CES}} h_{\theta}^{\gamma^{CES}+\rho^{CES}} \quad \forall \theta ,$$

while B and b are the cumulative distribution and the density function of β .

Proof. See Appendix A.12. □

Considering at first again the optimal marginal tax rates at the top and bottom income, we obtain a sharp distinction along the condition for strong relative bias. In particular, the optimal marginal tax at the top income is given by

$$\gamma^{CES}(1 - g_{\bar{\theta}}) + \rho^{CES}(1 - g_{\bar{\theta}}) .$$

In accordance with the general results from Proposition 2, this is positive if there is strong relative bias (i.e., $\gamma^{CES} + \rho^{CES} \geq 0$). Without strong relative bias, in contrast, the optimal marginal tax at the top

³⁴See also Sachs et al. (2020) and Stiglitz (1982) for versions of this result.

is always negative, following the logic of [Stiglitz \(1982\)](#). The reverse holds for the bottom income: the optimal marginal tax at the bottom is given by

$$\gamma^{CES}(1 - g_{\underline{\theta}}) + \rho^{CES}(1 - g_{\underline{\theta}}) ,$$

which is negative if and only if there is strong relative bias.

Next, we can specialize the optimal marginal tax formula from [Proposition 3](#) to the upper Pareto tail of the income distribution. In particular, suppose that, under some initial tax system with a constant marginal tax at the top, the Pareto coefficient of the income distribution is asymptotically constant for high incomes. Then, we obtain the following expression for the asymptotic optimal marginal tax on high incomes in the CES case.

Corollary 3. *Let the conditions of [Proposition 2](#) be satisfied and F and Φ take the CES form introduced in [Section 3.6](#). Suppose at a given tax $\bar{T}$, with $\bar{T}'(y) = \tau^{top}$ for all $y \geq \tilde{y}$ and some threshold $\tilde{y}$, the income distribution satisfies*

$$\lim_{\theta \rightarrow \bar{\theta}} \frac{1 - M_{y_{\theta}}}{m_{y_{\theta}} y_{\theta}} = \frac{1}{a}$$

for some $a > 1$. Moreover, let the disutility of labor be isoelastic with $e_{\theta} = e$ for all θ , and welfare weights satisfy

$$\lim_{\theta \rightarrow \bar{\theta}} g_{\theta} = g^{top}$$

at the optimal tax.

Then, the optimal tax T satisfies

$$\lim_{\theta \rightarrow \bar{\theta}} \frac{T'(y_{\theta})}{1 - T'(y_{\theta})} = \frac{1 - g^{top}}{ae} + \frac{a-1}{a} \gamma^{CES}(1 - g^{top}) + \frac{a-1}{a} \rho^{CES}(1 - g^{top}) . \quad (36)$$

Proof. See [Appendix A.12](#). □

The term $\frac{a-1}{a} \gamma^{CES}(1 - g^{top})$, which corresponds to $\overline{SE}$ in the general case, captures the impact of within-technology substitution effects and leads to a reduction in the optimal marginal tax. The term $\frac{a-1}{a} \rho^{CES}(1 - g^{top})$, corresponding to $\overline{TE}$ in the general case, captures directed technical change and causes an upwards adjustment of the marginal tax. If and only if there is strong relative bias, the directed technical change term dominates and the optimal marginal tax in the Pareto tail exceeds that obtained in a setting with exogenous wages.

This basically confirms the findings from the general case presented in [Corollary 2](#). Yet, in contrast to [Corollary 2](#), the results for the CES case do not require assumptions on the shape of the income distribution under the optimal tax. Instead, they hold under the empirically plausible condition that, under a tax with a constant marginal tax rate at the top, the upper tail of the income distribution resembles a Pareto distribution.

6.3 Comparison to Exogenous Technology Planner

While the previous results provide insightful formulas for the optimal tax with directed technical change, a rigorous assessment of the impact of directed technical change must define a benchmark to which the results are compared. To this end, I introduce the concept of an exogenous technology planner, who ignores directed technical change effects. I then contrast the tax perceived as optimal by the exogenous technology planner with the truly optimal tax, which accounts for directed technical change. To obtain precise analytical results, I restrict attention to the CES case in this comparison.

The exogenous technology planner observes the economy under some initial tax $\bar{T}$, correctly infers all parameters of the economy, but mistakenly believes that technology remains fixed at its current state $\phi^*(\bar{T})$, irrespectively of the tax schedule. Formally, the exogenous technology planner computes the optimal tax based on the equilibrium conditions (1), (2), (10), and (12), but replaces the equilibrium technology condition (11) by the “wrong” equation

$$\phi^*(l) = \phi^*(l(\bar{T})) = \operatorname{argmax}_{\phi \in \Phi} F(l(\bar{T}), \phi) \quad \forall l. \quad ^{35}$$

Proposition 4 in Appendix A.14 shows that the tax $T_{\bar{T}}^{ex}$ perceived as optimal by the exogenous technology planner satisfies the following condition:

$$\frac{T_{\bar{T}}^{ex'}(y_\theta)}{1 - T_{\bar{T}}^{ex'}(y_\theta)} = \left(1 + \frac{1}{e_\theta}\right) \frac{1 - \bar{B}_{\bar{\beta}_\theta}}{\bar{b}_{\bar{\beta}_\theta} \bar{\beta}_\theta} (1 - \tilde{g}_\theta) + \gamma^{CES} (1 - g_\theta), \quad (37)$$

where the function $\bar{\beta} : \theta \mapsto \bar{\beta}_\theta$ is given by

$$\bar{\beta}_\theta := \kappa_\theta^{1+\gamma^{CES}} h_\theta^{\gamma^{CES}} (\phi^*(\bar{T}))^{1+\gamma^{CES}} \quad \forall \theta,$$

while $\bar{B}$ and $\bar{b}$ are the cumulative distribution and the density function of $\bar{\beta}$.

Comparing the exogenous technology planner’s tax rates (37) with the optimal tax in equation (35), there are two differences. First, optimal taxes account for the directed technical change adjustment

$$\rho^{CES} (1 - g_\theta).$$

This term is increasing in θ (as welfare weights are decreasing in θ at the optimum) and in this sense necessitates a progressive adjustment of the tax schedule. The intuition for this adjustment is the same as for the top and bottom tax rate adjustments discussed above: lowering marginal tax rates at the bottom and raising them at the top induces technical change that compresses the wage distribution and hence improves equity.

³⁵Note that the set of optimal taxes computed by the exogenous technology planner for arbitrary initial taxes $\bar{T}$ strictly includes the self-confirming policy equilibrium of Rothschild and Scheuer (2013). Specifically, when setting $\bar{T}$ to the tax in the self-confirming policy equilibrium, the exogenous technology planner’s preferred tax is exactly the self-confirming policy equilibrium tax. Hence, comparing optimal taxes to those computed by the exogenous technology planner for arbitrary initial taxes includes the comparison to the self-confirming policy equilibrium.

The second difference is that the optimal tax formula features the hazard ratio of β whereas the exogenous technology planner uses the hazard ratio of $\bar{\beta}$. The function β can be interpreted as the degree of exogenous inequality in the model: if labor supply were identical across all workers, wages would be proportional to β . In contrast, the function $\bar{\beta}$ is the exogenous technology planner's wrong inference about the degree of exogenous inequality. The exogenous technology planner believes that, if all workers' labor supply were identical, wages would be proportional to $\bar{\beta}$ instead of β .

It can be shown that the exogenous technology planner's measure of exogenous inequality is higher than the true one (see Appendix A.15):

$$\frac{1 - \bar{B}_{\bar{\beta}\theta}}{\bar{b}_{\bar{\beta}\theta}\bar{\beta}_{\theta}} > \frac{1 - B_{\beta\theta}}{b_{\beta\theta}\beta_{\theta}} \quad \forall \theta. \quad (38)$$

This raises the exogenous technology planner's tax rates at all income levels relative to the optimum. The second adjustment due to directed technical change therefore reduces marginal tax rates everywhere.

Intuitively, the exogenous technology planner overestimates the degree of exogenous inequality in the economy because the planner mistakenly believes that the skill bias of the equilibrium technology under the initial tax $\bar{T}$ is exogenous. Since more exogenous inequality calls for higher marginal tax rates, the exogenous technology planner chooses elevated marginal tax rates everywhere.³⁶

At the bottom of the income distribution, the two adjustments point in the same direction. Hence, if we assume exogenous welfare weights, directed technical change calls for unambiguously lower marginal tax rates at the lowest income and, by continuity, in some neighborhood thereof.

At the top of the income distribution, the effects move in opposite directions. Yet, at the highest income, the ABC term (Diamond, 1998) in equations (35) and (37) vanishes, such that the only remaining difference is the (positive) term $\rho^{CES}(1 - g_{\bar{\theta}})$. Hence, with exogenous welfare weights, the optimal marginal tax at the top is unambiguously higher when accounting for directed technical change.

Moreover, when the upper tail of the income distribution has a Pareto shape as in Corollary 3, the exogenous technology planner's preferred tax satisfies (see Corollary 4 in Appendix A.14)

$$\lim_{\theta \rightarrow \bar{\theta}} \frac{T_{\bar{T}}^{ex'}(y_{\theta})}{1 - T_{\bar{T}}^{ex'}(y_{\theta})} = \frac{1 - g^{top}}{ae} + \frac{a - 1}{a} \gamma^{CES}(1 - g^{top}),$$

where a is the Pareto tail parameter of the income distribution and g^{top} the asymptotic welfare weight. This expression is strictly smaller than the optimal marginal tax in Corollary 3. Hence, conditional on the limit welfare weight g^{top} , directed technical change unambiguously raises the optimal marginal tax in the Pareto tail.

To summarize, accounting for directed technical change effects leads to lower marginal tax rates at the bottom of the income distribution and to higher marginal tax rates both at the very top and in the

³⁶An intuition for the positive impact of exogenous inequality on marginal tax rates is that a higher degree of exogenous inequality implies that the pre-tax income distribution will respond less strongly to rising tax rates, such that redistribution can be achieved with less distortions and hence at a lower efficiency loss.

Pareto tail of the income distribution.

7 Quantitative Analysis

To assess the quantitative relevance of directed technical change effects, I calibrate the CES version of the model to estimates from the empirical literature on directed technical change. I use the calibration to simulate the effects of tax reforms and to compute optimal taxes.

7.1 Calibration

The calibration proceeds as follows. First, I set the wage elasticity parameters γ^{CES} and ρ^{CES} (equivalently, σ and δ), the labor supply elasticity e (assuming isoelastic disutility of labor), and the initial tax function $\bar{T}$ (approximating the US income tax system in 2005) on the basis of existing empirical estimates. In the second step, I infer the exogenous technology parameter κ from the US earnings distribution in 2005.

Within-Technology Substitution Effects The within-technology substitution elasticity γ^{CES} and the directed technical change elasticity ρ^{CES} govern the response of relative wages to changes in relative labor inputs. Directed technical change effects are likely to arise with considerable delay, implying that to measure ρ^{CES} , one has to track relative wages over a long period of time after an exogenous change in labor inputs occurred. Within-technology substitution, in contrast, does not require firms to change their production technology, so its effects are likely to occur over a much shorter period of time. The timing of the effects therefore provides an opportunity to identify γ^{CES} and ρ^{CES} separately.

The empirical literature that aims to identify an elasticity of substitution between differentially skilled worker groups without explicit reference to directed technical change typically focuses on comparably short time periods of about one year or slightly more. I take these estimates to set γ^{CES} .

Besides the timing of the effects, an important property in which many empirical studies differ is the definition of the skill groups between which an elasticity of substitution is measured. Many studies focus on college graduates versus those without a college degree. Others consider high school graduates versus high school dropouts. [Dustmann, Frattini and Preston \(2013\)](#) stand out in that they estimate substitution elasticities between workers located at 20 different points in the wage distribution. They test for heterogeneity in these elasticities but find no evidence for it. In light of this result, the CES assumption, which imposes a single elasticity of substitution between any two disjoint groups of workers, seems an acceptable simplification. It implies that all estimates, irrespective of the definition of skill groups, are equally relevant for the calibration of γ^{CES} .

[Acemoglu \(2002\)](#) summarizes the consensus of the literature at that time as σ being somewhere between 1.4 and 2, which implies that γ^{CES} falls between -0.5 and -0.7 . The results of [Carneiro et al. \(2019\)](#) imply a short-run elasticity, measured within two years after a skill supply shock, of -0.5 (for a detailed description of [Carneiro et al., 2019](#), see below). This value falls within the consensus range observed by [Acemoglu \(2002\)](#). Moreover, [Carneiro et al. \(2019\)](#) is the only study that estimates

Study	Skill Groups			Time Horizon	Geographical Level		Cross-wage Effect
Carneiro et al. (2019)	College	vs.	non-college	2 years	Norwegian	municipalities	−0.55
Carneiro et al. (2019)	College	vs.	non-college	11 years	Norwegian	municipalities	0
Carneiro et al. (2019)	College	vs.	non-college	17 years	Norwegian	municipalities	0.5
Lewis (2011)	High school	vs.	high-school dropout	10 years	US metro areas		−0.14
Dustmann and Glitz (2015)	Postsecondary	vs.	vocational degree or apprenticeship versus no postsecondary education	10 years	German	local labor markets (aggregates of German counties)	−0.09
Morrow and Trefler (2017)	Some tertiary	versus	no tertiary education	Short (see description in Appendix B.1)	38 countries		−0.53
Morrow and Trefler (2017)	Some tertiary	versus	no tertiary education	Long (see description in Appendix B.1)	38 countries		−0.11

Table 1. The table shows estimates of the effect of relative skill supply changes on relative wages from a set of empirical studies. A brief outline of each study with an explanation of how the numbers in the last column are derived from the respective study’s results is provided in Appendix B.1.

wage responses at different points in time. Thereby, it provides estimates of γ^{CES} and ρ^{CES} obtained consistently within a single framework. For these reasons, I set $\gamma^{CES} = -0.5$, the estimate implied by [Carneiro et al. \(2019\)](#). The implied elasticity of substitution is $\sigma = 2$.

Directed Technical Change Effects A few studies measure the response of wages to skill supply shocks over substantially longer periods of time (about 10 years or more). Most of them explicitly reference directed technical change and provide evidence for technology adjustments being an important driver of the long-run wage responses. Since this applies only to a handful of papers, I give a brief overview over each of them in Appendix B.1.

Table 1 shows the results of these papers. The short-run estimates are -0.55 and -0.53 , which (further) motivates my choice of γ^{CES} . Estimates over a period of about 10 years are consistently close to zero, ranging from -0.1 to 0 . Finally, the estimate from [Carneiro et al. \(2019\)](#) for an adjustment period of 17 years shows an effect of 0.5 . These long-run effects are total effects, in the sense that they include both within-technology and between-technology (directed technical change) substitution. Hence, they map into the sum of γ^{CES} and ρ^{CES} .

Based on Table 1, I consider two cases. The first case, derived from the 10 year estimates, sets $\gamma^{CES} + \rho^{CES}$ to -0.1 , which, given $\gamma^{CES} = -0.5$, implies $\rho^{CES} = 0.4$. In this case, within-technology substitution and directed technical change effects are of a similar magnitude and almost offset each other (given that they work in opposite directions). In the second case, based on the 17 year estimate of [Carneiro et al. \(2019\)](#), I set $\gamma^{CES} + \rho^{CES}$ to 0.5 , such that $\rho^{CES} = 1$. In this case, directed technical change

dominates within-technology substitution, that is, there is strong relative bias. I call this the strong bias case.

The conservative case is supported by all four studies in Table 1. Moreover, there are at least two further papers that, for different reasons, do not provide estimates that could be used to infer γ^{CES} and ρ^{CES} , but nevertheless support the view of the conservative case that the long-run wage effects of skill supply shocks are close to zero. First, [Blundell, Green and Jin \(2020\)](#) document that a large and sudden increase in the share of individuals holding a college degree in the 1990s in the UK left the wage premium associated with college education basically unchanged. They provide empirical results suggesting that firms responded to the hike in the relative supply of college graduates by adopting production forms that granted higher degrees of autonomy and responsibility to their workers, which likely benefited highly qualified workers' productivity. They argue that these endogenous technology adjustments offset the negative within-technology substitution effect on the college premium. Second, [Clemens, Lewis and Postel \(2018\)](#) study the effect of the exclusion of half a million Mexican farm workers (braceros) from the US in 1965 on US farm workers' wages and find no evidence for differential wage changes following the event in states heavily exposed to the bracero exclusion relative to less exposed states. They provide striking evidence for rapid adoption of labor-replacing technologies on farms in heavily exposed states after the exclusion.

The strong bias case is supported directly only by [Carneiro et al. \(2019\)](#). Nevertheless, I believe that the case for strong bias is stronger than it might appear from this. The studies in Table 1 analyze the differential evolution of wages between often quite narrowly defined geographical areas, which were hit differentially by plausibly exogenous skill supply shocks. By construction, such estimates miss all directed technical change effects that appear on a higher geographical level. Since the relevant markets for innovative technologies are plausibly much larger than most of the geographical units listed in Table 1, the estimates are likely to capture mostly the effects of investments into adoption of already existing technologies, rather than the effects of re-directed inventive activity.³⁷ The model developed in Section 3 implies that endogenous adoption and innovation work in the same direction, cumulating in the total directed technical change effect that is represented by ρ^{CES} . Hence, the estimates of Table 1 likely miss part of ρ^{CES} and therefore underestimate it.³⁸

Another piece of evidence in favor of strong bias is provided by [Fadinger and Mayr \(2014\)](#). They show that in a cross section of countries, relative skill supply measures are negatively correlated with relative unemployment rates of more versus less skilled workers and with relative emigration rates of skilled workers. In a directed technical change model with frictional labor markets and endogenous migration, they show that both correlations can be interpreted as signs of strong relative bias.³⁹

³⁷See, for example, [Dechezlepretre, Hemous, Olsen and Zanella \(2019\)](#) and [San \(2019\)](#) for patent-based evidence that inventions respond to the structure of labor supply.

³⁸A potential source of upwards bias in directed technical change effects obtained by comparing small geographical units are Rybczynski effects: a rise in relative skill supply in one region increases the region's exports of skill-intensive goods, which raises skilled workers' wages. This wage increase may be mistakenly attributed to directed technical change. All the studies listed in Table 1, however, provide different forms of evidence suggesting that adjustments in the output mix of their observation units are not driving their results. See the respective papers for details.

³⁹Note at this point that there are at least two empirical studies that are in more or less open contradiction to the predictions of directed technical change theory. First, [Blum \(2010\)](#) finds that in a panel of countries, increases in the relative supply

Labor Supply Elasticity I assume an isoelastic disutility of labor of the form

$$v(l) = \frac{e}{e+1} l^{\frac{e+1}{e}}.$$

With this specification, the parameter e is equal to the labor supply elasticity for a linearized tax function, e_θ , introduced in Section 6.1. Chetty (2012) estimates a value of 0.33 for this elasticity. The survey by Saez, Slemrod and Giertz (2012), on the other hand, reports a preferred estimate of 0.25 for the elasticity of taxable income. I show in Appendix B.2 that the model matches an elasticity of taxable income of 0.25 if the elasticity e is set to 0.67. The range between 0.33 and 0.67 contains many other estimates from the literature (e.g. Meghir and Phillips, 2010). I choose an intermediate value of 0.5 for my baseline calibration.⁴⁰ The main insights are robust to choosing other values within this range.

Initial Tax System I set the initial tax system, denoted by $\bar{T}$, as an approximation to the US income tax in 2005. I follow Heathcote et al. (2017), who show that a schedule with a constant rate of progressivity provides a good approximation:

$$T(y) = y - \lambda y^{1-P}.$$

Heathcote et al. (2017) estimate the parameters of such a tax function on 2000 to 2005 income and tax data for the US and obtain values of $p = 0.181$ and $\lambda = 5.568$. I use these values in all simulations.

Exogenous Technology With the parameters γ^{CES} , ρ^{CES} , e , and $\bar{T}$ calibrated, the exogenous technology parameter κ is identified by the earnings distribution under the initial tax system $\bar{T}$. I approximate the earnings distribution by smoothly combining a lognormal distribution for incomes below \$200k and a Pareto distribution with tail parameter 1.5 above \$200k (Diamond and Saez, 2011). Moreover, I assume that the type distribution h is standard uniform on $[\underline{\theta}, \bar{\theta}] = [0, 1]$. In the CES case, this assumption is insubstantial, because the cross-type wage elasticity between any two distinct types of workers is independent of the types' locations in the type space. Given an estimate of the income distribution, it is straightforward to compute the function κ from workers' first-order condition (2) and the wage equation (15). The procedure is described in more detail in Appendix B.3.

Welfare Function Finally, I use a welfare function of the type

$$V(\{u_\theta\}_{\theta \in \Theta}) = \left(\int_{\underline{\theta}}^{\bar{\theta}} u_\theta^{1-r} h_\theta d\theta \right)^{\frac{1}{1-r}},$$

of skilled workers reduce their relative wages by more in the long-run than in the short-run. Second, Ciccone and Peri (2005) report long-run estimates for the elasticity of substitution between college graduates and non-college workers of about 1.5, which maps into a total wage elasticity $\gamma^{CES} + \rho^{CES}$ of -0.7 . With $\gamma^{CES} = -0.5$, this implies a negative ρ^{CES} , inconsistent with theory. These results should serve as a word of caution regarding the simulation results below. Yet, I do not respect them directly in the simulations. After all, calibrating a model to empirical results that contradict the qualitative predictions of the model makes no sense.

⁴⁰The same value is used, for example, by Mankiw, Weinzierl and Yagan (2009) and Heathcote et al. (2017).

where the relative inequality aversion parameter r allows to vary the strength of the preference for equity in a flexible way (Atkinson, 1970).

For the baseline calibration, I set $r \rightarrow 1$, such that the (income-weighted) average of optimal marginal tax rates in the conservative case (i.e., for $\rho^{CES} = 0.4$) is the same as in the US 2005 tax and transfer system.⁴¹

7.2 Simulation

Given the calibrated CES version of the model, I simulate the effect of tax reforms on wages and compute optimal taxes on the basis of the analytical results of Sections 5 and 6.

Tax Reforms I study a hypothetical tax reform that reverses the cumulative impact on tax progressivity of US income tax reforms from 1970 to 2005. As documented by Piketty and Saez (2007), the US income tax system underwent a series of regressive reforms in this period. Heathcote et al. (2017) estimate the decline in tax progressivity between 1970 and 2005 to be 0.034 when measured by the progressivity parameter of a constant-rate-of-progressivity tax schedule. Taking as a starting value the progressivity estimate of $p = 0.181$, I hence ask: what are the effects of raising the progressivity of a constant-rate-of-progressivity tax from 0.181 (its 2005 US value) to 0.215 (its 1970 US value) on the wage distribution?⁴²

In the case of strong bias, such a progressive reform will reduce wage inequality. A particular focus of the analysis will thus be whether taking back the regressive tax reforms of the past decades would result in a meaningful reduction in US wage inequality under the strong bias calibration described above. Reversely, the results are informative about whether regressive tax reforms played a role, via directed technical change effects, in the rise of US wage inequality observed over the period under consideration.

I compute the wage effects of the described reform using the expression provided by Corollary 1. For the exogenous technology planner, who ignores directed technical change, the effects are obtained by setting the directed technical change elasticity ρ^{CES} to zero.⁴³

The results are displayed in Figure 1. In the conservative case, the reform has almost no effect on wages. This was expected, because directed technical change and within-technology substitution effects approximately offset each other in this case. When ignoring directed technical change, the model predicts moderate wage decreases for low-skilled workers and even smaller gains for the high-skilled. In the strong bias case, wages for low-skilled workers rise by up to 3% while wages for high-skilled workers decrease by up to 1.5%.

To put the results into perspective, I compute the effect on the 90-10-percentile ratio of the wage distribution. In the conservative case, this ratio increases by 0.3%, whereas in the strong bias case the

⁴¹I report results for $r = 50$ (close to Rawlsian) in Appendix C.3.

⁴²I choose the post-reform value for the parameter λ (the second parameter of a constant-rate-of-progressivity tax function) such that, in the conservative case described above ($\rho^{CES} = 0.4$), the reform leaves tax revenue unchanged.

⁴³Corollary 1 provides a local approximation of the wage effects of tax reforms. I also compute the exact changes in wages due to the reform and find that the difference to the local approximation is negligible.




Figure 1. The figure displays the total wage changes in log points induced by the progressive tax reform described in the text. Wage changes are shown for workers at each percentile of the wage distribution. The red dashed line and the blue line are for the strong bias and, respectively, the conservative case of the baseline calibration described in the text. The black line indicates wage changes as predicted by the exogenous technology planner.

ratio falls by 2%. This reduction is almost exclusively driven by an increase in the wage at the 10th percentile of about 1.9%, while the wage at the 90th percentile is basically unchanged. With a labor supply elasticity of $e = 0.67$ (at the high end of the range discussed in Section 7.1), the reduction in the 90-10-percentile ratio becomes 2.8%.

These effects can be compared to the actually observed changes in the US wage distribution between 1970 and 2005. In this period, the 90-10-percentile ratio rose by about 30% (e.g. [Acemoglu and Autor, 2011](#)). My results then suggest that regressive tax reforms, in conjunction with directed technical change, can explain up to 9% of the total increase in the 90-10 ratio (2.8% out of a total of 30%).

Note at this point that the model likely underestimates the directed technical change effects of regressive reforms by omitting potentially relevant adjustment margins. In particular, one might expect that, in the long run, individuals' education and occupation choices respond to less progressivity in the tax schedule in a way that reinforces the labor supply responses analyzed here. Quantifying these effects in a richer model is an interesting step for future research but outside the scope of the present paper. Instead, the present paper focuses on the design of optimal non-linear taxes without ad hoc constraints on their shape, which requires a comparably simple setup.⁴⁴

⁴⁴Forces that amplify the labor supply responses to taxes reduce the level of optimal marginal tax rates, but a priori, there is little reason to suspect that they affect the way in which directed technical change modifies the slope of the optimal marginal tax schedule. Hence, the omission of the above mentioned elements does not seem critical for the main insights of the optimal tax analysis.




Figure 2. The figure displays optimal marginal tax rates by income level. The red dashed line and the blue line are for the strong bias and, respectively, the conservative case of the baseline calibration described in the text. The black lines are for the exogenous technology planner and for the case with a fully exogenous wage distribution. The vertical line indicates the US median earnings level in 2005 of about \$52k.

Optimal Taxes I compute optimal marginal tax rates according to Proposition 3 and, for comparison, the marginal tax rates preferred by the exogenous technology planner as given by equation (37). As a further benchmark, I include the marginal tax rates that are preferred when the entire wage distribution is perceived as exogenous and fixed at its state under the US 2005 tax system.

Figure 2 shows that, as predicted by theory, directed technical change effects reduce optimal marginal tax rates in the lower part of the income distribution and increase them in the upper part. The point where directed technical change effects reverse their sign is close to the US median earnings level in 2005, indicated by the vertical line. Below the median income, optimal marginal tax rates are reduced by up to 18 percentage points relative to the exogenous technology planner’s tax rates. Above the median, the increase in optimal marginal tax rates due to directed technical change becomes as large as 10 percentage points.

These changes occur over the bulk of the income distribution. At the 10th percentile, the optimal marginal tax falls by 17 percentage points (5 percentage points) in the strong bias (conservative) case relative to the exogenous technology planner’s tax. At the 90th percentile, the increase in the optimal marginal tax is 8 percentage points (3 percentage points) in the strong bias (conservative) case. For strong bias, this leads marginal tax rates to increase monotonically with income over large parts of the income distribution, whereas they follow a pronounced U-shape when ignoring directed technical change.⁴⁵

⁴⁵At the very bottom of the income distribution (below the 10th percentile), optimal marginal tax rates are high in all cases.

Compared to the exogenous wage benchmark studied extensively in the existing literature, optimal marginal tax rates can be more or less progressive, depending on whether there is strong relative bias. In the conservative case, directed technical change and within-technology substitution effects almost exactly offset each other, such that the optimal tax is very close to that obtained with exogenous wages. In the strong bias case, directed technical change dominates and the optimal tax is substantially more progressive.

The welfare gains from implementing the optimal tax relative to the exogenous technology planner's preferred tax can be sizable. In the strong bias case, they are equivalent to an increase in the lump-sum payment of \$430 annually, corresponding to 1.5% of the lump-sum payment or 0.6% of average income under the exogenous technology planner's tax. In the conservative case, the gains are much smaller, amounting to \$35 annually. On the other end of the spectrum, they become as large as \$850 annually under strong bias and with Rawlsian welfare (see Appendix C.3).

Remarkably, in the strong bias case, the welfare level achieved by the actual US 2005 tax and transfer scheme is between the optimum and the level achieved by the exogenous technology planner's tax. Hence, for reasonable parameter values, ignoring directed technical change can lead to tax reform proposals that deteriorate welfare in the present environment. This is due to the fact that, under strong bias, the optimal tax with monotonically increasing marginal tax rates is much closer in shape to the actual tax schedule than to the U-shaped tax preferred when ignoring directed technical change.

8 Conclusion

I develop a model with directed technical change and endogenous labor supply, in which the structure of labor supply determines the direction of technical change. Tax reforms affect the direction of technical change by altering the structure of labor supply.

Under certain conditions, any progressive income tax reform induces technical change that compresses the pre-tax wage distribution. As a consequence, when directed technical change is taken into account – as opposed to treating technology as exogenous – optimal marginal tax rates are higher in the upper tail and lower in the lower tail of the income distribution.

Simulating the cumulative (regressive) reforms of the US tax and transfer system between 1970 and 2005, the resulting wage effects, accounting for the induced technical change, are modest. In the calibration most favorable to directed technical change effects, the regressive tax reforms explain about 9% of the total contemporaneous increase in US wage inequality as measured by the 90-10-percentile ratio of the wage distribution. The impact of directed technical change on optimal marginal tax rates, however, is large. Optimal marginal tax rates for workers who earn about half of the 2005 US median income are reduced by more than 10 percentage points relative to the benchmark where technology is treated as exogenous.

This is because, for comparability with other quantitative optimal tax studies (e.g. [Mankiw et al., 2009](#); [Brewer, Saez and Shepard, 2010](#)), I assume that the income distribution has a mass point at zero. Consequently, the negative marginal tax results obtained from Proposition 3, which require strictly positive incomes for the least skilled workers, do not apply.

Future work may elaborate on the positive predictions of the theory. For a more complete understanding of the contribution of past tax reforms to changes in wage inequality, a richer model would be useful, including adjustments in occupational choices and education decisions to taxes.

References

- Acemoglu, Daron (1998) "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality," *Quarterly Journal of Economics*, **113** (4), 1055–1089.
- (2002) "Directed Technical Change," *Review of Economic Studies*, **69** (4), 781–809.
- (2007) "Equilibrium Bias of Technology," *Econometrica*, **75** (5), 1371–1409.
- (2010) "When Does Labor Scarcity Encourage Innovation?" *Journal of Political Economy*, **118** (6), 1037–1078.
- Acemoglu, Daron and David H. Autor (2011) "Skills, Tasks and Technologies: Implications for Employment and Earnings," *Handbook of Labor Economics*, **4**, 1043–1171.
- Acemoglu, Daron, Andrea Manera, and Pascual Restrepo (2020) "Does the US Tax Code Favor Automation?," NBER Working Paper 27052.
- Acemoglu, Daron and Pascual Restrepo (2018) "The Race Between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment," *American Economic Review*, **108** (6), 1488–1542.
- Akcigit, Ufuk, John R. Grigsby, Tom Nicholas, and Stefanie Stantcheva (2019a) "Taxation and Innovation in the 20th Century," NBER Working Paper 24982.
- Akcigit, Ufuk, Douglas Hanley, and Stefanie Stantcheva (2019b) "Optimal Taxation and R&D Policies," NBER Working Paper 22908.
- Akcigit, Ufuk and Stefanie Stantcheva (2020) "Taxation and Innovation: What Do We Know?" in Austan Goolsbee and Benjamin Jones eds. *Innovation and Public Policy*: University of Chicago Press.
- Ales, Laurence, Musab Kurnaz, and Christopher Sleet (2015) "Technical Change, Wage Inequality, and Taxes," *American Economic Review*, **105** (10), 3061–3101.
- Atkinson, Anthony B. (1970) "On the Measurement of Inequality," *Journal of Economic Theory*, **2**, 244–263.
- Blum, Bernardo S. (2010) "Endowments, Output, and the Bias of Directed Innovation," *Review of Economic Studies*, **77** (2), 534–559.
- Blundell, Richard, David Green, and Wenchao Jin (2020) "The UK as a Technological Follower: Higher Education Expansion, Technology Adoption, and the Labour Market," Working Paper.
- Brewer, Mike, Emmanuel Saez, and Andrew Shepard (2010) "Means-testing and Tax Rates on Earnings," in Stuart Adam, Tim Besley, Richard Blundell, Stephen Bond, Robert Chote, Malcolm Gammie, Paul Johnson, Gareth Myles, and James M. Poterba eds. *Dimensions of Tax Design: The Mirrlees Review*: Oxford University Press, 90–201.

- Carneiro, Pedro, Kai Liu, and Kjell G. Salvanes (2019) "The Supply of Skill and Endogenous Technical Change: Evidence From a College Expansion Reform," Working Paper.
- Chetty, Raj (2012) "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply," *Econometrica*, **80** (3), 96–1018.
- Ciccone, Antonio and Giovanni Peri (2005) "Long-run Substitutability Between More and Less Educated Workers: Evidence from US States 1950-1990," *Review of Economics and Statistics*, **87** (4), 652–663.
- Clemens, Michael A., Ethan G. Lewis, and Hannah M. Postel (2018) "Immigration Restrictions as Active Labor Market Policy: Evidence from the Mexican Bracero Exclusion," *American Economic Review*, **108** (6), 1468–1487.
- Costinot, Arnaud and Jonathan Vogel (2010) "Matching and Inequality in the World Economy," *Journal of Political Economy*, **118** (4), 747–786.
- Costinot, Arnaud and Ivan Werning (2020) "Robots, Trade, and Luddism: A Sufficient Statistic Approach to Optimal Technology Regulation," Working Paper.
- Dechezlepretre, Antoine, David Hemous, Morten Olsen, and Carlo Zanella (2019) "Automating Labor: Evidence from Firm-level Patent Data," Working Paper.
- Diamond, Peter A. (1998) "Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates," *American Economic Review*, **88** (1), 83–95.
- Diamond, Peter and Emmanuel Saez (2011) "The Case for a Progressive Tax: From Basic Research to Policy Recommendations," *Journal of Economic Perspectives*, **25** (4), 165–190.
- Dustmann, Christian, Tommaso Frattini, and Ian P. Preston (2013) "The Effect of Immigration along the Distribution of Wages," *Review of Economic Studies*, **80** (1), 145–173.
- Dustmann, Christian and Albrecht Glitz (2015) "How Do Industries and Firms Respond to Changes in Local Labor Supply?" *Journal of Labor Economics*, **33** (3), 711–750.
- Eaton, Jonathan and Samuel Kortum (2002) "Technology, Geography, and Trade," *Econometrica*, **70** (5), 1741–1779.
- Fadinger, Harald and Karin Mayr (2014) "Skill-Biased Technological Change, Unemployment and Brain Drain," *Journal of the European Economic Association*, **12** (2), 397–431.
- Goldin, Claudia and Lawrence F. Katz (2008) *The Race Between Education and Technology*: Harvard University Press.
- Guerreiro, Joao, Sergio Rebelo, and Pedro Teles (2018) "Should Robots be Taxed?," NBER Working Paper 23806.
- Heathcote, Jonathan, Kjetil Storesletten, and Gianluca Violante (2017) "Optimal Tax Progressivity: An Analytical Framework," *Quarterly Journal of Economics*, **132** (4), 1693–1754.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2020) "How Should Tax Progressivity Respond to Rising Income Inequality?" *Journal of the European Economic Association*, forthcoming.
- Jacobs, Bas and Uwe Thuemmel (2018a) "Optimal Linear Income Taxation and Education Subsidies under Skill-Biased Technical Change," Working Paper.

- (2018b) “Optimal Taxation of Income and Human Capital and Skill-Biased Technical Change,” Working Paper.
- Jagadeesan, Ravi (2019) “Optimal Taxation with an Endogenous Growth Rate,” Working Paper.
- Jones, Charles I. (2019) “Taxing Top Incomes in a World of Ideas,” Working Paper.
- Kiley, Michael T. (1999) “The Supply of Skilled Labour and Skill-Biased Technological Progress,” *Economic Journal*, **109** (458), 708–724.
- Lewis, Ethan (2011) “Immigration, Skill Mix, and Capital Skill Complementarity,” *Quarterly Journal of Economics*, **126** (2), 1029–1069.
- Lockwood, Benjamin B., Charles G. Nathanson, and E. Glen Weyl (2017) “Taxation and the Allocation of Talent,” *Journal of Political Economy*, **125** (5), 1635–1682.
- Loebbing, Jonas (2016) “A LeChatelier Principle for Relative Demand and Implications for Directed Technical Change,” Working Paper.
- (2020) “An Elementary Theory of Directed Technical Change and Wage Inequality,” Working Paper.
- Mankiw, N. Gregory, Matthew Weinzierl, and Danny Yagan (2009) “Optimal Taxation in Theory and Practice,” *Journal of Economic Perspectives*, **23** (4), 147–74.
- Meghir, Costas and David Phillips (2010) “Labour Supply and Taxes,” in Stuart Adam, Tim Besley, Richard Blundell, Stephen Bond, Robert Chote, Malcolm Gammie, Paul Johnson, Gareth Myles, and James M. Poterba eds. *Dimensions of Tax Design: The Mirrlees Review*: Oxford University Press, 202–274.
- Milgrom, Paul and Chris Shannon (1994) “Monotone Comparative Statics,” *Econometrica*, **62** (1), 157–180.
- Mirrlees, James A. (1971) “An Exploration in the Theory of Optimum Income Taxation,” *Review of Economic Studies*, **38** (2), 175–208.
- Morrow, Peter M. and Daniel Trefler (2017) “Endowments, Skill-Biased Technology, and Factor Prices: A Unified Approach to Trade,” NBER Working Paper 24078.
- Naito, Hisahiro (1999) “Re-Examination of Uniform Commodity Taxes Under a Non-Linear Income Tax System and Its Implications for Production Efficiency,” *Journal of Public Economics*, **71** (2), 165–188.
- Piketty, Thomas and Emmanuel Saez (2007) “How Progressive is the US Federal Tax System? A Historical and International Perspective,” *Journal of Economic Perspectives*, **21** (1), 3–24.
- Romer, Paul M. (1990) “Endogenous Technological Change,” *Journal of Political Economy*, **98** (5), 71–102.
- (1994) “The Origins of Endogenous Growth,” *Journal of Economic Perspectives*, **8** (1), 3–22.
- Rothschild, Casey and Florian Scheuer (2013) “Redistributive Taxation in the Roy Model,” *Quarterly Journal of Economics*, **128**, 623–668.
- (2016) “Optimal Taxation with Rent-Seeking,” *Review of Economic Studies*, **83** (3), 1225–1262.

- Sachs, Dominik, Aleh Tsyvinski, and Nicolas Werquin (2020) “Nonlinear Tax Incidence and Optimal Taxation in General Equilibrium,” *Econometrica*, **88** (2), 469–493.
- Saez, Emmanuel (2001) “Using Elasticities to Derive Optimal Income Tax Rates,” *Review of Economic Studies*, **68** (1), 205–229.
- Saez, Emmanuel, Joel Slemrod, and Seth H. Giertz (2012) “The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review,” *Journal of Economic Perspectives*, **50** (1), 3–50.
- San, Shmuel (2019) “Labor Supply and Directed Technical Change: Evidence from the Abrogation of the Bracero Program in 1964,” Working Paper.
- Slavik, Ctirad and Hakki Yazici (2014) “Machines, Buildings, and Optimal Dynamic Taxes,” *Journal of Monetary Economics*, **66**, 47–61.
- Stiglitz, Joseph E. (1982) “Self-Selection and Pareto Efficient Taxation,” *Journal of Public Economics*, **17** (2), 213–240.
- Thuemmel, Uwe (2018) “Optimal Taxation of Robots,” Working Paper.

A Proofs and Derivations

This appendix contains all proofs and derivations omitted from the main text.

A.1 Gateaux Derivatives

Here, I provide a rigorous definition of the Gateaux derivative D_{l_θ} in direction of the labor input of a given type θ . The derivative D_{u_θ} used in Section 6 is defined analogously.

Let $x : (l, z) \mapsto x(l, z)$ be a function of the (infinite-dimensional) labor input vector l and, potentially, further variables z . The Gateaux derivative of x in direction of l_θ , for any given type θ , is defined as

$$D_{l_\theta} x(l, z) := \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \frac{dx(l + \mu \tilde{l}_{\Delta, \theta}, z)}{d\mu} \Big|_{\mu=0},$$

where $\tilde{l}_{\Delta, \theta} : \tilde{\theta} \mapsto \tilde{l}_{\Delta, \theta, \tilde{\theta}}$, with $\Delta \in \mathbb{R}_{++}$, is a real-valued function on the type space defined as follows. For interior types $\theta \in (\underline{\theta}, \bar{\theta})$, it is given by

$$\tilde{l}_{\Delta, \theta, \tilde{\theta}} = \begin{cases} 0 & \text{for } \tilde{\theta} < \theta - \Delta \\ \frac{\tilde{\theta} - \theta + \Delta}{\Delta} & \text{for } \tilde{\theta} \in [\theta - \Delta, \theta] \\ \frac{\theta - \tilde{\theta} + \Delta}{\Delta} & \text{for } \tilde{\theta} \in [\theta, \theta + \Delta] \\ 0 & \text{for } \tilde{\theta} > \theta + \Delta. \end{cases}$$

For the highest type $\bar{\theta}$, it is given by

$$\tilde{l}_{\Delta, \bar{\theta}, \tilde{\theta}} = \begin{cases} 0 & \text{for } \tilde{\theta} < \bar{\theta} - \Delta \\ \frac{2(\tilde{\theta} - \bar{\theta} + \Delta)}{\Delta} & \text{for } \tilde{\theta} \in [\bar{\theta} - \Delta, \bar{\theta}] ; \end{cases}$$

and for the lowest type $\underline{\theta}$, it becomes

$$\tilde{l}_{\Delta, \underline{\theta}, \tilde{\theta}} = \begin{cases} \frac{2(\underline{\theta} - \tilde{\theta} + \Delta)}{\Delta} & \text{for } \tilde{\theta} \in [\underline{\theta}, \underline{\theta} + \Delta] \\ 0 & \text{for } \tilde{\theta} > \underline{\theta} + \Delta . \end{cases}$$

Intuitively, the derivative is obtained by perturbing the labor supply function continuously in a neighborhood of type θ and letting this neighborhood converge to θ .

To demonstrate that the thus defined derivative works as expected, I now derive the labor demand equation (3) in detail. Note to this end that the definition of D_{l_θ} extends exactly to the derivative with respect to aggregate labor inputs, D_{L_θ} . We start from final good firms' profits, given by

$$\tilde{G}(L, \phi, q) = \int_{\underline{\theta}}^{\bar{\theta}} w_\theta L_\theta d\theta - \sum_{j=1}^J \int_0^1 p_{j,k} q_{j,k} dk .$$

Taking the derivative D_{L_θ} and equating it with zero yields

$$D_{L_\theta} \tilde{G}(L, \phi, q) = D_{L_\theta} \int_{\underline{\theta}}^{\bar{\theta}} w_{\tilde{\theta}} L_{\tilde{\theta}} d\tilde{\theta} .$$

The remaining task is to show that

$$D_{L_\theta} \int_{\underline{\theta}}^{\bar{\theta}} w_{\tilde{\theta}} L_{\tilde{\theta}} d\tilde{\theta} = w_\theta .$$

I derive this equality for interior types $\theta \in (\underline{\theta}, \bar{\theta})$ only. For the boundary types $\underline{\theta}$ and $\bar{\theta}$, the derivation is analogous.

By definition of D_{L_θ} , we obtain:

$$D_{L_\theta} \int_{\underline{\theta}}^{\bar{\theta}} w_{\tilde{\theta}} L_{\tilde{\theta}} d\tilde{\theta} = \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\bar{\theta}} \frac{d}{d\mu} w_{\tilde{\theta}} \left(L_{\tilde{\theta}} + \mu \tilde{L}_{\Delta, \theta, \tilde{\theta}} \right) \Big|_{\mu=0} d\tilde{\theta} .$$

Moreover, by definition of $\tilde{L}_{\Delta, \theta}$ (defined analogously to $\tilde{l}_{\Delta, \theta}$):

$$\begin{aligned} \int_{\underline{\theta}}^{\bar{\theta}} \frac{d}{d\mu} w_{\tilde{\theta}} \left(L_{\tilde{\theta}} + \mu \tilde{L}_{\Delta, \theta, \tilde{\theta}} \right) \Big|_{\mu=0} d\tilde{\theta} &= \int_{\theta-\Delta}^{\theta} \frac{d}{d\mu} w_{\tilde{\theta}} \left(L_{\tilde{\theta}} + \mu \frac{\tilde{\theta} - \theta + \Delta}{\Delta} \right) \Big|_{\mu=0} d\tilde{\theta} \\ &\quad + \int_{\theta}^{\theta+\Delta} \frac{d}{d\mu} w_{\tilde{\theta}} \left(L_{\tilde{\theta}} + \mu \frac{\theta - \tilde{\theta} + \Delta}{\Delta} \right) \Big|_{\mu=0} d\tilde{\theta} . \end{aligned}$$

Hence:

$$\int_{\underline{\theta}}^{\bar{\theta}} \frac{d}{d\mu} w_{\tilde{\theta}} \left(L_{\tilde{\theta}} + \mu \tilde{L}_{\Delta, \theta, \tilde{\theta}} \right) \Big|_{\mu=0} d\tilde{\theta} = \int_{\theta-\Delta}^{\theta} w_{\tilde{\theta}} \frac{\tilde{\theta} - \theta + \Delta}{\Delta} d\tilde{\theta} + \int_{\theta}^{\theta+\Delta} w_{\tilde{\theta}} \frac{\theta - \tilde{\theta} + \Delta}{\Delta} d\tilde{\theta} .$$

By L'Hôpital's rule, this implies:

$$\lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\bar{\theta}} \frac{d}{d\mu} w_{\tilde{\theta}} \left(L_{\tilde{\theta}} + \mu \tilde{L}_{\Delta, \theta, \tilde{\theta}} \right) \Big|_{\mu=0} d\tilde{\theta} = \lim_{\Delta \rightarrow 0} \frac{1}{2\Delta} \int_{\theta-\Delta}^{\theta} w_{\tilde{\theta}} d\tilde{\theta} + \lim_{\Delta \rightarrow 0} \frac{1}{2\Delta} \int_{\theta}^{\theta+\Delta} w_{\tilde{\theta}} d\tilde{\theta} .$$

Finally, applying L'Hôpital's rule again, we obtain:

$$\lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\bar{\theta}} \frac{d}{d\mu} w_{\tilde{\theta}} \left(L_{\tilde{\theta}} + \mu \tilde{L}_{\Delta, \theta, \tilde{\theta}} \right) \Big|_{\mu=0} d\tilde{\theta} = \lim_{\Delta \rightarrow 0} \frac{w_{\theta-\Delta}}{2} + \lim_{\Delta \rightarrow 0} \frac{w_{\theta+\Delta}}{2} = w_{\theta} ,$$

where the last equality requires that w is continuous in θ , which I assume is the case in equilibrium.

A.2 Labor Supply Elasticities

In this section, I first provide a rigorous definition of the elasticity of labor supply with respect to the marginal retention rate, ϵ_{θ}^R . Then, I show how the elasticities ϵ_{θ}^R and ϵ_{θ}^w , used in the tax reform analysis of Section 5, relate to the elasticity e_{θ} used in the optimal tax analysis of Section 6. These relations are used for the calibration of the model, as described in detail in Appendix B.2.

For an arbitrary type θ , let $l_{\theta}(T, w_{\theta})$ denote labor supply as a function of the tax T and the wage w_{θ} , as given by workers' first-order condition (2). The elasticity of labor supply with respect to the marginal retention rate is then defined as

$$\epsilon_{\theta}^R(T, l, w) := \frac{R'_T(w_{\theta} l_{\theta})}{l_{\theta}} \frac{dl_{\theta}(T + \mu \tilde{\tau}, w_{\theta})}{d\mu} \Big|_{\mu=0} ,$$

where the auxiliary tax reform $\tilde{\tau}$ is chosen such that, as the scaling factor μ of the reform goes to zero, it raises the marginal retention rate by one infinitesimal unit:

$$\forall y : \tilde{\tau}(y) = -y, \quad \text{and thus:} \quad (y - (T(y) + \mu \tilde{\tau}(y)))' = 1 - T'(y) + \mu .$$

Defined in this way, the elasticity ϵ_{θ}^R gives exactly the local response of individual labor supply to a one unit increase in the marginal retention rate.

To express the elasticities ϵ_{θ}^R and ϵ_{θ}^w in terms of primitives of the utility and the tax function – which will also clarify their relation to the elasticity e_{θ} from Section 6 – start from workers' first-order condition:

$$v'(l_{\theta}(T, w_{\theta})) = R'_T(w_{\theta} l_{\theta}(T, w_{\theta})) w_{\theta} .$$

Perturbing both sides of the equation in direction of the above defined tax reform $\tilde{\tau}$, we obtain:

$$v''(l_\theta(T, w_\theta)) \left. \frac{dl_\theta(T + \mu\tilde{\tau}, w_\theta)}{d\mu} \right|_{\mu=0} = w_\theta \left. \frac{d(1 - T'(w_\theta l_\theta(T, w_\theta)) + \mu)}{d\mu} \right|_{\mu=0} - T''(w_\theta l_\theta(T, w_\theta)) w_\theta^2 \left. \frac{dl_\theta(T + \mu\tilde{\tau}, w_\theta)}{d\mu} \right|_{\mu=0}$$

and hence:

$$\left. \frac{dl_\theta(T + \mu\tilde{\tau}, w_\theta)}{d\mu} \right|_{\mu=0} = \frac{w_\theta}{v''(l_\theta(T, w_\theta)) + T''(w_\theta l_\theta(T, w_\theta)) w_\theta^2}.$$

By definition of ϵ_θ^R we obtain

$$\epsilon_\theta^R = \frac{\frac{w_\theta(1 - T'(w_\theta l_\theta(T, w_\theta)))}{v''(l_\theta(T, w_\theta)) l_\theta(T, w_\theta)}}{1 + \frac{T''(w_\theta l_\theta(T, w_\theta)) w_\theta l_\theta(T, w_\theta)}{1 - T'(w_\theta l_\theta(T, w_\theta))} \frac{(1 - T'(w_\theta l_\theta(T, w_\theta))) w_\theta}{v''(l_\theta(T, w_\theta)) l_\theta(T, w_\theta)}}.$$

Now, use the first-order condition to replace $w(1 - T'(wl))$ by $v'(l)$ and the definition of the rate of progressivity of a tax schedule, $P_T(y)$, to obtain

$$\epsilon_\theta^R = \frac{\frac{v'(l_\theta(T, w_\theta))}{v''(l_\theta(T, w_\theta)) l_\theta(T, w_\theta)}}{1 + P_T(w_\theta l_\theta(T, w_\theta)) \frac{v'(l_\theta(T, w_\theta))}{v''(l_\theta(T, w_\theta)) l_\theta(T, w_\theta)}}.$$

Finally, with the definition of the elasticity e_θ , the following relation between the two elasticities arises:

$$\epsilon_\theta^R(T, l, w) = \frac{e_\theta(l)}{1 + e_\theta(l) P_T(w_\theta l_\theta)} . \quad (39)$$

The difference between the two elasticities is that ϵ_θ^R accounts for the non-linearity of the tax function, while e_θ applies along a linearized tax function. Consequently, if the tax function is locally linear, that is, $P_T(w_\theta, l_\theta) = 0$, the elasticities coincide.

Next, for the elasticity of labor supply with respect to the wage, differentiate workers' first-order condition with respect to w_θ on both sides,

$$v''(l_\theta(T, w_\theta)) \frac{\partial l_\theta(T, w_\theta)}{\partial w_\theta} = 1 - T'(w_\theta l_\theta(T, w_\theta)) - T''(w_\theta l_\theta(T, w_\theta)) w_\theta^2 \frac{\partial l_\theta(T, w_\theta)}{\partial w_\theta} - T''(w_\theta, l_\theta(T, w_\theta)) w_\theta l_\theta(T, w_\theta) ,$$

and rearrange it to obtain

$$\frac{\partial l_\theta(T, w_\theta)}{\partial w_\theta} = \frac{1 - T'(w_\theta l_\theta(T, w_\theta)) - T''(w_\theta l_\theta(T, w_\theta))}{v''(l_\theta(T, w_\theta)) + T''(w_\theta l_\theta(T, w_\theta)) w_\theta^2}.$$

Then, use the definition of ϵ_θ^w to get

$$\epsilon_\theta^w = \frac{\left(1 - \frac{T''(w_\theta l_\theta(T, w_\theta))w_\theta l_\theta(T, w_\theta)}{1 - T'(w_\theta l_\theta(T, w_\theta))}\right) \frac{(1 - T'(w_\theta l_\theta(T, w_\theta)))w_\theta}{v''(l_\theta(T, w_\theta))l_\theta(T, w_\theta)}}{1 + \frac{T''(w_\theta l_\theta(T, w_\theta))w_\theta l_\theta(T, w_\theta)}{1 - T'(w_\theta l_\theta(T, w_\theta))} \frac{(1 - T'(w_\theta l_\theta(T, w_\theta)))w_\theta}{v''(l_\theta(T, w_\theta))l_\theta(T, w_\theta)}} .$$

Replacing $w(1 - T'(wl))$ by $v'(l)$ and using the definition of the rate of progressivity of T yields:

$$\epsilon_\theta^w = \frac{(1 - P_T(w_\theta l_\theta(T, w_\theta))) \frac{v'(l_\theta(T, w_\theta))}{v''(l_\theta(T, w_\theta))l_\theta(T, w_\theta)}}{1 + P_T(w_\theta l_\theta(T, w_\theta)) \frac{v'(l_\theta(T, w_\theta))}{v''(l_\theta(T, w_\theta))l_\theta(T, w_\theta)}} .$$

With the definition of e_θ , this implies:

$$\epsilon_\theta^w(T, l, w) = \frac{(1 - P_T(w_\theta l_\theta))e_\theta(l)}{1 + e_\theta(l)P_T(w_\theta l_\theta)} . \quad (40)$$

Note at this point that the second-order condition of workers' utility maximization is

$$v''(l_\theta) + T'(w_\theta l_\theta)w_\theta^2 \geq 0 .$$

At the utility maximum, this is equivalent to (using workers' first-order condition)

$$\frac{T''(w_\theta l_\theta)w_\theta l_\theta}{1 - T'(w_\theta l_\theta)} \frac{v'(l_\theta)}{v''(l_\theta)l_\theta} = P_T(w_\theta l_\theta)e_\theta(l_\theta) \geq -1 .$$

Assumption 1 in the main text ensures that this inequality is satisfied strictly. Hence, workers' second-order condition is satisfied strictly and the elasticities ϵ_θ^R and ϵ_θ^w are well defined.

A.3 Derivation of Aggregate Production, Wages, and Wage Elasticities in the CES Case

In this section, I derive the expressions for aggregate production, wages, and wage elasticities for the CES special case introduced in Section 3.6.

Aggregate Production To derive the aggregate production function $F(l, \phi)$ as given by equation (13) in the main text, start from the definition of F :

$$F(l, \phi) = \max_{\{q_\theta\}_{\theta \in \Theta}} \left\{ \tilde{G}(\{h_\theta l_\theta\}_{\theta \in \Theta}, \tilde{\phi}, \{q_\theta\}_{\theta \in \Theta}) - \int_{\underline{\theta}}^{\bar{\theta}} \eta_\theta q_\theta d\theta \right\} .$$

With the CES form of final good firms' production function $\tilde{G}$, the first-order conditions for the maximization with respect to q are:

$$\tilde{G}^{\frac{1}{\sigma}} \left(\tilde{\kappa}_\theta \tilde{\phi}_\theta h_\theta^{1-\alpha} l_\theta^{1-\alpha} \right)^{\frac{\bar{\sigma}-1}{\sigma}} \alpha q_\theta^{\frac{\alpha\bar{\sigma}-\alpha-\bar{\sigma}}{\sigma}} = \eta_\theta \quad \forall \theta ,$$

which can be rearranged to yield an explicit expression for the maximizer:

$$q_\theta = \left(\frac{\alpha}{\eta_\theta} \right)^{\frac{\tilde{\sigma}}{\alpha + \tilde{\sigma} - \alpha\tilde{\sigma}}} \left(\tilde{\kappa}_\theta \tilde{\phi}_\theta h_\theta^{1-\alpha} l_\theta^{1-\alpha} \right)^{\frac{\tilde{\sigma}-1}{\alpha + \tilde{\sigma} - \alpha\tilde{\sigma}}} \tilde{G}^{\frac{1}{\alpha + \tilde{\sigma} - \alpha\tilde{\sigma}}} \quad \forall \theta. \quad (41)$$

Denoting this maximizer by q^* and inserting it into $\tilde{G}$ yields

$$\tilde{G}(\{h_\theta l_\theta\}_{\theta \in \Theta}, \tilde{\phi}, q^*) = \left[\int_{\underline{\theta}}^{\bar{\theta}} \left(\frac{\alpha}{\eta_\theta} \right)^{\frac{\alpha(\tilde{\sigma}-1)}{\alpha + \tilde{\sigma} - \alpha\tilde{\sigma}}} \left(\tilde{\kappa}_\theta \tilde{\phi}_\theta h_\theta^{1-\alpha} l_\theta^{1-\alpha} \right)^{\frac{\tilde{\sigma}-1}{\alpha + \tilde{\sigma} - \alpha\tilde{\sigma}}} \tilde{G}^{\frac{\alpha(\tilde{\sigma}-1)}{(\alpha + \tilde{\sigma} - \alpha\tilde{\sigma})\tilde{\sigma}}} d\theta \right]^{\frac{\tilde{\sigma}}{\tilde{\sigma}-1}},$$

which can be solved for $\tilde{G}$:

$$\tilde{G}(\{h_\theta l_\theta\}_{\theta \in \Theta}, \tilde{\phi}, q^*) = \alpha^{\frac{\alpha}{1-\alpha}} \left[\int_{\underline{\theta}}^{\bar{\theta}} \left(\eta_\theta^{-\frac{\alpha}{1-\alpha}} \tilde{\kappa}_\theta^{\frac{1}{1-\alpha}} \tilde{\phi}_\theta^{\frac{1}{1-\alpha}} h_\theta l_\theta \right)^{\frac{(1-\alpha)\tilde{\sigma}}{\alpha + \tilde{\sigma} - \alpha\tilde{\sigma}} \frac{\tilde{\sigma}-1}{\tilde{\sigma}}} d\theta \right]^{\frac{\alpha + \tilde{\sigma} - \alpha\tilde{\sigma}}{(1-\alpha)\tilde{\sigma}} \frac{\tilde{\sigma}}{\tilde{\sigma}-1}}. \quad (42)$$

This provides an expression for gross aggregate production. Using the maximizer q^* from equation (41) again, the part of gross output that goes into the production of intermediate goods becomes

$$\begin{aligned} \int_{\underline{\theta}}^{\bar{\theta}} \eta_\theta q_\theta^* d\theta &= \int_{\underline{\theta}}^{\bar{\theta}} \eta_\theta^{\frac{\alpha - \alpha\tilde{\sigma}}{\alpha + \tilde{\sigma} - \alpha\tilde{\sigma}}} \alpha^{\frac{\tilde{\sigma}}{\alpha + \tilde{\sigma} - \alpha\tilde{\sigma}}} \left(\tilde{\kappa}_\theta \tilde{\phi}_\theta h_\theta^{1-\alpha} l_\theta^{1-\alpha} \right)^{\frac{\tilde{\sigma}-1}{\alpha + \tilde{\sigma} - \alpha\tilde{\sigma}}} \tilde{G}^{\frac{1}{\alpha + \tilde{\sigma} - \alpha\tilde{\sigma}}} d\theta \\ &= \alpha^{\frac{\tilde{\sigma}}{\alpha + \tilde{\sigma} - \alpha\tilde{\sigma}}} \tilde{G}^{\frac{1}{\alpha + \tilde{\sigma} - \alpha\tilde{\sigma}}} \int_{\underline{\theta}}^{\bar{\theta}} \left(\eta_\theta^{-\frac{\alpha}{1-\alpha}} \tilde{\kappa}_\theta^{\frac{1}{1-\alpha}} \tilde{\phi}_\theta^{\frac{1}{1-\alpha}} h_\theta l_\theta \right)^{\frac{(1-\alpha)\tilde{\sigma}}{\alpha + \tilde{\sigma} - \alpha\tilde{\sigma}} \frac{\tilde{\sigma}-1}{\tilde{\sigma}}} d\theta \\ &= \alpha \tilde{G}. \end{aligned} \quad (43)$$

Combining equations (42) and (43), we obtain the net aggregate production F as follows:

$$\begin{aligned} F(l, \phi) &= (1 - \alpha) \tilde{G}(\{h_\theta l_\theta\}_{\theta \in \Theta}, \tilde{\phi}, q^*) \\ &= \left[\int_{\underline{\theta}}^{\bar{\theta}} \left((1 - \alpha) \alpha^{\frac{\alpha}{1-\alpha}} \eta_\theta^{-\frac{\alpha}{1-\alpha}} \tilde{\kappa}_\theta^{\frac{1}{1-\alpha}} \tilde{\phi}_\theta^{\frac{1}{1-\alpha}} h_\theta l_\theta \right)^{\frac{(1-\alpha)\tilde{\sigma}}{\alpha + \tilde{\sigma} - \alpha\tilde{\sigma}} \frac{\tilde{\sigma}-1}{\tilde{\sigma}}} d\theta \right]^{\frac{\alpha + \tilde{\sigma} - \alpha\tilde{\sigma}}{(1-\alpha)\tilde{\sigma}} \frac{\tilde{\sigma}}{\tilde{\sigma}-1}}. \end{aligned}$$

Defining

$$\begin{aligned} \frac{\sigma - 1}{\sigma} &:= \frac{(1 - \alpha)\tilde{\sigma}}{\alpha + \tilde{\sigma} - \alpha\tilde{\sigma}} \frac{\tilde{\sigma} - 1}{\tilde{\sigma}} \\ \kappa_\theta &:= (1 - \alpha) \alpha^{\frac{\alpha}{1-\alpha}} \eta_\theta^{-\frac{\alpha}{1-\alpha}} \tilde{\kappa}_\theta^{\frac{1}{1-\alpha}} \quad \forall \theta \\ \phi_\theta &:= \tilde{\phi}_\theta^{\frac{1}{1-\alpha}} \quad \forall \theta, \end{aligned}$$

net aggregate production becomes

$$F(l, \phi) = \left[\int_{\underline{\theta}}^{\bar{\theta}} (\kappa_{\theta} \phi_{\theta} l_{\theta} h_{\theta})^{\frac{\sigma-1}{\sigma}} d\theta \right]^{\frac{\sigma}{\sigma-1}},$$

which is equation (13) from the main text.

Wages I derive expression (15) for interior types $\theta \in (\underline{\theta}, \bar{\theta})$. For the boundary types $\underline{\theta}$ and $\bar{\theta}$ the derivation proceeds analogously and yields the same result.

Consider the derivative $D_{l_{\theta}} F(l, \phi)$. Using the definition of $D_{l_{\theta}}$ given in Appendix A.1 and the CES form of F , this derivative becomes

$$\begin{aligned} D_{l_{\theta}} F(l, \phi) &= \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} F(l, \phi)^{\frac{1}{\sigma}} \left[\int_{\theta-\Delta}^{\theta} (\kappa_{\tilde{\theta}} \phi_{\tilde{\theta}} h_{\tilde{\theta}})^{\frac{\sigma-1}{\sigma}} l_{\tilde{\theta}}^{-\frac{1}{\sigma}} \frac{\tilde{\theta} - \theta + \Delta}{\Delta} d\tilde{\theta} + \int_{\theta}^{\theta+\Delta} (\kappa_{\tilde{\theta}} \phi_{\tilde{\theta}} h_{\tilde{\theta}})^{\frac{\sigma-1}{\sigma}} l_{\tilde{\theta}}^{-\frac{1}{\sigma}} \frac{\theta - \tilde{\theta} + \Delta}{\Delta} d\tilde{\theta} \right]. \end{aligned}$$

Applying L'Hôpital's rule yields:

$$\lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\theta-\Delta}^{\theta} (\kappa_{\tilde{\theta}} \phi_{\tilde{\theta}} h_{\tilde{\theta}})^{\frac{\sigma-1}{\sigma}} l_{\tilde{\theta}}^{-\frac{1}{\sigma}} \frac{\tilde{\theta} - \theta + \Delta}{\Delta} d\tilde{\theta} = \frac{1}{2} (\kappa_{\theta} \phi_{\theta} h_{\theta})^{\frac{\sigma-1}{\sigma}} l_{\theta}^{-\frac{1}{\sigma}}$$

and

$$\lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\theta}^{\theta+\Delta} (\kappa_{\tilde{\theta}} \phi_{\tilde{\theta}} h_{\tilde{\theta}})^{\frac{\sigma-1}{\sigma}} l_{\tilde{\theta}}^{-\frac{1}{\sigma}} \frac{\theta - \tilde{\theta} + \Delta}{\Delta} d\tilde{\theta} = \frac{1}{2} (\kappa_{\theta} \phi_{\theta} h_{\theta})^{\frac{\sigma-1}{\sigma}} l_{\theta}^{-\frac{1}{\sigma}},$$

where I used continuity of κ , h , ϕ , and l in θ . The former two are continuous by assumption; ϕ is continuous in equilibrium if l is continuous, as evident from equation (45) below; and continuity of l is presumed in all equilibria under consideration.

Finally, combine the two previous expressions to obtain

$$D_{l_{\theta}} F(l, \phi) = (\kappa_{\theta} \phi_{\theta} h_{\theta})^{\frac{\sigma-1}{\sigma}} l_{\theta}^{-\frac{1}{\sigma}} F(l, \phi)^{\frac{1}{\sigma}}$$

and therewith

$$w_{\theta}(l, \phi) = (\kappa_{\theta} \phi_{\theta})^{\frac{\sigma-1}{\sigma}} (l_{\theta} h_{\theta})^{-\frac{1}{\sigma}} F(l, \phi)^{\frac{1}{\sigma}}.$$

Wage Elasticities Again, I focus on the derivations for interior types. Given expression (15), the own-type substitution elasticity is simply the elasticity of w_{θ} with respect to l_{θ} :

$$\gamma_{\theta}^{own} = -\frac{1}{\sigma}.$$

The cross-type substitution elasticity $\gamma_{\theta, \tilde{\theta}}$ is

$$\begin{aligned}\gamma_{\theta, \tilde{\theta}} &= \frac{l_{\tilde{\theta}}}{w_{\theta}} D_{l_{\tilde{\theta}}} (\kappa_{\theta} \phi_{\theta})^{\frac{\sigma-1}{\sigma}} (l_{\theta} h_{\theta})^{-\frac{1}{\sigma}} F(l, \phi)^{\frac{1}{\sigma}} \\ &= \frac{l_{\tilde{\theta}}}{w_{\theta}} (\kappa_{\theta} \phi_{\theta})^{\frac{\sigma-1}{\sigma}} (l_{\theta} h_{\theta})^{-\frac{1}{\sigma}} \frac{1}{\sigma} F(l, \phi)^{\frac{1}{\sigma}-1} w_{\tilde{\theta}} h_{\tilde{\theta}} \\ &= \frac{1}{\sigma} \frac{w_{\tilde{\theta}}(l, \phi) h_{\tilde{\theta}} l_{\tilde{\theta}}}{F(l, \phi)}.\end{aligned}$$

For the directed technical change elasticities, consider first the determination of equilibrium technology described by equation (11). In the CES case, the first-order condition for the maximization problem in equation (11) is

$$\delta \phi_{\theta}^{*\delta-1} \lambda = (\kappa_{\theta} h_{\theta} l_{\theta})^{\frac{\sigma-1}{\sigma}} \phi_{\theta}^{*-1/\sigma} F(l, \phi)^{\frac{1}{\sigma}} \quad \forall \theta, \quad (44)$$

where λ is the Lagrange multiplier for the R&D resource constraint. The conditions equate the marginal R&D cost of raising ϕ_{θ} , converted into units of final good by λ , with the marginal gain in production. The latter is given by the derivative $D_{\phi_{\theta}} F$, which is computed analogously to $D_{l_{\theta}} F$ above.

Solving the first-order conditions for ϕ_{θ} yields

$$\phi_{\theta}^* = (\delta \lambda)^{\frac{-\sigma}{(\delta-1)\sigma+1}} (\kappa_{\theta} h_{\theta} l_{\theta})^{\frac{\sigma-1}{(\delta-1)\sigma+1}} F(l, \phi)^{\frac{1}{(\delta-1)\sigma+1}} \quad \forall \theta. \quad (45)$$

With this expression, we can use the R&D resource constraint

$$\int_{\underline{\theta}}^{\bar{\theta}} \phi_{\theta}^{*\delta} d\theta = \bar{C}$$

to solve for the Lagrange multiplier:

$$\lambda = \frac{1}{\delta} \bar{C}^{-\frac{(\delta-1)\sigma+1}{\delta\sigma}} \left[\int_{\underline{\theta}}^{\bar{\theta}} (\kappa_{\theta} h_{\theta} l_{\theta})^{\frac{(\sigma-1)\delta}{(\delta-1)\sigma+1}} d\theta \right]^{\frac{(\delta-1)\sigma+1}{\sigma\delta}} F(l, \phi)^{\frac{1}{\sigma}}.$$

Plugging this into equation (45), we obtain the following expression for the equilibrium technology ϕ^* :

$$\phi_{\theta}^* = \bar{C}^{\frac{1}{\delta}} (\kappa_{\theta} h_{\theta} l_{\theta})^{\frac{\sigma-1}{(\delta-1)\sigma+1}} \left[\int_{\underline{\theta}}^{\bar{\theta}} (\kappa_{\theta} h_{\theta} l_{\theta})^{\frac{(\sigma-1)\delta}{(\delta-1)\sigma+1}} d\theta \right]^{-\frac{1}{\delta}}. \quad (46)$$

We can now use equation (46) to derive the directed technical change elasticities. The own-type directed technical change elasticity is simply derived from equations (15) and (46) as

$$\begin{aligned}\rho_{\theta}^{own} &= \frac{\phi_{\theta}^*}{w_{\theta}} \frac{\partial w_{\theta}}{\partial \phi_{\theta}^*} \frac{l_{\theta}}{\phi_{\theta}^*} \frac{\partial \phi_{\theta}^*}{\partial l_{\theta}} \\ &= \frac{\sigma-1}{\sigma} \frac{\sigma-1}{(\delta-1)\sigma+1}.\end{aligned}$$

For the cross-type directed technical change elasticity, start from its definition:

$$\begin{aligned}\rho_{\theta,\tilde{\theta}} &= \frac{l_{\tilde{\theta}}}{w_{\theta}} D_{\phi,l_{\tilde{\theta}}} (\kappa_{\theta} \phi_{\theta}^*(l))^{\frac{\sigma-1}{\sigma}} (l_{\theta} h_{\theta})^{-\frac{1}{\sigma}} F(l, \phi^*(l))^{\frac{1}{\sigma}} \\ &= \frac{\sigma-1}{\sigma} \frac{l_{\tilde{\theta}}}{\phi_{\theta}^*} D_{l_{\tilde{\theta}}} \phi_{\theta}^*(l) + \frac{l_{\tilde{\theta}}}{w_{\theta}} (\kappa_{\theta} \phi_{\theta}^*(l))^{\frac{\sigma-1}{\sigma}} (l_{\theta} h_{\theta})^{-\frac{1}{\sigma}} D_{\phi,l_{\tilde{\theta}}} F(l, \phi^*(l))^{\frac{1}{\sigma}} .\end{aligned}\quad (47)$$

Here, the notation $D_{\phi,l_{\tilde{\theta}}}$ is introduced to capture only the effect of an increase in $l_{\tilde{\theta}}$ that is mediated by the adjustment of ϕ^* . Formally, if $x : (l, \phi) \mapsto x(l, \phi)$ is a function of labor input l and technology ϕ , the derivative $D_{\phi,l_{\theta}} x(l, \phi)$ is defined as

$$D_{\phi,l_{\theta}} x(l, \phi) := \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \left. \frac{dx(l, \phi^*(l + \mu \tilde{l}_{\Delta, \theta}))}{d\mu} \right|_{\mu=0} ,$$

where $\tilde{l}_{\Delta, \theta}$ is the perturbation function used in the definition of $D_{l_{\theta}}$ in Appendix A.1. The definition is analogous to that of $D_{l_{\theta}}$, with the exception that for $D_{\phi,l_{\theta}}$ the perturbation $\tilde{l}_{\Delta, \theta}$ affects x via ϕ^* instead of directly.

The second term of the sum in (47) is zero by the envelope theorem. So, we obtain

$$\rho_{\theta,\tilde{\theta}} = \frac{\sigma-1}{\sigma} \frac{l_{\tilde{\theta}}}{\phi_{\theta}^*} D_{l_{\tilde{\theta}}} \phi_{\theta}^*(l) .$$

We can compute $D_{l_{\tilde{\theta}}} \phi_{\theta}^*(l)$ analogously to the computation of $D_{l_{\theta}} F(l, \phi)$ above, using equation (46):

$$\begin{aligned}D_{l_{\tilde{\theta}}} \phi_{\theta}^*(l) &= -\frac{\sigma-1}{(\delta-1)\sigma+1} \bar{C}^{\frac{1}{\delta}} (\kappa_{\theta} h_{\theta} l_{\theta})^{\frac{\sigma-1}{(\delta-1)\sigma+1}} \left[\int_{\underline{\theta}}^{\bar{\theta}} (\kappa_{\theta} h_{\theta} l_{\theta})^{\frac{(\sigma-1)\delta}{(\delta-1)\sigma+1}} d\theta \right]^{-\frac{1}{\delta}} \\ &\quad \times \left[\int_{\underline{\theta}}^{\bar{\theta}} (\kappa_{\theta} h_{\theta} l_{\theta})^{\frac{(\sigma-1)\delta}{(\delta-1)\sigma+1}} d\theta \right]^{-1} (\kappa_{\tilde{\theta}} h_{\tilde{\theta}} l_{\tilde{\theta}})^{\frac{(\sigma-1)\delta}{(\delta-1)\sigma+1}} l_{\tilde{\theta}}^{-1} \\ &= -\frac{\sigma-1}{(\delta-1)\sigma+1} \frac{\phi_{\theta}^*(l)}{l_{\tilde{\theta}}} \frac{\phi_{\tilde{\theta}}^{*\delta}(l)}{\bar{C}} .\end{aligned}$$

Thereby,

$$\rho_{\theta,\tilde{\theta}} = -\frac{\sigma-1}{\sigma} \frac{\sigma-1}{(\delta-1)\sigma+1} \frac{\phi_{\tilde{\theta}}^{*\delta}(l)}{\bar{C}} .\quad (48)$$

To derive the expression from Section 5.2 in the main text, note that we can rewrite the first-order condition (44) as

$$\delta \phi_{\theta}^{*\delta-1} \lambda = w_{\theta} h_{\theta} l_{\theta} \phi_{\theta}^{*-1} ,$$

which implies

$$\phi_{\theta}^{*\delta} = \frac{1}{\lambda \delta} w_{\theta} h_{\theta} l_{\theta} .\quad (49)$$

We now integrate this over θ and use Euler's homogeneous function theorem to obtain

$$\bar{C} = \frac{1}{\lambda\delta} F.$$

Using this to eliminate λ in equation (49), we obtain

$$\phi_{\theta}^{*\delta} = \bar{C} \frac{w_{\theta} h_{\theta} l_{\theta}}{F}.$$

Finally, combining this with equation (48) yields

$$\rho_{\theta, \tilde{\theta}} = -\frac{\sigma-1}{\sigma} \frac{\sigma-1}{(\delta-1)\sigma+1} \frac{w_{\tilde{\theta}}(l, \phi) h_{\tilde{\theta}} l_{\tilde{\theta}}}{F(l, \phi)},$$

which is the expression given in the main text.

A.4 Linear Homogeneity of Aggregate Production

In the main text, I assume that final good firms' production function $\tilde{G}$ is linear homogeneous in the rival inputs l and q . Here, I show that the aggregate production function F and its equilibrium version F^* (defined below) inherit this property. These results will be useful in the proofs below.

Lemma 4. *The aggregate production function F defined in (9) is linear homogeneous in l .*

Proof. Aggregate production $F(l, \phi)$ for some labor input l and some technology ϕ is defined as

$$\max_q \left\{ \tilde{G}(\{h_{\theta} l_{\theta}\}_{\theta \in \Theta}, \phi, q) - \sum_{j=1}^J \eta_j q_j \right\}.$$

Let $q^*(l, \phi)$ denote a solution to this maximization problem.

Consider now the labor input λl for some $\lambda > 0$ and the intermediate input $\lambda q^*(l, \phi)$. Since $\tilde{G}$ is linear homogeneous in l and q , the first-order conditions of the maximization problem are satisfied at λl , $\lambda q^*(l, \phi)$, and ϕ . Since $\tilde{G}$ is concave in l and q , first-order conditions are sufficient for a maximum, and $\lambda q^*(l, \phi)$ is a maximizer of $\tilde{G}$ at λl and ϕ . So, using linear homogeneity of $\tilde{G}$ again,

$$\begin{aligned} F(\lambda l, \phi) &= \tilde{G}(\{\lambda h_{\theta} l_{\theta}\}_{\theta \in \Theta}, \phi, \lambda q^*(l, \phi)) - \sum_{j=1}^J \eta_j \lambda q_j^*(l, \phi) \\ &= \lambda \tilde{G}(\{h_{\theta} l_{\theta}\}_{\theta \in \Theta}, \phi, q^*(l, \phi)) - \sum_{j=1}^J \eta_j \lambda q_j^*(l, \phi) \\ &= \lambda F(l, \phi). \end{aligned}$$

□

Consider next the equilibrium aggregate production function

$$F^*(l) := F(l, \phi^*(l)) . \quad (50)$$

Lemma 5. *The equilibrium aggregate production function F^* is linear homogeneous in l .*

Proof. By the condition for equilibrium technology $\phi^*(l)$, the equilibrium aggregate production function satisfies

$$F^*(l) = \max_{\phi \in \Phi} F(l, \phi) .$$

Then, by linear homogeneity of F in l (see Lemma 4):

$$\begin{aligned} F^*(\lambda l) &= \max_{\phi \in \Phi} F(\lambda l, \phi) \\ &= \max_{\phi \in \Phi} \lambda F(l, \phi) \\ &= \lambda \max_{\phi \in \Phi} F(l, \phi) \\ &= \lambda F^*(l) \end{aligned}$$

for any $\lambda > 0$. □

A.5 Proof of Lemma 1: Weak Relative Bias

Lemma 1 is a local version of Theorem 2 in [Loebbing \(2020\)](#). Yet, it is not strictly covered by the theorem, because, as described in the main text and footnote 18, I use a slightly unusual definition of quasisupermodularity, which allows me to dispense with the lattice structure of Φ . Still, the proof follows closely the proof of Theorem 2 in [Loebbing \(2020\)](#).

Take any two labor inputs l and $\tilde{l}$ such that $\tilde{l}$ has greater relative skill supply, that is, $\tilde{l}_\theta / \tilde{l}_{\bar{\theta}} \geq l_\theta / l_{\bar{\theta}}$ for all $\theta \geq \bar{\theta}$. Since F is linear homogeneous in labor (Lemma 4), wages are independent of the scale of the labor input. So, for the purpose of Lemma 1, we can always scale l up or down such that $F(l, \phi^*(\tilde{l})) = F(\tilde{l}, \phi^*(\tilde{l}))$. In words, we scale l such that it is contained in the (exogenous-technology) isoquant of F through $(\tilde{l}, \phi^*(\tilde{l}))$.

Moreover, by definition of the equilibrium technology ϕ^* , we have $F(l, \phi^*(l)) \geq F(l, \underline{\phi})$ for all $\underline{\phi} \preceq^{sb} \phi^*(l), \phi^*(\tilde{l})$. Quasisupermodularity then implies that there is a $\bar{\phi} \succeq^{sb} \phi^*(l), \phi^*(\tilde{l})$ such that $F(l, \bar{\phi}) \geq F(l, \phi^*(\tilde{l}))$.

Now assume, to derive a contradiction, that $\phi^*(\tilde{l}) \not\succeq^{sb} \phi^*(l)$. Then, $\bar{\phi} \neq \phi^*(\tilde{l})$ and, by uniqueness of $\arg\max_{\phi \in \Phi} F(\tilde{l}, \phi)$ (Assumption 1), we must have $F(\tilde{l}, \phi^*(\tilde{l})) > F(\tilde{l}, \bar{\phi})$.

Combining the previous results, we obtain

$$F(l, \bar{\phi}) \geq F(l, \phi^*(\tilde{l})) = F(\tilde{l}, \phi^*(\tilde{l})) > F(\tilde{l}, \bar{\phi}) . \quad (51)$$

In words, increasing relative skill supply by moving from l to $\tilde{l}$ leaves output unchanged at $\phi^*(\tilde{l})$ but

reduces output at $\bar{\phi}$. Intuitively, this is incompatible with $\bar{\phi}$ being more skill complementary than $\phi^*(\tilde{l})$, which is what I show formally in the following.

To that end, consider a monotonic and differentiable path $l(\tau)$ from l to $\tilde{l}$ such that $l(0) = l$, $l(1) = \tilde{l}$ and $F(l(\tau), \phi^*(\tilde{l})) = F(l, \phi^*(\tilde{l}))$ for all $\tau \in [0, 1]$. (Monotonic means here that $l_\theta(\tau)$ is monotonic in τ for all θ .) Applying the mean value theorem, the inequalities in (51) imply that there is a $\tilde{\tau} \in (0, 1)$ such that

$$\int_{\underline{\theta}}^{\bar{\theta}} D_{l_\theta} F(l(\tilde{\tau}), \bar{\phi}) \frac{dl_\theta(\tilde{\tau})}{d\tau} d\theta < 0. \quad (52)$$

Let $\tilde{\theta}$ denote a skill level such that $l_\theta \leq \tilde{l}_\theta$ for all $\theta \leq \tilde{\theta}$ and $l_\theta \geq \tilde{l}_\theta$ for all $\theta > \tilde{\theta}$. Such a skill level exists because $\tilde{l}$ has greater relative skill supply than l . Noting that

$$\int_{\underline{\theta}}^{\bar{\theta}} D_{l_\theta} F(l(\tilde{\tau}), \phi^*(\tilde{l})) \frac{dl_\theta(\tilde{\tau})}{d\tau} d\theta = 0,$$

we can now extend inequality (52) to

$$\int_{\underline{\theta}}^{\bar{\theta}} \left(D_{l_\theta} F(l(\tilde{\tau}), \bar{\phi}) - \frac{D_{l_\theta} F(l(\tilde{\tau}), \bar{\phi})}{D_{l_\theta} F(l(\tilde{\tau}), \phi^*(\tilde{l}))} D_{l_\theta} F(l(\tilde{\tau}), \phi^*(\tilde{l})) \right) \frac{dl_\theta(\tilde{\tau})}{d\tau} d\theta < 0. \quad (53)$$

By definition of $\tilde{\theta}$ and monotonicity of $l(\tau)$, we know that $dl_\theta(\tilde{\tau})/d\tau$ is positive for all $\theta > \tilde{\theta}$ and negative for all $\theta \leq \tilde{\theta}$. Moreover, since $\bar{\phi} \succeq^{sb} \phi^*(\tilde{l})$, the difference

$$D_{l_\theta} F(l(\tilde{\tau}), \bar{\phi}) - \frac{D_{l_\theta} F(l(\tilde{\tau}), \bar{\phi})}{D_{l_\theta} F(l(\tilde{\tau}), \phi^*(\tilde{l}))} D_{l_\theta} F(l(\tilde{\tau}), \phi^*(\tilde{l}))$$

is also positive for $\theta > \tilde{\theta}$ and negative for $\theta \leq \tilde{\theta}$. This implies that the left-hand side of (53) must be (weakly) positive, a contradiction.

Hence, we have shown that $\phi^*(\tilde{l}) \succeq^{sb} \phi^*(l)$, that is, the equilibrium technology is more skill-biased under $\tilde{l}$ than under l . So, the increase in relative skill supply induces skill-biased technical change. The local implication of this global result is Lemma 1.

A.6 Proof of Lemma 2: Progressive Tax Reforms

I consider the case of a progressive reform. For regressive reforms, all signs are reversed, but the proof is otherwise unchanged.

For a progressive reform (τ, μ) , we can restate the inequality in Lemma 2 as

$$\frac{R'_T(y) - \mu\tau'(y)}{R'_T(y)} \leq \frac{R'_T(\tilde{y}) - \mu\tau'(\tilde{y})}{R'_T(\tilde{y})} \quad \forall y \geq \tilde{y}.$$

Taking logs and rearranging yields

$$\log R'_T(y) - \log R'_T(\tilde{y}) \leq \log R'_T(\tilde{y}) - \log R'_T(\tilde{y}) \quad \forall y \geq \tilde{y}.$$

In words, the difference between the log marginal retention rates under the new and the old tax function is decreasing in income. Equivalently, the derivative of this difference with respect to income is negative:

$$\frac{d \left(\log R'_{\tilde{T}}(y) - \log R'_T(y) \right)}{dy} \leq 0 \quad \forall y.$$

But this, using the definition of the rate of progressivity P of a tax function, is equivalent to

$$-\frac{1}{y} (P_{\tilde{T}}(y) - P_T(y)) \leq 0 \quad \forall y$$

and hence,

$$P_{\tilde{T}}(y) \geq P_T \quad \forall y.$$

So, the inequality in Lemma 2 is equivalent to Definition 2 of a progressive tax reform.

A.7 Proof of Lemma 3: Labor Input Effects of Tax Reforms

The proof proceeds in five steps.

Step 1. It is easy to see that

$$\hat{l}_{\theta,\tau}^{(n)} = \widetilde{TE}_{\theta,\tau}^{(n)} + \widetilde{SE}_{\theta,\tau}^{(n)}$$

for all $n \geq 1$. Hence, the two expressions (24) and (25) are equal.

Step 2. Suppose for now that all the series in expressions (24) and (25) converge. Then, take expression (24) and insert it into the fixed point equation (21):

$$\begin{aligned} \sum_{n=1}^{\infty} \hat{l}_{\theta,\tau}^{(n)} &= -\epsilon_{\theta}^R \frac{\tau'(y_{\theta}(T))}{1 - T'(y_{\theta}(T))} + \epsilon_{\theta}^w \zeta_{\theta}^{own} \sum_{n=1}^{\infty} \hat{l}_{\theta,\tau}^{(n)} + \epsilon_{\theta}^w \int_{\underline{\theta}}^{\bar{\theta}} \zeta_{\theta,\tilde{\theta}} \sum_{n=1}^{\infty} \hat{l}_{\tilde{\theta},\tau}^{(n)} d\tilde{\theta} \\ &= -\epsilon_{\theta}^R \frac{\tau'(y_{\theta}(T))}{1 - T'(y_{\theta}(T))} + \sum_{n=1}^{\infty} \left[\epsilon_{\theta}^w \zeta_{\theta}^{own} \hat{l}_{\theta,\tau}^{(n)} + \epsilon_{\theta}^w \int_{\underline{\theta}}^{\bar{\theta}} \zeta_{\theta,\tilde{\theta}} \hat{l}_{\tilde{\theta},\tau}^{(n)} d\tilde{\theta} \right] \\ &= -\epsilon_{\theta}^R \frac{\tau'(y_{\theta}(T))}{1 - T'(y_{\theta}(T))} + \sum_{n=1}^{\infty} \hat{l}_{\theta,\tau}^{(n+1)} \\ &= -\epsilon_{\theta}^R \frac{\tau'(y_{\theta}(T))}{1 - T'(y_{\theta}(T))} + \sum_{n=2}^{\infty} \hat{l}_{\theta,\tau}^{(n)} \\ &= \sum_{n=1}^{\infty} \hat{l}_{\theta,\tau}^{(n)}. \end{aligned}$$

This proves that, conditional upon convergence of the series, expression (24) solves the fixed point equation (21). Then, by Step 1, expression (25) also solves the fixed point equation conditional upon convergence.

Step 3. Regarding convergence, consider expression (24) first. Start from the definition of $\hat{l}_{\theta,\tau}^{(n)}$ and

take the square of both sides of the equation:

$$\left(\widehat{l}_{\theta,\tau}^{(n)}\right)^2 = (\epsilon_{\theta}^w \zeta_{\theta}^{own})^2 \left(\widehat{l}_{\theta,\tau}^{(n-1)}\right)^2 + (\epsilon_{\theta}^w)^2 \left(\int_{\underline{\theta}}^{\bar{\theta}} \zeta_{\theta,\tilde{\theta}} \widehat{l}_{\tilde{\theta},\tau}^{(n-1)} d\tilde{\theta}\right)^2 + 2\epsilon_{\theta}^w \zeta_{\theta}^{own} \widehat{l}_{\theta,\tau}^{(n-1)} \epsilon_{\theta}^w \int_{\underline{\theta}}^{\bar{\theta}} \zeta_{\theta,\tilde{\theta}} \widehat{l}_{\tilde{\theta},\tau}^{(n-1)} d\tilde{\theta}.$$

By the Cauchy-Schwarz inequality,

$$\begin{aligned} \left(\widehat{l}_{\theta,\tau}^{(n)}\right)^2 &\leq (\epsilon_{\theta}^w \zeta_{\theta}^{own})^2 \left(\widehat{l}_{\theta,\tau}^{(n-1)}\right)^2 + (\epsilon_{\theta}^w)^2 \int_{\underline{\theta}}^{\bar{\theta}} \zeta_{\theta,\tilde{\theta}}^2 d\tilde{\theta} \int_{\underline{\theta}}^{\bar{\theta}} \left(\widehat{l}_{\tilde{\theta},\tau}^{(n-1)}\right)^2 d\tilde{\theta} \\ &\quad + 2\epsilon_{\theta}^w \zeta_{\theta}^{own} \widehat{l}_{\theta,\tau}^{(n-1)} \epsilon_{\theta}^w \sqrt{\int_{\underline{\theta}}^{\bar{\theta}} \zeta_{\theta,\tilde{\theta}}^2 d\tilde{\theta}} \sqrt{\int_{\underline{\theta}}^{\bar{\theta}} \left(\widehat{l}_{\tilde{\theta},\tau}^{(n-1)}\right)^2 d\tilde{\theta}}, \end{aligned}$$

and after integrating over θ ,

$$\begin{aligned} \int_{\underline{\theta}}^{\bar{\theta}} \left(\widehat{l}_{\theta,\tau}^{(n)}\right)^2 d\theta &\leq \int_{\underline{\theta}}^{\bar{\theta}} (\epsilon_{\theta}^w \zeta_{\theta}^{own})^2 \left(\widehat{l}_{\theta,\tau}^{(n-1)}\right)^2 d\theta + \int_{\underline{\theta}}^{\bar{\theta}} (\epsilon_{\theta}^w)^2 \int_{\underline{\theta}}^{\bar{\theta}} \zeta_{\theta,\tilde{\theta}}^2 d\tilde{\theta} d\theta \int_{\underline{\theta}}^{\bar{\theta}} \left(\widehat{l}_{\tilde{\theta},\tau}^{(n-1)}\right)^2 d\tilde{\theta} \\ &\quad + 2 \int_{\underline{\theta}}^{\bar{\theta}} \epsilon_{\theta}^w \zeta_{\theta}^{own} \widehat{l}_{\theta,\tau}^{(n-1)} \epsilon_{\theta}^w \sqrt{\int_{\underline{\theta}}^{\bar{\theta}} \zeta_{\theta,\tilde{\theta}}^2 d\tilde{\theta}} d\theta \sqrt{\int_{\underline{\theta}}^{\bar{\theta}} \left(\widehat{l}_{\tilde{\theta},\tau}^{(n-1)}\right)^2 d\tilde{\theta}}. \end{aligned}$$

Taking the supremum of $\epsilon_{\theta}^w \zeta_{\theta}^{own}$ in the first term and applying the Cauchy-Schwarz inequality again to the last term yields:

$$\begin{aligned} \int_{\underline{\theta}}^{\bar{\theta}} \left(\widehat{l}_{\theta,\tau}^{(n)}\right)^2 d\theta &\leq \sup_{\theta \in \Theta} [(\epsilon_{\theta}^w \zeta_{\theta}^{own})^2] \int_{\underline{\theta}}^{\bar{\theta}} \left(\widehat{l}_{\theta,\tau}^{(n-1)}\right)^2 d\theta + \int_{\underline{\theta}}^{\bar{\theta}} \int_{\underline{\theta}}^{\bar{\theta}} (\epsilon_{\theta}^w \zeta_{\theta,\tilde{\theta}})^2 d\tilde{\theta} d\theta \int_{\underline{\theta}}^{\bar{\theta}} \left(\widehat{l}_{\tilde{\theta},\tau}^{(n-1)}\right)^2 d\tilde{\theta} \\ &\quad + 2 \sqrt{\int_{\underline{\theta}}^{\bar{\theta}} \int_{\underline{\theta}}^{\bar{\theta}} (\epsilon_{\theta}^w \zeta_{\theta}^{own} \epsilon_{\theta}^w \zeta_{\theta,\tilde{\theta}})^2 d\tilde{\theta} d\theta} \int_{\underline{\theta}}^{\bar{\theta}} \left(\widehat{l}_{\tilde{\theta},\tau}^{(n-1)}\right)^2 d\tilde{\theta}. \end{aligned}$$

The coefficients of $\int_{\underline{\theta}}^{\bar{\theta}} \left(\widehat{l}_{\theta,\tau}^{(n-1)}\right)^2 d\theta$ on the right-hand side of the inequality amount to

$$\sup_{\theta \in \Theta} (\epsilon_{\theta}^w \zeta_{\theta}^{own})^2 + \int_{\underline{\theta}}^{\bar{\theta}} (\epsilon_{\theta}^w \zeta_{\theta,\tilde{\theta}})^2 d\tilde{\theta} d\theta + 2 \sqrt{\int_{\underline{\theta}}^{\bar{\theta}} (\epsilon_{\theta}^w \zeta_{\theta}^{own})^2 (\epsilon_{\theta}^w \zeta_{\theta,\tilde{\theta}})^2 d\tilde{\theta} d\theta},$$

which is strictly smaller than one by condition (23). Hence, the term $\int_{\underline{\theta}}^{\bar{\theta}} \left(\widehat{l}_{\theta,\tau}^{(n-1)}\right)^2 d\theta$ is dominated by a geometric sequence converging to zero.

Regarding $\widehat{l}_{\theta,\tau}^{(n)}$, the Cauchy-Schwarz inequality implies

$$\widehat{l}_{\theta,\tau}^{(n)} \leq \epsilon_{\theta}^w \zeta_{\theta}^{own} \widehat{l}_{\theta,\tau}^{(n-1)} + \epsilon_{\theta}^w \sqrt{\int_{\underline{\theta}}^{\bar{\theta}} \zeta_{\theta,\tilde{\theta}}^2 d\tilde{\theta}} \sqrt{\int_{\underline{\theta}}^{\bar{\theta}} \left(\widehat{l}_{\tilde{\theta},\tau}^{(n-1)}\right)^2 d\tilde{\theta}}. \quad (54)$$

Suppose now, to derive a contradiction, that $\widehat{l}_{\theta,\tau}^{(n)}$ is not dominated by any geometric sequence that

converges to zero. Then, for any $c \in (\epsilon_\theta^w \zeta_\theta^{own}, 1)$ and for any $N \in \mathbb{N}$, there must exist $\bar{N}_N > N$ such that

$$\frac{|\hat{l}_{\theta,\tau}^{(\bar{N}_N)}|}{|\hat{l}_{\theta,\tau}^{(\bar{N}_N-1)}|} > c .$$

At the same time, since $\int_{\underline{\theta}}^{\bar{\theta}} \left(\hat{l}_{\theta,\tau}^{(n-1)}\right)^2 d\theta$ is dominated by a geometric sequence converging to zero, we must have

$$\frac{|\epsilon_{\theta,\tau}^w \sqrt{\int_{\underline{\theta}}^{\bar{\theta}} \zeta_{\theta,\bar{\theta}}^2 d\tilde{\theta}} \sqrt{\int_{\underline{\theta}}^{\bar{\theta}} \left(\hat{l}_{\theta,\tau}^{(n-1)}\right)^2 d\tilde{\theta}}|}{\hat{l}_{\theta,\tau}^{(\bar{N}_N)}} \rightarrow 0 \quad \text{as } N \rightarrow \infty .$$

But with equation (54) this implies, as $N \rightarrow \infty$,

$$\frac{|\hat{l}_{\theta,\tau}^{(\bar{N}_N)}|}{|\hat{l}_{\theta,\tau}^{(\bar{N}_N-1)}|} \rightarrow |\epsilon_\theta^w \zeta_\theta^{own}| < c ,$$

a contradiction.

So, $\hat{l}_{\theta,\tau}^{(n)}$ is dominated by a geometric sequence converging to zero and the series $\sum_{n=1}^{\infty} \hat{l}_{\theta,\tau}^{(n)}$ indeed exists.

Step 4. For convergence of the series $\widetilde{TE}_{\theta,\tau}$ and $\widetilde{SE}_{\theta,\tau}$, consider $\widetilde{TE}_{\theta,\tau}$ first. Replacing $\zeta_{\theta,\bar{\theta}}$ by $\rho_{\theta,\bar{\theta}}$, the reasoning in step 3 implies that $\widetilde{TE}_{\theta,\tau}$ converges. Second, note that

$$\widetilde{SE}_{\theta,\tau} = \sum_{n=1}^{\infty} \hat{l}_{\theta,\tau}^{(n)} - \widetilde{TE}_{\theta,\tau} .$$

Since we have already shown that both series on the right-hand side converge, the same must hold for $\widetilde{SE}_{\theta,\tau}$.

Step 5. The final step is to prove that, with constant labor supply elasticities, the terms $\widetilde{DE}_{\theta,\tau}$ and $\widetilde{TE}_{\theta,\tau}$ are decreasing in θ for any progressive reform τ and that this extends to the total effect $\hat{l}_{\theta,\tau}$ in the case of strong relative bias.

First, the result that $\widetilde{DE}_{\theta,\tau}$ decreases in θ if ϵ_θ^R is constant in θ follows immediately from the characterization of progressive tax reforms in Lemma 2.

Next, we can prove inductively that the term $\widetilde{TE}_{\theta,\tau}$ is decreasing in θ . Since the direct effect $-\epsilon_\theta^R \tau'(y_\theta(T)) / (1 - T'(y_\theta(T)))$ decreases in θ , Lemma 1 on weak relative bias implies that

$$\rho_\theta^{own}(-\epsilon_\theta^R) \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} + \int_{\underline{\theta}}^{\bar{\theta}} \rho_{\theta,\bar{\theta}}(-\epsilon_{\bar{\theta}}^R) \frac{\tau'(y_{\bar{\theta}}(T))}{1 - T'(y_{\bar{\theta}}(T))} d\tilde{\theta}$$

is decreasing in θ . With ϵ_θ^w constant in θ , this implies that the first element of $\widetilde{TE}_{\theta,\tau}$,

$$\widetilde{TE}_{\theta,\tau}^{(1)} = \epsilon_\theta^w \rho_\theta^{own}(-\epsilon_\theta^R) \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} + \epsilon_\theta^w \int_{\underline{\theta}}^{\bar{\theta}} \rho_{\theta,\bar{\theta}}(-\epsilon_{\bar{\theta}}^R) \frac{\tau'(y_{\bar{\theta}}(T))}{1 - T'(y_{\bar{\theta}}(T))} d\tilde{\theta} ,$$

decreases in θ .

Suppose now that $\widetilde{TE}_{\theta,\tau}^{(n)}$ is decreasing in θ . Then, again by Lemma 1 and because ϵ_θ^w is constant in θ , we have that

$$\epsilon_\theta^w \rho_\theta^{own} \widetilde{TE}_{\theta,\tau}^{(n)} + \epsilon_\theta^w \int_{\underline{\theta}}^{\bar{\theta}} \rho_{\theta,\tilde{\theta}} \widetilde{TE}_{\theta,\tau}^{(n)} d\tilde{\theta}$$

decreases in θ , which is equal to $\widetilde{TE}_{\theta,\tau}^{(n+1)}$. So, by induction, $\widetilde{TE}_{\theta,\tau}^{(n)}$ decreases in θ for all $n \geq 1$. The sum $\sum_{n=1}^{\infty} \widetilde{TE}_{\theta,\tau}^{(n)}$ must then also decrease in θ , which yields the desired result.

Finally, note that the same inductive reasoning applies to the total labor input effect $\widehat{l}_{\theta,\tau}$ in the case of strong relative bias. In particular, the first element $\widehat{l}_{\theta,\tau}^{(0)}$ decreases in θ as an immediate consequence of Lemma 2 and the constancy of ϵ_θ^R . Next, with constancy of ϵ_θ^w and under strong relative bias, if $\widehat{l}_{\theta,\tau}^{(n)}$ decreases in θ , so must $\widehat{l}_{\theta,\tau}^{(n+1)}$. Thus, by induction, the infinite sum $\sum_{n=0}^{\infty} \widehat{l}_{\theta,\tau}^{(n)}$ decreases in θ as well.

A.8 Proof of Proposition 1: Wage Effects of Tax Reforms

Equation (26) is obtained immediately by inserting equation (25) from Lemma 3 into (19).

To sign the slope of $TE_{\theta,\tau}$ for progressive reforms, note that, with constant labor supply elasticities, the sum $\widetilde{DE}_{\theta,\tau} + \widetilde{TE}_{\theta,\tau}$ is decreasing in θ for any progressive reform by Lemma 3. Then, by weak relative bias (Lemma 1), the term

$$\rho_\theta^{own} \left(\widetilde{DE}_{\theta,\tau}(T) + \widetilde{TE}_{\theta,\tau}(T) \right) + \int_{\underline{\theta}}^{\bar{\theta}} \rho_{\theta,\tilde{\theta}} \left(\widetilde{DE}_{\tilde{\theta},\tau}(T) + \widetilde{TE}_{\tilde{\theta},\tau}(T) \right) d\tilde{\theta},$$

which is equal to $TE_{\theta,\tau}$, must also decrease in τ for any progressive reform.

Next, if there is strong relative bias and labor supply elasticities are constant, the total labor input response $\widehat{l}_{\theta,\tau}$ is decreasing in θ for any progressive reform by Lemma 3. Then, by definition of strong relative bias, the total wage effect

$$(\gamma_\theta^{own} + \rho_\theta^{own}) \widehat{l}_{\theta,\tau} + \int_{\underline{\theta}}^{\bar{\theta}} (\gamma_{\theta,\tilde{\theta}} + \rho_{\theta,\tilde{\theta}}) \widehat{l}_{\tilde{\theta},\tau} d\tilde{\theta}$$

must also decrease in θ .

A.9 Proof of Corollary 1: Wage Effects of Tax Reforms in the CES Case

I first derive specific expressions for the labor input responses to tax reforms for the CES case.

Lemma 6. Fix an initial tax T and assume that F and Φ are CES as introduced in Section 3.6. Moreover, let the elasticities ϵ_θ^w and ϵ_θ^R be constant in θ , that is, $\epsilon_\theta^w = \epsilon^w$ and $\epsilon_\theta^R = \epsilon^R$ for all $\theta \in \Theta$. Then, the effect of tax reform τ on labor inputs can be written as

$$\widehat{l}_{\theta,\tau}(T) = -\bar{\epsilon}^R \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} - (\gamma^{CES} + \rho^{CES}) \epsilon^w \int_{\underline{\theta}}^{\bar{\theta}} \frac{w_{\tilde{\theta}} l_{\tilde{\theta}} h_{\tilde{\theta}}}{F(l(T), \phi^*(T))} (-\bar{\epsilon}^R) \frac{\tau'(y_{\tilde{\theta}}(T))}{1 - T'(y_{\tilde{\theta}}(T))} d\tilde{\theta}, \quad (55)$$

where

$$\bar{\epsilon}^R := \frac{\epsilon^R}{1 - (\gamma^{CES} + \rho^{CES})\epsilon^w}.$$

For any progressive reform τ , the relative effect on labor inputs $\hat{l}_{\theta,\tau}$ is decreasing in θ .

Proof. The fastest way to prove equation (55) is to check that it satisfies the fixed point equation (21). In the CES case and with ϵ_θ^w constant, this equation becomes

$$\hat{l}_{\theta,\tau}(T) = -\epsilon_\theta^R \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} + \epsilon^w(\gamma^{CES} + \rho^{CES})\hat{l}_{\theta,\tau}(T) - \epsilon^w(\gamma^{CES} + \rho^{CES}) \int_{\underline{\theta}}^{\bar{\theta}} \frac{w_{\tilde{\theta}} l_{\tilde{\theta}} h_{\tilde{\theta}}}{F(l(T), \phi^*(T))} \hat{l}_{\tilde{\theta},\tau} d\tilde{\theta}.$$

Inserting equation (55) yields:

$$\begin{aligned} \hat{l}_{\theta,\tau}(T) &= -\epsilon_\theta^R \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} + \epsilon^w(\gamma^{CES} + \rho^{CES})(-\bar{\epsilon}_\theta^R) \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} \\ &\quad - \epsilon^w(\gamma^{CES} + \rho^{CES})(\gamma^{CES} + \rho^{CES})\epsilon^w \int_{\underline{\theta}}^{\bar{\theta}} \frac{w_{\tilde{\theta}} l_{\tilde{\theta}} h_{\tilde{\theta}}}{F(l(T), \phi^*(T))} (-\bar{\epsilon}_{\tilde{\theta}}^R) \frac{\tau'(y_{\tilde{\theta}}(T))}{1 - T'(y_{\tilde{\theta}}(T))} d\tilde{\theta} \\ &\quad - \epsilon^w(\gamma^{CES} + \rho^{CES}) \int_{\underline{\theta}}^{\bar{\theta}} \frac{w_{\tilde{\theta}} l_{\tilde{\theta}} h_{\tilde{\theta}}}{F(l(T), \phi^*(T))} (-\bar{\epsilon}_{\tilde{\theta}}^R) \frac{\tau'(y_{\tilde{\theta}}(T))}{1 - T'(y_{\tilde{\theta}}(T))} d\tilde{\theta} \\ &\quad + \epsilon^w(\gamma^{CES} + \rho^{CES}) \int_{\underline{\theta}}^{\bar{\theta}} \frac{w_{\tilde{\theta}} l_{\tilde{\theta}} h_{\tilde{\theta}}}{F(l(T), \phi^*(T))} (\gamma^{CES} + \rho^{CES})\epsilon^w \int_{\underline{\theta}}^{\bar{\theta}} \frac{w_{\hat{\theta}} l_{\hat{\theta}} h_{\hat{\theta}}}{F(l(T), \phi^*(T))} (-\bar{\epsilon}_{\hat{\theta}}^R) \frac{\tau'(y_{\hat{\theta}}(T))}{1 - T'(y_{\hat{\theta}}(T))} d\hat{\theta} d\tilde{\theta} \\ &= -\bar{\epsilon}_\theta^R \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} - \epsilon^w(\gamma^{CES} + \rho^{CES}) \int_{\underline{\theta}}^{\bar{\theta}} \frac{w_{\tilde{\theta}} l_{\tilde{\theta}} h_{\tilde{\theta}}}{F(l(T), \phi^*(T))} (-\bar{\epsilon}_{\tilde{\theta}}^R) \frac{\tau'(y_{\tilde{\theta}}(T))}{1 - T'(y_{\tilde{\theta}}(T))} d\tilde{\theta} \\ &\quad - (\epsilon^w)^2(\gamma^{CES} + \rho^{CES})^2 \int_{\underline{\theta}}^{\bar{\theta}} \frac{w_{\tilde{\theta}} l_{\tilde{\theta}} h_{\tilde{\theta}}}{F(l(T), \phi^*(T))} (-\bar{\epsilon}_{\tilde{\theta}}^R) \frac{\tau'(y_{\tilde{\theta}}(T))}{1 - T'(y_{\tilde{\theta}}(T))} d\tilde{\theta} \\ &\quad + (\epsilon^w)^2(\gamma^{CES} + \rho^{CES})^2 \int_{\underline{\theta}}^{\bar{\theta}} \frac{w_{\tilde{\theta}} l_{\tilde{\theta}} h_{\tilde{\theta}}}{F(l(T), \phi^*(T))} (-\bar{\epsilon}_{\tilde{\theta}}^R) \frac{\tau'(y_{\tilde{\theta}}(T))}{1 - T'(y_{\tilde{\theta}}(T))} \int_{\underline{\theta}}^{\bar{\theta}} \frac{w_{\hat{\theta}} l_{\hat{\theta}} h_{\hat{\theta}}}{F(l(T), \phi^*(T))} d\hat{\theta} d\tilde{\theta} \\ &= -\bar{\epsilon}_\theta^R \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} - \epsilon^w(\gamma^{CES} + \rho^{CES}) \int_{\underline{\theta}}^{\bar{\theta}} \frac{w_{\tilde{\theta}} l_{\tilde{\theta}} h_{\tilde{\theta}}}{F(l(T), \phi^*(T))} (-\bar{\epsilon}_{\tilde{\theta}}^R) \frac{\tau'(y_{\tilde{\theta}}(T))}{1 - T'(y_{\tilde{\theta}}(T))} d\tilde{\theta}, \end{aligned}$$

where the last equality follows from the fact that F is linear homogeneous in l (see Lemma 4) and Euler's homogeneous function theorem. So, equation (55) solves the fixed point equation (21).

The remainder of Lemma 6 follows from the observation that the second term on the right-hand side of equation (55) is independent of θ and that, by Lemma 2, the first term decreases in θ if and only if the reform τ is progressive. $\square$

Using Lemma 6, we can now prove Corollary 1. To derive equation (30), note that, since the equilibrium aggregate production function F^* is linear homogeneous in l (see Lemma 5 in Appendix A.4), the total wage effects of a proportional change in all types' labor inputs are zero:

$$(\gamma_\theta^{own} + \rho_\theta^{own}) \hat{l}_{\theta,\tau}(T) + \int_{\underline{\theta}}^{\bar{\theta}} (\gamma_{\theta,\tilde{\theta}} + \rho_{\theta,\tilde{\theta}}) \hat{l}_{\tilde{\theta},\tau}(T) d\tilde{\theta} = 0 \quad \text{for all } \theta$$

if $\widehat{l}_{\theta,\tau}(T)$ is constant in θ . Hence, inserting equation (55) into equation (19), the second term of equation (55) vanishes. This leaves

$$\begin{aligned} \widehat{w}_{\theta,\tau}(T) = & \left(\gamma^{CES} + \rho^{CES} \right) (-\bar{\epsilon}^R) \frac{\tau'(y_{\theta}(T))}{1 - T'(y_{\theta}(T))} \\ & - \left(\gamma^{CES} + \rho^{CES} \right) \int_{\underline{\theta}}^{\bar{\theta}} \frac{w_{\tilde{\theta}} l_{\tilde{\theta}} h_{\tilde{\theta}}}{F(l(T), \phi^*(T))} (-\bar{\epsilon}^R) \frac{\tau'(y_{\tilde{\theta}}(T))}{1 - T'(y_{\tilde{\theta}}(T))} d\tilde{\theta}, \quad (56) \end{aligned}$$

which is equation (30).

To sign the slope of the total wage effect, note that the second term on the right-hand side of equation (56) is independent of θ . Thus, the total wage effect inherits its dependency on θ from the first term. By Lemma 2, $\tau'(y_{\theta}(T))/(1 - T'(y_{\theta}(T)))$ is increasing in θ for any progressive reform. Thus, first, the directed technical change effect $\rho^{CES}(-\bar{\epsilon}^R)\tau'(y_{\theta}(T))/(1 - T'(y_{\theta}(T)))$ is decreasing for any progressive reform. Second, if there is strong relative bias, $\gamma^{CES} + \rho^{CES} \geq 0$ and the entire first term in equation (56) decreases in θ . Without strong bias, these signs are reversed.

A.10 Proof of Proposition 2: Optimal Tax Formula

The optimal tax formula in Proposition 2 is obtained by maximizing welfare $W(c, l)$ subject to the resource constraint (31) and the incentive compatibility constraint (32). The derivation proceeds along the following steps: first eliminate consumption from the welfare maximization problem, then derive first-order conditions, use workers' first-order condition to reintroduce marginal tax rates into the equations, and finally derive the signs of the term TE_{θ}^* at the bottom and the top of the type distribution from directed technical change theory.

As a first step, the following lemma shows how to eliminate consumption from the welfare maximization problem.

Lemma 7. *The pair of consumption and labor inputs (c, l) satisfies the resource and incentive compatibility constraints (31) and (32) if and only if $c = c^*(l)$, where $c^*(l) = \{c_{\theta}^*(l)\}_{\theta \in \Theta}$ is determined by*

$$c_{\theta}^*(l) = F(l, \phi^*(l)) - \int_{\underline{\theta}}^{\bar{\theta}} v(l_{\tilde{\theta}}) h_{\tilde{\theta}} d\tilde{\theta} - \int_{\underline{\theta}}^{\bar{\theta}} v'(l_{\tilde{\theta}}) l_{\tilde{\theta}} (1 - H_{\tilde{\theta}}) \widehat{w}_{\tilde{\theta}} d\tilde{\theta} + \int_{\underline{\theta}}^{\theta} v'(l_{\tilde{\theta}}) l_{\tilde{\theta}} \widehat{w}_{\tilde{\theta}} d\tilde{\theta} + v(l_{\theta}) \quad (57)$$

for all θ .

Proof. ($\Rightarrow$) I first show that the constraints (31) and (32) imply equation (57). For that, write consumption as

$$c_{\theta} = c_{\underline{\theta}} + \int_{\underline{\theta}}^{\theta} c'_{\tilde{\theta}} d\tilde{\theta}.$$

By the incentive compatibility constraint (32), this implies for all θ :

$$\begin{aligned} c_\theta &= c_{\underline{\theta}} + \int_{\underline{\theta}}^{\theta} v'(l_{\tilde{\theta}})(w'_{\tilde{\theta}}l_{\tilde{\theta}} + w_{\tilde{\theta}}l'_{\tilde{\theta}}) \frac{1}{w_{\tilde{\theta}}} d\tilde{\theta} \\ &= c_{\underline{\theta}} + \int_{\underline{\theta}}^{\theta} v'(l_{\tilde{\theta}})\widehat{w}_{\tilde{\theta}}l_{\tilde{\theta}} d\tilde{\theta} + v(l_\theta) - v(l_{\underline{\theta}}) . \end{aligned} \quad (58)$$

Combining this with the resource constraint (31), we obtain the following expression for $c_{\underline{\theta}}$:

$$c_{\underline{\theta}} = F(l, \phi^*(l)) - \int_{\underline{\theta}}^{\bar{\theta}} \int_{\underline{\theta}}^{\theta} v'(l_{\tilde{\theta}})\widehat{w}_{\tilde{\theta}}l_{\tilde{\theta}} d\tilde{\theta} h_\theta d\theta - \int_{\underline{\theta}}^{\bar{\theta}} v(l_\theta)h_\theta d\theta + v(l_{\underline{\theta}}) .$$

Using integration by parts to solve the double integral yields:

$$c_{\underline{\theta}} = F(l, \phi^*(l)) - \int_{\underline{\theta}}^{\bar{\theta}} v'(l_\theta)\widehat{w}_\theta l_\theta (1 - H_\theta) d\theta - \int_{\underline{\theta}}^{\bar{\theta}} v(l_\theta)h_\theta d\theta + v(l_{\underline{\theta}}) .$$

Inserting this back into equation (58), we obtain:

$$c_\theta = F(l, \phi^*(l)) - \int_{\underline{\theta}}^{\bar{\theta}} v'(l_{\tilde{\theta}})\widehat{w}_{\tilde{\theta}}l_{\tilde{\theta}}(1 - H_{\tilde{\theta}}) d\tilde{\theta} - \int_{\underline{\theta}}^{\bar{\theta}} v(l_{\tilde{\theta}})h_{\tilde{\theta}} d\tilde{\theta} + \int_{\underline{\theta}}^{\theta} v'(l_{\tilde{\theta}})\widehat{w}_{\tilde{\theta}}l_{\tilde{\theta}} d\tilde{\theta} + v(l_\theta) ,$$

which is equation (57) defining c^* above.

($\Leftarrow$) Differentiating c^* with respect to θ shows immediately that equation (57) implies the incentive compatibility constraint (32). Similarly, after multiplying c_θ^* by h_θ and integrating over $[\underline{\theta}, \bar{\theta}]$, standard computations show that

$$\int_{\underline{\theta}}^{\bar{\theta}} c_\theta^*(l) d\theta = F(l, \phi^*(l)) ,$$

which proves that equation (57) also implies the resource constraint (31). $\square$

Lemma 7 provides an equivalent representation of resource and incentive compatibility constraints, which is explicitly solved for c . Hence, instead of maximizing welfare subject to the two constraints, we can study the unconstrained maximization of

$$\widehat{W}(l) := W(c^*(l), l)$$

with l being the only choice variable. The first part of the proof now uses the first-order conditions of this unconstrained problem to derive the condition for optimal marginal tax rates provided in Proposition 2.

Part 1. The first-order conditions are given by

$$D_{l_\theta} \widehat{W}(l) = 0 \quad \text{for all } \theta .$$

We hence study the derivative $D_{l_\theta} \widehat{W}(l)$ first. Using the notation for welfare weights introduced in the

main text and the definition of D_{l_θ} from Appendix A.1, the derivative can be written as

$$\begin{aligned}
D_{l_\theta} \widehat{W}(l) &= w_\theta h_\theta - v'(l_\theta) h_\theta \\
&\quad - (v''(l_\theta) l_\theta + v'(l_\theta)) (1 - H_\theta) \widehat{w}_\theta + \widetilde{g}_\theta (1 - H_\theta) (v''(l_\theta) l_\theta + v'(l_\theta)) \\
&\quad - \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\bar{\theta}} v'(l_{\widetilde{\theta}}) l_{\widetilde{\theta}} (1 - H_{\widetilde{\theta}}) \left. \frac{d\widehat{w}_\theta^*(l + \mu \widetilde{l}_{\Delta, \theta})}{d\mu} \right|_{\mu=0} d\widetilde{\theta} \\
&\quad + \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\bar{\theta}} \widetilde{g}_{\widetilde{\theta}} h_{\widetilde{\theta}} \int_{\underline{\theta}}^{\widetilde{\theta}} v'(l_{\widehat{\theta}}) l_{\widehat{\theta}} \left. \frac{d\widehat{w}_\theta^*(l + \mu \widetilde{l}_{\Delta, \theta})}{d\mu} \right|_{\mu=0} d\widehat{\theta} d\widetilde{\theta}
\end{aligned} \tag{59}$$

for all θ , where the terms in the first two lines were derived following the procedure detailed in Sections A.1 and A.3, which uses continuity of l and $\widehat{w}$ in θ .⁴⁶ Following the notation introduced in Section 5.1, the expression $d\widehat{w}_\theta^*(l + \mu \widetilde{l}_{\Delta, \theta})/d\mu|_{\mu=0}$ denotes the total derivative of the return to skill in the direction of l_θ , accounting both for the substitution and the directed technical change effects:

$$\left. \frac{d\widehat{w}_\theta^*(l + \mu \widetilde{l}_{\Delta, \theta})}{d\mu} \right|_{\mu=0} = \left. \frac{d\widehat{w}_{\widetilde{\theta}}(l + \mu \widetilde{l}_{\Delta, \theta}, \phi^*(l))}{d\mu} \right|_{\mu=0} + \left. \frac{d\widehat{w}_{\widetilde{\theta}}(l, \phi^*(l + \mu \widetilde{l}_{\Delta, \theta}))}{d\mu} \right|_{\mu=0}.$$

Using integration by parts to solve the double integral in equation (59), the derivative of the welfare function becomes

$$\begin{aligned}
D_{l_\theta} \widehat{W}(l) &= w_\theta h_\theta - v'(l_\theta) h_\theta - (1 - \widetilde{g}_\theta) (v''(l_\theta) l_\theta + v'(l_\theta)) (1 - H_\theta) \widehat{w}_\theta \\
&\quad - \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\bar{\theta}} v'(l_{\widetilde{\theta}}) l_{\widetilde{\theta}} (1 - H_{\widetilde{\theta}}) \left. \frac{d\widehat{w}_\theta^*(l + \mu \widetilde{l}_{\Delta, \theta})}{d\mu} \right|_{\mu=0} d\widetilde{\theta} \\
&\quad + \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\bar{\theta}} \widetilde{g}_{\widetilde{\theta}} (1 - H_{\widetilde{\theta}}) v'(l_{\widetilde{\theta}}) l_{\widetilde{\theta}} \left. \frac{d\widehat{w}_\theta^*(l + \mu \widetilde{l}_{\Delta, \theta})}{d\mu} \right|_{\mu=0} d\widetilde{\theta} \\
&= w_\theta h_\theta - v'(l_\theta) h_\theta - (1 - \widetilde{g}_\theta) (v''(l_\theta) l_\theta + v'(l_\theta)) (1 - H_\theta) \widehat{w}_\theta \\
&\quad - \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\bar{\theta}} (1 - \widetilde{g}_{\widetilde{\theta}}) v'(l_{\widetilde{\theta}}) l_{\widetilde{\theta}} (1 - H_{\widetilde{\theta}}) \left. \frac{d\widehat{w}_\theta^*(l + \mu \widetilde{l}_{\Delta, \theta})}{d\mu} \right|_{\mu=0} d\widetilde{\theta}.
\end{aligned} \tag{60}$$

We now use workers' first-order condition (2) to introduce marginal tax rates into the equation. In particular, condition (2) implies

$$v'(l_\theta) l_\theta = (1 - T'(y_\theta)) y_\theta \tag{61}$$

and

$$v''(l_\theta) l_\theta \widehat{w}_\theta + v'(l_\theta) \widehat{w}_\theta = \left(1 + \frac{1}{e_\theta}\right) (1 - T'(y_\theta)) w'_\theta. \tag{62}$$

⁴⁶The return to skill $\widehat{w}$ is continuous in θ because l is C^1 by hypothesis of Proposition 2.

Using equations (2), (61), and (62) in equation (60), we obtain

$$D_{l_\theta} \widehat{W}(l) = T'(y_\theta) y_\theta h_\theta - (1 - T'(y_\theta)) \left(1 + \frac{1}{e_\theta}\right) (1 - \tilde{g}_\theta)(1 - H_\theta) w'_\theta l_\theta \\ - \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\bar{\theta}} (1 - T'(y_{\tilde{\theta}})) y_{\tilde{\theta}} (1 - \tilde{g}_{\tilde{\theta}})(1 - H_{\tilde{\theta}}) \frac{d\widehat{w}_\theta^*(l + \mu \tilde{l}_{\Delta, \theta})}{d\mu} \Big|_{\mu=0} d\tilde{\theta}.$$

Splitting up the total derivative $d\widehat{w}_\theta^*(l + \mu \tilde{l}_{\Delta, \theta})/d\mu \Big|_{\mu=0}$ into its substitution and directed technical change components, this becomes:

$$D_{l_\theta} \widehat{W}(l) = T'(y_\theta) w_\theta h_\theta - (1 - T'(y_\theta)) \left(1 + \frac{1}{e_\theta}\right) (1 - \tilde{g}_\theta)(1 - H_\theta) w'_\theta \\ - \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\bar{\theta}} (1 - T'(y_{\tilde{\theta}})) y_{\tilde{\theta}} (1 - \tilde{g}_{\tilde{\theta}})(1 - H_{\tilde{\theta}}) \frac{d\widehat{w}_{\tilde{\theta}}(l + \mu \tilde{l}_{\Delta, \theta}, \phi^*(l))}{d\mu} \Big|_{\mu=0} d\tilde{\theta} \\ - \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\bar{\theta}} (1 - T'(y_{\tilde{\theta}})) y_{\tilde{\theta}} (1 - \tilde{g}_{\tilde{\theta}})(1 - H_{\tilde{\theta}}) \frac{d\widehat{w}_{\tilde{\theta}}(l, \phi^*(l + \mu \tilde{l}_{\Delta, \theta}))}{d\mu} \Big|_{\mu=0} d\tilde{\theta}.$$

Equating the derivative to zero, dividing by $1 - T'(y_\theta)$, and rearranging yields:

$$\frac{T'(y_\theta)}{1 - T'(y_\theta)} = \left(1 + \frac{1}{e_\theta}\right) (1 - \tilde{g}_\theta) \frac{1 - H_\theta}{h_\theta} \widehat{w}_\theta \\ - \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\bar{\theta}} \frac{1 - T'(y_{\tilde{\theta}})}{1 - T'(y_\theta)} \frac{1 - H_{\tilde{\theta}}}{h_\theta w_\theta} (1 - \tilde{g}_{\tilde{\theta}}) y_{\tilde{\theta}} \frac{d\widehat{w}_{\tilde{\theta}}(l + \mu \tilde{l}_{\Delta, \theta}, \phi^*(l))}{d\mu} \Big|_{\mu=0} d\tilde{\theta} \\ - \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\bar{\theta}} \frac{1 - T'(y_{\tilde{\theta}})}{1 - T'(y_\theta)} \frac{1 - H_{\tilde{\theta}}}{h_\theta w_\theta} (1 - \tilde{g}_{\tilde{\theta}}) y_{\tilde{\theta}} \frac{d\widehat{w}_{\tilde{\theta}}(l, \phi^*(l + \mu \tilde{l}_{\Delta, \theta}))}{d\mu} \Big|_{\mu=0} d\tilde{\theta}. \quad (63)$$

Finally, let n_w and N_w denote the density and cumulative distribution functions of the distribution of wages and use the change-of-variable $h_\theta = n_{w_\theta} w'_\theta$ to obtain the condition for marginal tax rates from Proposition 2:

$$\frac{T'(y_\theta)}{1 - T'(y_\theta)} = \left(1 + \frac{1}{e_\theta}\right) (1 - \tilde{g}_\theta) \frac{1 - N_{w_\theta}}{n_{w_\theta} w_\theta} \\ - \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\bar{\theta}} \frac{1 - T'(y_{\tilde{\theta}})}{1 - T'(y_\theta)} \frac{1 - H_{\tilde{\theta}}}{h_\theta w_\theta} (1 - \tilde{g}_{\tilde{\theta}}) y_{\tilde{\theta}} \frac{d\widehat{w}_{\tilde{\theta}}(l + \mu \tilde{l}_{\Delta, \theta}, \phi^*(l))}{d\mu} \Big|_{\mu=0} d\tilde{\theta} \\ - \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\bar{\theta}} \frac{1 - T'(y_{\tilde{\theta}})}{1 - T'(y_\theta)} \frac{1 - H_{\tilde{\theta}}}{h_\theta w_\theta} (1 - \tilde{g}_{\tilde{\theta}}) y_{\tilde{\theta}} \frac{d\widehat{w}_{\tilde{\theta}}(l, \phi^*(l + \mu \tilde{l}_{\Delta, \theta}))}{d\mu} \Big|_{\mu=0} d\tilde{\theta}.$$

Part 2. The second part of the proof is to show that $TE_{\underline{\theta}}^* \leq 0$ and $TE_{\bar{\theta}}^* \geq 0$. We only consider $TE_{\bar{\theta}}^*$ because the proof for $TE_{\underline{\theta}}^*$ is analogous.

Consider first the derivative $d\widehat{w}_{\bar{\theta}}(l, \phi^*(l + \mu \tilde{l}_{\Delta, \bar{\theta}}))/d\mu \Big|_{\mu=0}$. It measures the local directed technical

change effect of the labor input change $\tilde{l}_{\Delta,\bar{\theta}}$ (defined in Appendix A.1) on the return to skill $\hat{w}_{\bar{\theta}}$.

For $\theta \leq \bar{\theta} - \Delta$, the labor input change is zero by definition. On $(\bar{\theta} - \Delta, \bar{\theta}]$ it varies in θ according to

$$\frac{1}{\tilde{l}_{\Delta,\bar{\theta},\theta}} \frac{d\tilde{l}_{\Delta,\bar{\theta},\theta}}{d\theta} = \frac{2\Delta}{2\Delta(\theta - \bar{\theta} + \Delta)} = \frac{1}{\theta - \bar{\theta} + \Delta} \geq \frac{1}{\Delta}.$$

Hence, given the optimal labor input l , we can find an $\epsilon > 0$ such that for all $\Delta < \epsilon$ and for all $\theta \in (\bar{\theta} - \Delta, \bar{\theta})$:

$$\frac{1}{\tilde{l}_{\Delta,\bar{\theta},\theta}} \frac{d\tilde{l}_{\Delta,\bar{\theta},\theta}}{d\theta} \geq \frac{1}{l_{\theta}} \frac{dl_{\theta}}{d\theta}.^{47}$$

So, for $\Delta < \epsilon$, the relative labor input change $\tilde{l}_{\Delta,\bar{\theta},\theta}/l_{\theta}$ increases in θ . Thus, by Lemma 1 (weak relative bias), we obtain that

$$\left. \frac{d\hat{w}_{\bar{\theta}}(l, \phi^*(l + \mu\tilde{l}_{\Delta,\bar{\theta}}))}{d\mu} \right|_{\mu=0} \geq 0$$

for all $\bar{\theta}$ if $\Delta < \epsilon$. Hence, for $\Delta < \epsilon$,

$$\int_{\underline{\theta}}^{\bar{\theta}} \frac{1 - T'(y_{\bar{\theta}})}{1 - T'(y_{\bar{\theta}})} \frac{1 - H_{\bar{\theta}}}{h_{\bar{\theta}}w_{\bar{\theta}}} (1 - \tilde{g}_{\bar{\theta}})y_{\bar{\theta}} \left. \frac{d\hat{w}_{\bar{\theta}}(l, \phi^*(l + \mu\tilde{l}_{\Delta,\bar{\theta}}))}{d\mu} \right|_{\mu=0} d\bar{\theta} \geq 0$$

and therefore

$$TE_{\bar{\theta}}^* = \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\bar{\theta}} \frac{1 - T'(y_{\bar{\theta}})}{1 - T'(y_{\bar{\theta}})} \frac{1 - H_{\bar{\theta}}}{h_{\bar{\theta}}w_{\bar{\theta}}} (1 - \tilde{g}_{\bar{\theta}})y_{\bar{\theta}} \left. \frac{d\hat{w}_{\bar{\theta}}(l, \phi^*(l + \mu\tilde{l}_{\Delta,\bar{\theta}}))}{d\mu} \right|_{\mu=0} d\bar{\theta} \geq 0.$$

Part 3. Finally, we show that, if there is strong relative bias, $SE_{\underline{\theta}}^* + TE_{\underline{\theta}}^* \leq 0$ and $SE_{\bar{\theta}}^* + TE_{\bar{\theta}}^* \geq 0$. Again, the proof is analogous for both statements and we focus on the latter.

Note first that $SE_{\bar{\theta}}^* + TE_{\bar{\theta}}^*$ can be written as

$$\lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\bar{\theta}} \frac{1 - T'(y_{\theta})}{1 - T'(y_{\bar{\theta}})} \frac{1 - H_{\theta}}{h_{\bar{\theta}}w_{\bar{\theta}}} (1 - \tilde{g}_{\theta})y_{\theta} \left. \frac{d\hat{w}_{\bar{\theta}}^*(l + \mu\tilde{l}_{\Delta,\bar{\theta}})}{d\mu} \right|_{\mu=0} d\theta.$$

We have already shown in part 2 that, for sufficiently small Δ , the relative labor input change $\tilde{l}_{\Delta,\bar{\theta},\theta}/l_{\theta}$ increases in θ . By definition of strong relative bias (see equation (17)), this implies

$$\left. \frac{d\hat{w}_{\bar{\theta}}^*(l + \mu\tilde{l}_{\Delta,\bar{\theta}})}{d\mu} \right|_{\mu=0} \geq 0$$

for all θ and small Δ . Analogously to part 2, it follows that $SE_{\bar{\theta}}^* + TE_{\bar{\theta}}^* \geq 0$.

⁴⁷Here we use that, by hypothesis, $\limsup_{\theta \rightarrow \bar{\theta}} l'_{\theta} < \infty$.

A.11 Proof of Corollary 2: Optimal Tax Formula on the Pareto Tail

The limit expression for the optimal marginal tax rate follows immediately from the convergence assumptions in Corollary 2. So, the only statements in need of a proof are the sign restrictions on $\overline{TE}$ and $\overline{TE} + \overline{SE}$.

Part 1. I start with the proof of $\overline{TE} \geq 0$. First, note that

$$\begin{aligned} & \int_{\theta}^{\bar{\theta}} (1 - T'(y_{\tilde{\theta}})) h_{\tilde{\theta}} w_{\tilde{\theta}} l_{\tilde{\theta}} TE_{\tilde{\theta}}^* d\tilde{\theta} \\ &= \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\bar{\theta}} (1 - T'(y_{\hat{\theta}}))(1 - H_{\hat{\theta}})(1 - \tilde{g}_{\hat{\theta}}) y_{\hat{\theta}} \int_{\theta}^{\bar{\theta}} l_{\tilde{\theta}} \frac{d\hat{w}_{\hat{\theta}}(l, \phi^*(l + \mu \tilde{l}_{\Delta, \hat{\theta}}))}{d\mu} \Big|_{\mu=0} d\tilde{\theta} d\hat{\theta} \\ &= \int_{\underline{\theta}}^{\bar{\theta}} (1 - T'(y_{\hat{\theta}}))(1 - H_{\hat{\theta}})(1 - \tilde{g}_{\hat{\theta}}) y_{\hat{\theta}} \frac{d\hat{w}_{\hat{\theta}}(l, \phi^*(l + \mu dl))}{d\mu} \Big|_{\mu=0} d\hat{\theta}, \end{aligned}$$

where the labor input change dl is given by

$$dl_{\tilde{\theta}} = \begin{cases} 0 & \text{if } \tilde{\theta} < \theta \\ l_{\tilde{\theta}} & \text{if } \theta \leq \tilde{\theta}. \end{cases}$$

Since $dl_{\tilde{\theta}}/l_{\tilde{\theta}}$ increases in $\tilde{\theta}$, that is, dl is an increase in relative skill supply, we have by Lemma 1:

$$\frac{d\hat{w}_{\hat{\theta}}(l, \phi^*(l + \mu dl))}{d\mu} \Big|_{\mu=0} \geq 0$$

for all $\hat{\theta}$ and hence:

$$\int_{\theta}^{\bar{\theta}} (1 - T'(y_{\tilde{\theta}})) h_{\tilde{\theta}} w_{\tilde{\theta}} l_{\tilde{\theta}} TE_{\tilde{\theta}}^* d\tilde{\theta} \geq 0$$

for all θ . Therefore, we obtain the following result:

$$\begin{aligned} 0 &\leq \lim_{\theta \rightarrow \bar{\theta}} \frac{\int_{\theta}^{\bar{\theta}} (1 - T'(y_{\tilde{\theta}})) h_{\tilde{\theta}} w_{\tilde{\theta}} l_{\tilde{\theta}} TE_{\tilde{\theta}}^* d\tilde{\theta}}{\int_{\theta}^{\bar{\theta}} (1 - T'(y_{\tilde{\theta}})) h_{\tilde{\theta}} w_{\tilde{\theta}} l_{\tilde{\theta}} d\tilde{\theta}} \\ &= \lim_{\theta \rightarrow \bar{\theta}} \frac{-(1 - T'(y_{\theta})) h_{\theta} w_{\theta} l_{\theta} TE_{\theta}^*}{-(1 - T'(y_{\theta})) h_{\theta} w_{\theta} l_{\theta}} \\ &= \lim_{\theta \rightarrow \bar{\theta}} TE_{\theta}^*, \end{aligned}$$

where the second line uses L'Hôpital's rule.

Part 2. The proof that $\overline{TE} + \overline{SE} \geq 0$ under strong relative bias proceeds along the same lines as part 1. In particular, we can show analogously to part 1 that, if there is strong bias,

$$\int_{\theta}^{\bar{\theta}} (1 - T'(y_{\tilde{\theta}})) h_{\tilde{\theta}} w_{\tilde{\theta}} l_{\tilde{\theta}} (TE_{\tilde{\theta}}^* + SE_{\tilde{\theta}}^*) d\tilde{\theta} \geq 0$$

for all θ . It follows, as in part 1, that

$$\lim_{\theta \rightarrow \bar{\theta}} (TE_{\theta}^* + SE_{\theta}^*) \geq 0 .$$

A.12 Proof of Proposition 3: Optimal Tax Formula for the CES Case

To derive expression (35) for optimal marginal tax rates in the CES case, I start by specializing the terms TE_{θ}^* and SE_{θ}^* .

Lemma 8. *Suppose the conditions of Proposition 2 are satisfied and F and Φ take the CES form introduced in Section 3.6. Then, the terms TE_{θ}^* and SE_{θ}^* take the following form for every θ :*

$$SE_{\theta}^* = (1 - g_{\theta})\gamma^{CES} - \left(1 + \frac{1}{e_{\theta}}\right) \frac{1 - H_{\theta}}{h_{\theta}} (1 - \tilde{g}_{\theta}) \hat{l}_{\theta} \gamma^{CES} \quad (64)$$

$$TE_{\theta}^* = (1 - g_{\theta})\rho^{CES} - \left(1 + \frac{1}{e_{\theta}}\right) \frac{1 - H_{\theta}}{h_{\theta}} (1 - \tilde{g}_{\theta}) \hat{l}_{\theta} \rho^{CES} . \quad (65)$$

Proof. I focus on the expression for TE_{θ}^* , because the derivation of SE_{θ}^* is analogous, with γ^{CES} in the place of ρ^{CES} .

The central step is to obtain an expression for the directed technical change derivative of the return to skill $\hat{w}_{\tilde{\theta}}$ that features in TE_{θ}^* . From equation (15) we obtain

$$\hat{w}_{\tilde{\theta}} = \frac{\sigma - 1}{\sigma} \hat{\kappa}_{\tilde{\theta}} + \frac{\sigma - 1}{\sigma} \hat{\phi}_{\tilde{\theta}} - \frac{1}{\sigma} \hat{l}_{\tilde{\theta}} - \frac{1}{\sigma} \hat{h}_{\tilde{\theta}} \quad (66)$$

and from equation (45):

$$\hat{\phi}_{\tilde{\theta}}^* = \frac{\sigma - 1}{(\delta - 1)\sigma + 1} \left(\hat{\kappa}_{\tilde{\theta}} + \hat{l}_{\tilde{\theta}} + \hat{h}_{\tilde{\theta}} \right) . \quad (67)$$

Hence, the partial effect of the perturbation $\tilde{l}_{\Delta,\theta}$ on $\hat{w}_{\tilde{\theta}}$ is

$$\begin{aligned} \left. \frac{d\hat{w}_{\tilde{\theta}}(l + \mu\tilde{l}_{\Delta,\theta}, \phi^*(l))}{d\mu} \right|_{\mu=0} &= -\frac{1}{\sigma} \frac{d}{d\mu} \left(\widehat{l_{\tilde{\theta}} + \mu\tilde{l}_{\Delta,\theta,\tilde{\theta}}} \right) \Big|_{\mu=0} \\ &= -\frac{1}{\sigma} \frac{d}{d\mu} \frac{l'_{\tilde{\theta}} + \mu\tilde{l}'_{\Delta,\theta,\tilde{\theta}}}{l_{\tilde{\theta}} + \mu\tilde{l}_{\Delta,\theta,\tilde{\theta}}} \Big|_{\mu=0} \\ &= \gamma^{CES} \left(\frac{\tilde{l}'_{\Delta,\theta,\tilde{\theta}}}{l_{\tilde{\theta}}} - \hat{l}_{\tilde{\theta}} \frac{\tilde{l}_{\Delta,\theta,\tilde{\theta}}}{l_{\tilde{\theta}}} \right) . \end{aligned}$$

Analogously, the directed technical change effect is given by

$$\begin{aligned} \left. \frac{d\widehat{w}_{\tilde{\theta}}(l, \phi^*(l + \mu\tilde{l}_{\Delta, \theta}))}{d\mu} \right|_{\mu=0} &= \frac{(\sigma - 1)^2}{(\delta - 1)\sigma^2 + \sigma} \frac{d}{d\mu} (\widehat{l_{\tilde{\theta}} + \mu\tilde{l}_{\Delta, \theta, \tilde{\theta}}}) \Big|_{\mu=0} \\ &= \rho^{CES} \left(\frac{\tilde{l}'_{\Delta, \theta, \tilde{\theta}}}{\tilde{l}_{\tilde{\theta}}} - \widehat{l_{\tilde{\theta}}} \frac{\tilde{l}_{\Delta, \theta, \tilde{\theta}}}{\tilde{l}_{\tilde{\theta}}} \right). \end{aligned}$$

Using the last expression and the definition of $\tilde{l}_{\Delta, \theta}$, the term TE_{θ}^* becomes

$$\begin{aligned} TE_{\theta}^* &= \lim_{\Delta \rightarrow 0} \frac{1}{\Delta^2} \int_{\theta-\Delta}^{\theta} \frac{1 - T'(y_{\tilde{\theta}})}{1 - T'(y_{\theta})} \frac{1 - H_{\tilde{\theta}}}{h_{\theta} w_{\theta}} (1 - \tilde{g}_{\tilde{\theta}}) w_{\tilde{\theta}} \rho^{CES} \left(1 - \widehat{l_{\tilde{\theta}}}(\tilde{\theta} - \theta + \Delta) \right) d\tilde{\theta} \\ &\quad + \frac{1}{\Delta^2} \int_{\theta}^{\theta+\Delta} \frac{1 - T'(y_{\tilde{\theta}})}{1 - T'(y_{\theta})} \frac{1 - H_{\tilde{\theta}}}{h_{\theta} w_{\theta}} (1 - \tilde{g}_{\tilde{\theta}}) w_{\tilde{\theta}} \rho^{CES} \left(-1 - \widehat{l_{\tilde{\theta}}}(\theta - \tilde{\theta} + \Delta) \right) d\tilde{\theta}. \end{aligned}$$

Applying L'Hôpital's rule yields:

$$\begin{aligned} TE_{\theta}^* &= \lim_{\Delta \rightarrow 0} \frac{1}{2\Delta} \frac{1 - T'(y_{\theta-\Delta})}{1 - T'(y_{\theta})} \frac{1 - H_{\theta-\Delta}}{h_{\theta} w_{\theta}} (1 - \tilde{g}_{\theta-\Delta}) w_{\theta-\Delta} \rho^{CES} \\ &\quad - \frac{1}{2\Delta} \int_{\theta-\Delta}^{\theta} \frac{1 - T'(y_{\tilde{\theta}})}{1 - T'(y_{\theta})} \frac{1 - H_{\tilde{\theta}}}{h_{\theta} w_{\theta}} (1 - \tilde{g}_{\tilde{\theta}}) w_{\tilde{\theta}} \rho^{CES} \widehat{l_{\tilde{\theta}}} d\tilde{\theta} \\ &\quad - \frac{1}{2\Delta} \frac{1 - T'(y_{\theta+\Delta})}{1 - T'(y_{\theta})} \frac{1 - H_{\theta+\Delta}}{h_{\theta} w_{\theta}} (1 - \tilde{g}_{\theta+\Delta}) w_{\theta+\Delta} \rho^{CES} \\ &\quad - \frac{1}{2\Delta} \int_{\theta}^{\theta+\Delta} \frac{1 - T'(y_{\tilde{\theta}})}{1 - T'(y_{\theta})} \frac{1 - H_{\tilde{\theta}}}{h_{\theta} w_{\theta}} (1 - \tilde{g}_{\tilde{\theta}}) w_{\tilde{\theta}} \rho^{CES} \widehat{l_{\tilde{\theta}}} d\tilde{\theta}. \end{aligned}$$

Rearranging and replacing marginal retention rates by workers' first-order condition (2), we obtain:

$$\begin{aligned} TE_{\theta}^* &= \lim_{\Delta \rightarrow 0} \frac{1}{2\Delta} \rho^{CES} \left[\frac{v'(l_{\theta-\Delta})}{v'(l_{\theta})} \frac{1 - H_{\theta-\Delta}}{h_{\theta}} (1 - \tilde{g}_{\theta-\Delta}) - \frac{v'(l_{\theta+\Delta})}{v'(l_{\theta})} \frac{1 - H_{\theta+\Delta}}{h_{\theta}} (1 - \tilde{g}_{\theta+\Delta}) \right] \\ &\quad - \frac{1}{2\Delta} \int_{\theta-\Delta}^{\theta} \frac{v'(l_{\tilde{\theta}})}{v'(l_{\theta})} \frac{1 - H_{\tilde{\theta}}}{h_{\theta}} (1 - \tilde{g}_{\tilde{\theta}}) \rho^{CES} \widehat{l_{\tilde{\theta}}} d\tilde{\theta} \\ &\quad - \frac{1}{2\Delta} \int_{\theta}^{\theta+\Delta} \frac{v'(l_{\tilde{\theta}})}{v'(l_{\theta})} \frac{1 - H_{\tilde{\theta}}}{h_{\theta}} (1 - \tilde{g}_{\tilde{\theta}}) \rho^{CES} \widehat{l_{\tilde{\theta}}} d\tilde{\theta}. \end{aligned}$$

Next, we apply L'Hôpital's rule a second time and obtain:

$$TE_{\theta}^* = -\frac{v''(l_{\theta})}{v'(l_{\theta})} \frac{1 - H_{\theta}}{h_{\theta}} (1 - \tilde{g}_{\theta}) l'_{\theta} \rho^{CES} + (1 - \tilde{g}_{\theta}) \rho^{CES} + \frac{1 - H_{\theta}}{h_{\theta}} \tilde{g}'_{\theta} \rho^{CES} - \frac{1 - H_{\theta}}{h_{\theta}} (1 - \tilde{g}_{\theta}) \widehat{l_{\theta}} \rho^{CES}.$$

Using the definition of the elasticity e_{θ} yields:

$$TE_{\theta}^* = -\frac{1}{e_{\theta}} \frac{1 - H_{\theta}}{h_{\theta}} (1 - \tilde{g}_{\theta}) \widehat{l_{\theta}} \rho^{CES} + (1 - \tilde{g}_{\theta}) \rho^{CES} + \frac{1 - H_{\theta}}{h_{\theta}} \tilde{g}'_{\theta} \rho^{CES} - \frac{1 - H_{\theta}}{h_{\theta}} (1 - \tilde{g}_{\theta}) \widehat{l_{\theta}} \rho^{CES}.$$

Finally, it is straightforward to show that

$$\tilde{g}'_\theta = (\tilde{g}_\theta - g_\theta) \frac{h_\theta}{1 - H_\theta}.$$

Inserting this into the previous expression for TE_θ^* , we obtain:

$$TE_\theta^* = -\frac{1}{e_\theta} \frac{1 - H_\theta}{h_\theta} (1 - \tilde{g}_\theta) \hat{l}_\theta \rho^{CES} + (1 - \tilde{g}_\theta) \rho^{CES} + (\tilde{g}_\theta - g_\theta) \rho^{CES} - \frac{1 - H_\theta}{h_\theta} (1 - \tilde{g}_\theta) \hat{l}_\theta \rho^{CES}$$

and after rearranging:

$$TE_\theta^* = (1 - g_\theta) \rho^{CES} - \left(1 + \frac{1}{e_\theta}\right) \frac{1 - H_\theta}{h_\theta} (1 - \tilde{g}_\theta) \hat{l}_\theta \rho^{CES},$$

which is the desired expression. $\square$

With the expressions from Lemma 8, we are now in a position to derive equation (35). We start with equation (63) from the proof of Proposition 2 and replace SE_θ^* and TE_θ^* by the expressions from Lemma 8. This yields:

$$\frac{T'(y_\theta)}{1 - T'(y_\theta)} = \left(1 + \frac{1}{e_\theta}\right) \frac{1 - H_\theta}{h_\theta} (1 - \tilde{g}_\theta) \hat{w}_\theta \quad (68)$$

$$+ (\gamma^{CES} + \rho^{CES}) \left[(1 - g_\theta) - \left(1 + \frac{1}{e_\theta}\right) \frac{1 - H_\theta}{h_\theta} (1 - \tilde{g}_\theta) \hat{l}_\theta \right]$$

$$= \left(1 + \frac{1}{e_\theta}\right) \frac{1 - H_\theta}{h_\theta} (1 - \tilde{g}_\theta) \left(\hat{w}_\theta - (\gamma^{CES} + \rho^{CES}) \hat{l}_\theta \right) \quad (69)$$

$$+ \gamma^{CES} (1 - g_\theta) + \rho^{CES} (1 - g_\theta). \quad (70)$$

The return to skill $\hat{w}_\theta$ can be computed from equations (66) and (67) as

$$\hat{w}_\theta = (1 + \gamma^{CES} + \rho^{CES}) \hat{\kappa}_\theta + (\gamma^{CES} + \rho^{CES}) \hat{h}_\theta + (\gamma^{CES} + \rho^{CES}) \hat{l}_\theta.$$

Using this in the previous expression for marginal tax rates yields

$$\frac{T'(y_\theta)}{1 - T'(y_\theta)} = \left(1 + \frac{1}{e_\theta}\right) \frac{1 - H_\theta}{h_\theta} (1 - \tilde{g}_\theta) \left[(1 + \gamma^{CES} + \rho^{CES}) \hat{\kappa}_\theta + (\gamma^{CES} + \rho^{CES}) \hat{h}_\theta \right]$$

$$+ \gamma^{CES} (1 - g_\theta) + \rho^{CES} (1 - g_\theta).$$

Now we use the definition of β ,

$$\beta_\theta := \kappa_\theta^{1 + \gamma^{CES} + \rho^{CES}} h_\theta^{\gamma^{CES} + \rho^{CES}},$$

to note that

$$(1 + \gamma^{CES} + \rho^{CES}) \hat{\kappa}_\theta + (\gamma^{CES} + \rho^{CES}) \hat{h}_\theta = \hat{\beta}_\theta$$

and hence

$$\frac{T'(y_\theta)}{1 - T'(y_\theta)} = \left(1 + \frac{1}{e_\theta}\right) \frac{1 - H_\theta}{h_\theta} (1 - \tilde{g}_\theta) \hat{\beta}_\theta + \gamma^{CES} (1 - g_\theta) + \rho^{CES} (1 - g_\theta) .$$

Finally, with the change-of-variable $h_\theta = b_{\beta_\theta} \beta'_\theta$, we obtain equation (35):

$$\frac{T'(y_\theta)}{1 - T'(y_\theta)} = \left(1 + \frac{1}{e_\theta}\right) \frac{1 - B_{\beta_\theta}}{b_{\beta_\theta} \beta_\theta} (1 - \tilde{g}_\theta) + \gamma^{CES} (1 - g_\theta) + \rho^{CES} (1 - g_\theta) .$$

A.13 Proof Corollary 3: Optimal Tax Formula on the Pareto Tail for the CES Case

Corollary 3 starts from the assumption that, under some initial tax $\bar{T}$ with constant marginal top tax rate, the income distribution has the Pareto property (i.e., its inverse hazard ratio is asymptotically constant). We trace the Pareto property of the income distribution back to the distribution of the exogenous inequality measure β . Inserting this distribution into the optimal tax formula (35) from Proposition 3 then yields the desired result.

First, by two changes-of-variable we obtain

$$\frac{1 - B_{\beta_\theta}}{b_{\beta_\theta} \beta_\theta} = \frac{1 - H_\theta}{h_\theta} \hat{\beta}_\theta = \frac{1 - M_{y_\theta}}{m_{y_\theta} y_\theta} \frac{\hat{\beta}_\theta}{\hat{y}_\theta} , \quad (71)$$

where all incomes are assessed at the given tax $\bar{T}$. Now we use

$$\hat{y}_\theta = \hat{w}_\theta + \hat{l}_\theta = (1 + \epsilon_\theta^w) \hat{w}_\theta$$

to express the growth rate of income as a function of the growth rate of wages, again assessing all endogenous variables at equilibrium under the given tax $\bar{T}$. For $\hat{\beta}_\theta$ we obtain

$$\hat{\beta}_\theta = \hat{w}_\theta - (\gamma^{CES} + \rho^{CES}) \hat{l}_\theta = \left(1 - (\gamma^{CES} + \rho^{CES}) \epsilon_\theta^w\right) \hat{w}_\theta .$$

It follows that

$$\frac{\hat{\beta}_\theta}{\hat{y}_\theta} = \frac{1 - (\gamma^{CES} + \rho^{CES}) \epsilon_\theta^w}{1 + \epsilon_\theta^w}$$

and, with equation (71),

$$\frac{1 - B_{\beta_\theta}}{b_{\beta_\theta} \beta_\theta} = \frac{1 - M_{y_\theta}}{m_{y_\theta} y_\theta} \frac{1 - (\gamma^{CES} + \rho^{CES}) \epsilon_\theta^w}{1 + \epsilon_\theta^w} , \quad (72)$$

where incomes and the labor supply elasticity ϵ_θ^w are assessed under the tax $\bar{T}$. In particular, using equation (40) from Appendix A.2 on the relationship between ϵ_θ^w and e_θ , we obtain that

$$\epsilon_\theta^w = \frac{(1 - P_{\bar{T}}(y_\theta)) e_\theta}{1 + e_\theta P_{\bar{T}}(y_\theta)} .$$

Since the tax $\bar{T}$ features a constant top tax rate, we have $\lim_{\theta \rightarrow \bar{\theta}} P_{\bar{T}}(y_{\theta}) = 0$ and hence

$$\lim_{\theta \rightarrow \bar{\theta}} \epsilon_{\theta}^w = e_{\bar{\theta}} = e ,$$

where the last equality reflects the assumption that the disutility of labor is isoelastic. Moreover, we know by hypothesis that

$$\lim_{\theta \rightarrow \bar{\theta}} \frac{1 - M_{y_{\theta}}}{m_{y_{\theta}} y_{\theta}} = \frac{1}{a} .$$

Combining these limits, we obtain

$$\lim_{\theta \rightarrow \bar{\theta}} \frac{1 - B_{\beta_{\theta}}}{b_{\beta_{\theta}} \beta_{\theta}} = \lim_{\theta \rightarrow \bar{\theta}} \frac{1 - M_{y_{\theta}}}{m_{y_{\theta}} y_{\theta}} \frac{1 - (\gamma^{CES} + \rho^{CES}) \epsilon_{\theta}^w}{1 + \epsilon_{\theta}^w} = \frac{1 - (\gamma^{CES} + \rho^{CES}) e}{a(1 + e)} .$$

In words, from the observed Pareto tail of the income distribution under tax $\bar{T}$ we can infer that the exogenous inequality measure β must also have a Pareto tail with tail parameter given by the previous equation. Using this parameter in the optimal tax equation (35) from Proposition 3, we obtain the following expression for the optimal marginal tax rate in the upper tail of the income distribution:

$$\begin{aligned} \lim_{\theta \rightarrow \bar{\theta}} \frac{T'(y_{\theta})}{1 - T'(y_{\theta})} &= \left(1 + \frac{1}{e}\right) \frac{1 - (\gamma^{CES} + \rho^{CES}) e}{a(1 + e)} (1 - g^{top}) + \gamma^{CES} (1 - g^{top}) + \rho^{CES} (1 - g^{top}) \\ &= \frac{1 - g^{top}}{ae} + \frac{a - 1}{a} \gamma^{CES} (1 - g^{top}) + \frac{a - 1}{a} \rho^{CES} (1 - g^{top}) , \end{aligned}$$

where g^{top} is the asymptotic welfare weight defined in the Corollary.

A.14 Tax Formula for the Exogenous Technology Planner

As described in the main text, the exogenous technology planner believes that the economy works according to all equilibrium conditions from Section 3.5 with the exception of the condition for the equilibrium technology (11). Instead of following equation (11), the exogenous technology planner believes that technology remains fixed at its equilibrium value under a given initial tax $\bar{T}$, $\phi^*(l(\bar{T}))$. The idea is that the planner observes the economy under the tax $\bar{T}$ when computing optimal taxes and believes technology to be exogenous.

The exogenous technology planner's optimal tax $T_{\bar{T}}^{ex}$ then satisfies the conditions provided by the following Proposition.

Proposition 4. *Suppose the conditions of Proposition 2 are satisfied and F and Φ take the CES form introduced in Section 3.6. Suppose that all equilibrium variables are determined according to conditions (2), (1), (10), and (12), plus the (exogenous) technology equation*

$$\phi^*(l) = \phi^*(l(\bar{T})) = \operatorname{argmax}_{\phi \in \Phi} F(l(\bar{T}), \phi) \quad \forall l ,$$

where $\bar{T}$ is a given initial tax function.

Then, at every type θ , the exogenous technology planner's preferred tax $T_{\bar{T}}^{ex}$ satisfies the following conditions.

$$\frac{T_{\bar{T}}^{ex'}(y_\theta)}{1 - T_{\bar{T}}^{ex'}(y_\theta)} = \left(1 + \frac{1}{e_\theta}\right) \frac{1 - \bar{B}_{\bar{\beta}_\theta}}{\bar{b}_{\bar{\beta}_\theta} \bar{\beta}_\theta} (1 - \tilde{g}_\theta) + \gamma^{CES} (1 - g_\theta) ,$$

where all variables satisfy the conditions listed above under the tax $T_{\bar{T}}^{ex}$; the function $\bar{\beta} : \theta \mapsto \bar{\beta}_\theta$ is given by

$$\bar{\beta}_\theta := \kappa_\theta^{1+\gamma^{CES}} h_\theta^{\gamma^{CES}} (\phi^*(\bar{T}))^{1+\gamma^{CES}} \quad \forall \theta ;$$

and $\bar{B}$ and $\bar{b}$ are the cumulative distribution and the density function of $\bar{\beta}$.

Proof. It can be verified that all steps in the proof of Proposition 2 hold for the case of the exogenous technology planner when imposing

$$\left. \frac{d\hat{w}_\theta(l, \phi^*(l + \mu \tilde{l}_{\Delta, \theta}))}{d\mu} \right|_{\mu=0} = 0 .$$

With this constraint, we can derive a counterpart to equation (63) for the exogenous technology planner:

$$\frac{T'(y_\theta)}{1 - T'(y_\theta)} = \left(1 + \frac{1}{e_\theta}\right) (1 - \tilde{g}_\theta) \frac{1 - H_\theta}{h_\theta} \hat{w}_\theta - SE_\theta^* .$$

Using Lemma 8 to replace SE_θ^* , we obtain:

$$\begin{aligned} \frac{T'(y_\theta)}{1 - T'(y_\theta)} &= \left(1 + \frac{1}{e_\theta}\right) \frac{1 - H_\theta}{h_\theta} (1 - \tilde{g}_\theta) \hat{w}_\theta \\ &\quad + \gamma^{CES} \left[(1 - g_\theta) - \left(1 + \frac{1}{e_\theta}\right) \frac{1 - H_\theta}{h_\theta} (1 - \tilde{g}_\theta) \hat{l}_\theta \right] \\ &= \left(1 + \frac{1}{e_\theta}\right) \frac{1 - H_\theta}{h_\theta} (1 - \tilde{g}_\theta) (\hat{w}_\theta - \gamma^{CES} \hat{l}_\theta) + \gamma^{CES} (1 - g_\theta) . \end{aligned}$$

From the perspective of the exogenous technology planner, the return to skill is now given by

$$\hat{w}_\theta = (1 + \gamma^{CES}) \hat{\kappa}_\theta + (1 + \gamma^{CES}) \hat{\phi}_\theta^*(\bar{T}) + \gamma^{CES} \hat{h}_\theta + \gamma^{CES} \hat{l}_\theta ,$$

where $\hat{\phi}_\theta^*(\bar{T})$ denotes the growth rate of technology that prevails in equilibrium under the initial tax system $\bar{T}$. Using this in the previous expression for the exogenous technology planner's optimal tax rates and applying the definition of $\bar{\beta}$ yields:

$$\frac{T'(y_\theta)}{1 - T'(y_\theta)} = \left(1 + \frac{1}{e_\theta}\right) \frac{1 - H_\theta}{h_\theta} (1 - \tilde{g}_\theta) \hat{\beta}_\theta + \gamma^{CES} (1 - g_\theta) .$$

With the change of variable $h_\theta = \bar{b}_{\bar{\beta}_\theta} \bar{\beta}_\theta'$, we obtain equation (37),

$$\frac{T'(y_\theta)}{1 - T'(y_\theta)} = \left(1 + \frac{1}{e_\theta}\right) \frac{1 - \bar{B}_{\bar{\beta}_\theta}}{\bar{b}_{\bar{\beta}_\theta} \bar{\beta}_\theta} (1 - \tilde{g}_\theta) + \gamma^{CES} (1 - g_\theta),$$

which completes the proof. $\square$

For the marginal tax rate in the upper Pareto tail of the income distribution, the exogenous technology planner obtains the following characterization.

Corollary 4. *Suppose the conditions of Proposition 2 are satisfied and F and Φ take the CES form introduced in Section 3.6. Suppose equilibrium variables were determined by conditions (2), (1), (10), and (12), and by the (exogenous) technology equation*

$$\phi^*(l) = \phi^*(l(\bar{T})) = \operatorname{argmax}_{\phi \in \Phi} F(l(\bar{T}), \phi) \quad \forall l,$$

where $\bar{T}$ is a given initial tax function with $\bar{T}'(y) = \tau^{top}$ for all $y \geq \tilde{y}$ and some threshold $\tilde{y}$.

Moreover, assume that under the tax $\bar{T}$ the income distribution satisfies

$$\lim_{\theta \rightarrow \bar{\theta}} \frac{1 - M_{y_\theta}}{m_{y_\theta} y_\theta} = \frac{1}{a}$$

for some $a > 1$. Finally, let the disutility of labor be isoelastic with $e_\theta = e$ for all θ , and let welfare weights satisfy

$$\lim_{\theta \rightarrow \bar{\theta}} g_\theta = g^{top}$$

at the exogenous technology planner's preferred tax.

Then, the exogenous technology planner's preferred tax $T_{\bar{T}}^{ex}$ satisfies

$$\lim_{\theta \rightarrow \bar{\theta}} \frac{T_{\bar{T}}^{ex'}(y_\theta)}{1 - T_{\bar{T}}^{ex'}(y_\theta)} = \frac{1 - g^{top}}{ae} + \frac{a - 1}{a} \gamma^{CES} (1 - g^{top}).$$

Proof. The corollary can be derived from Proposition 4 in the same way as Corollary 3 is derived from Proposition 3. In particular, consider first the implications of the Pareto shape of the income distribution under tax $\bar{T}$ for the exogenous inequality measure $\bar{\beta}$. Using changes-of-variable, we obtain

$$\frac{1 - \bar{B}_{\bar{\beta}_\theta}}{\bar{b}_{\bar{\beta}_\theta} \bar{\beta}_\theta} = \frac{1 - H_\theta}{h_\theta} \hat{\bar{\beta}}_\theta = \frac{1 - M_{y_\theta}}{m_{y_\theta} y_\theta} \frac{\hat{\bar{\beta}}_\theta}{\hat{y}_\theta}, \quad (73)$$

where all incomes are assessed at the given tax $\bar{T}$. The exogenous technology planner's measure of exogenous inequality $\bar{\beta}$ now evolves over the type space according to

$$\hat{\bar{\beta}}_\theta = \hat{w}_\theta - \gamma^{CES} \hat{l}_\theta = (1 - \gamma^{CES} \epsilon_\theta^w) \hat{w}_\theta,$$

while

$$\hat{y}_\theta = \hat{w}_\theta + \hat{l}_\theta = (1 + \epsilon_\theta^w) \hat{w}_\theta ,$$

where all endogenous variables are assessed at equilibrium under the given tax $\bar{T}$.⁴⁸ Combining the previous expressions, we find that

$$\frac{\hat{\beta}_\theta}{\hat{y}_\theta} = \frac{1 - \gamma^{CES} \epsilon_\theta^w}{1 + \epsilon_\theta^w}$$

and, with equation (73),

$$\frac{1 - \bar{B}_{\bar{\beta}_\theta}}{\bar{b}_{\bar{\beta}_\theta} \bar{\beta}_\theta} = \frac{1 - M_{y_\theta}}{m_{y_\theta} y_\theta} \frac{1 - \gamma^{CES} \epsilon_\theta^w}{1 + \epsilon_\theta^w} , \quad (74)$$

where incomes and the labor supply elasticity ϵ_θ^w are assessed under the tax $\bar{T}$. Inserting this expression into equation (37) for optimal marginal tax rates computed by the exogenous technology planner and taking limits yields:⁴⁹

$$\begin{aligned} \lim_{\theta \rightarrow \bar{\theta}} \frac{T_{\bar{T}}^{ex'}(y_\theta)}{1 - T_{\bar{T}}^{ex'}(y_\theta)} &= \left(1 + \frac{1}{e}\right) \frac{1 - \gamma^{CES} e}{a(1 + e)} (1 - g^{top}) + \gamma^{CES} (1 - g^{top}) \\ &= \frac{1 - g^{top}}{ae} + \frac{a - 1}{a} \gamma^{CES} (1 - g^{top}) . \end{aligned}$$

□

A.15 Comparison between Optimal Taxes and Exogenous Technology Planner

As discussed in the main text, the conditions for optimal taxes and for the exogenous technology planner's preferred taxes feature two differences. First, the exogenous technology planner neglects the progressive term $\rho^{CES}(1 - g_\theta)$. Second, he uses $\bar{\beta}$ instead of β to measure the degree of exogenous inequality in the economy. Here, I show that the exogenous technology planner thereby overestimates the degree of exogenous inequality: the function $\bar{\beta}$ progresses at a higher rate in θ than the function β , such that inequality (38) from the main text holds.

First, let $\bar{l}$ be the equilibrium labor input under the initial tax system $\bar{T}$. Then, by construction, the growth rates of β and $\bar{\beta}$ must satisfy

$$\hat{\beta}_\theta = \hat{w}_\theta(\bar{l}, \phi^*(\bar{l})) - (\gamma^{CES} + \rho^{CES}) \hat{l}_\theta$$

and

$$\hat{\bar{\beta}}_\theta = \hat{w}_\theta(\bar{l}, \phi^*(\bar{l})) - \gamma^{CES} \hat{l}_\theta .$$

Moreover, by Assumption 1, the marginal rate of tax $\bar{T}$ is strictly below 1 everywhere and, hence, its rate of progressivity is below 1 as well. Then, equation (40) implies that $\epsilon_\theta^w > 0$ and hence $\hat{l}_\theta =$

⁴⁸Note that under the initial tax $\bar{T}$ the exogenous and the endogenous planner agree about the equilibrium and in particular about the equilibrium technology. Hence, there is no need to distinguish the equilibrium values of the endogenous variables under tax $\bar{T}$ as perceived by the exogenous technology planner and their true equilibrium values.

⁴⁹The limit computations are analogous to those in the proof of Corollary 3. I therefore omit the details here.

$\epsilon_\theta^w \widehat{w}_\theta(\bar{l}, \phi^*(\bar{l})) > 0$ for all θ . Combining this with the expressions for $\widehat{\beta}_\theta$ and $\widehat{\bar{\beta}}_\theta$, we find that $\widehat{\beta}_\theta < \widehat{\bar{\beta}}_\theta$. Now use a change-of-variable to obtain

$$b_{\beta_\theta} \beta_\theta = \frac{f_\theta}{\widehat{\beta}_\theta} > \frac{f_\theta}{\widehat{\bar{\beta}}_\theta} = \bar{b}_{\bar{\beta}_\theta} \bar{\beta}_\theta,$$

and hence

$$\frac{1 - B_{\beta_\theta}}{b_{\beta_\theta} \beta_\theta} < \frac{1 - \bar{B}_{\bar{\beta}_\theta}}{\bar{b}_{\bar{\beta}_\theta} \bar{\beta}_\theta},$$

which proves inequality (38) in the main text.

Intuitively, since labor supply increases in the skill level θ under the tax $\bar{T}$, the equilibrium technology under this tax must be skill biased. The exogenous technology planner falsely believes that this skill bias is exogenous and persists irrespective of changes in labor supply. Thus, the exogenous technology planner overestimates the degree of exogenous inequality in the economy.

B Calibration Details

This Appendix provides more detailed information on the calibration of the directed technical change elasticity ρ^{CES} , the labor supply elasticity e , and the exogenous technology parameter $\kappa : \theta \mapsto \kappa_\theta$.

B.1 Calibration of Directed Technical Change Effects

The parameter ρ^{CES} , which controls the strength of directed technical change effects, is calibrated on the basis of the empirical estimates summarized in Table 1. Here, I give a brief overview over each of the studies listed in Table 1 and explain how I obtain the estimates in the last column of the table.

Carneiro et al. (2019) Carneiro et al. (2019) estimate the responses of relative supply and relative wages of college versus non-college workers to plausibly exogenous college openings in Norwegian municipalities in the 1970s, using synthetic control methods. They find that relative supply in a municipality starts rising shortly after the college opening and follows an upwards trend throughout the observation period of up to 17 years, compared to the synthetic control municipality. The relative wage first declines and then reverses its trend, surpassing the relative wage in the control municipality slightly more than 10 years after the college opening (see Figures 4 and 5 in Carneiro et al., 2019).

The numbers in Table 1 are derived from the estimates presented in Carneiro et al. (2019) as follows. First, measuring relative supply and relative wage changes two years after a college opening, Carneiro et al. (2019) estimate an elasticity of the relative wage with respect to relative supply of -0.549 , reported in column 1 of their Table 2. This produces the first row of Table 1 in the present paper.

Second, Figures 5 and 6 imply that after 10 years, relative wages in the treated municipalities and their synthetic controls were equal. Hence, when measured after 10 years, there is a zero effect of the

exogenous relative supply increase in the relative wage, leading to the second row of Table 1 in the present paper.

Finally, the third row of Table 1 is obtained from the plots presented in Figure 5 in [Carneiro et al. \(2019\)](#) as follows. The plots show that after 17 years, the log change in the relative wage, compared to the synthetic control municipality, is

$$\log \left(\frac{w_c^{17}}{w_{nc}^{17}} \right) - \log \left(\frac{w_c^0}{w_{nc}^0} \right) \approx 0.02 .$$

At the same time, the log change in the share of college workers in the total workforce was

$$\log \left(\frac{l_c^{17}}{l_c^{17} + l_{nc}^{17}} \right) - \log \left(\frac{l_c^0}{l_c^0 + l_{nc}^0} \right) \approx 0.04 .$$

To map this change into the change in the ratio of college over non-college workers, I rewrite the log change as follows:

$$\begin{aligned} \log \left(\frac{l_c^{17}}{l_c^{17} + l_{nc}^{17}} \right) - \log \left(\frac{l_c^0}{l_c^0 + l_{nc}^0} \right) &= \log(l_c^{17}) - \log(l_{nc}^{17}) - \log(l_c^0) + \log(l_{nc}^0) \\ &\quad - \log(l_c^{17} + l_{nc}^{17}) + \log(l_{nc}^{17}) + \log(l_c^0 + l_{nc}^0) - \log(l_{nc}^0) . \end{aligned}$$

[Carneiro et al. \(2019\)](#) report that the share of college workers was close to zero in most of the treated municipalities at the beginning of the observation period and still small at the end of the period. Hence, I apply the approximations

$$\log(l_c^{17} + l_{nc}^{17}) - \log(l_{nc}^{17}) \approx \log(l_c^0 + l_{nc}^0) - \log(l_{nc}^0) \approx 0$$

to obtain

$$\log \left(\frac{l_c^{17}}{l_c^{17} + l_{nc}^{17}} \right) - \log \left(\frac{l_c^0}{l_c^0 + l_{nc}^0} \right) \approx \log(l_c^{17}) - \log(l_{nc}^{17}) - \log(l_c^0) + \log(l_{nc}^0) \approx 0.04 .$$

Finally, relating the change in relative supply to the change in relative wages, I obtain

$$\frac{\log \left(\frac{w_c^{17}}{w_{nc}^{17}} \right) - \log \left(\frac{w_c^0}{w_{nc}^0} \right)}{\log \left(\frac{l_c^{17}}{l_{nc}^{17}} \right) - \log \left(\frac{l_c^0}{l_{nc}^0} \right)} \approx 0.5 ,$$

which is the estimate used in Table 1.

Note that relative supply did not change instantaneously at the beginning of the observation period but steadily increased throughout (Figure 4 in [Carneiro et al., 2019](#)). Hence, part of the relative supply increase occurred only shortly before the relative wage increase is measured at year 17 of the observation period. To the extent that technology adjustments to the more recent part of the rise in relative supply are not yet reflected in the measured increase in relative wages, the above procedure

underestimates the actual long-run effect of an exogenous increase in relative supply on relative wages.

Lewis (2011) Lewis (2011) uses plausibly exogenous variation in immigrant inflows across US metropolitan areas in the 1980s and 1990s to estimate the relationship between the relative supply of high-school graduates versus high-school dropouts on their relative wages. He studies changes over 10 year intervals, thereby capturing a rather long-run elasticity. He also provides evidence showing that firms' decisions to adopt a range of automation technologies in the manufacturing sector respond to the (exogenous component) of changes in relative supply in the way predicted by theory. This supports the view that the estimated long-run wage elasticity captures directed technical change effects.

In column 1 of Table VIII, Lewis (2011) reports a wage elasticity estimate of -0.136 . This is the estimate I use in Table 1.

Dustmann and Glitz (2015) Dustmann and Glitz (2015) exploit the arguably exogenous component of immigration inflows to German regions between 1985 and 1995 to analyze how regions absorb changes in relative skill supply. They decompose the change in relative employment levels between skill groups into a component due to between-firm scale adjustments and within-firm factor intensity adjustments. The latter turns out vastly more important, suggesting that Rybcinsky type output mix adjustments are small. Moreover, they find that relative wages hardly respond to relative supply changes. This leaves technology adjustments biased towards the skill group that becomes more abundant as the main margin of adjustment.

The authors distinguish between workers without postsecondary education (low-skilled), with postsecondary vocational or apprenticeship degrees (medium-skilled), and with college education (high-skilled). Due to extensive right-censoring of wages in the data, they consider their results for college workers less reliable and focus mainly on medium- and low-skilled workers.

For the relative wage of medium- versus low-skilled workers, Dustmann and Glitz (2015) estimate an elasticity with respect to relative supply of -0.091 (row 2, column 4, Table 2) over a period of ten years. This estimate uses data for the tradable goods sector (which includes, but is not limited to, the manufacturing sector). For the non-tradable sector, the authors find a much smaller wage elasticity. Yet, when they pool all industries, results are close to those for the tradable goods sector again (see description on page 727, Dustmann and Glitz, 2015). Hence, I use the estimate for the tradable sector in Table 1.

Morrow and Trefler (2017) Morrow and Trefler (2017) start from a detailed neoclassical model of international trade building on Eaton and Kortum (2002). They estimate their model on sectoral factor input and price data for a cross-section of 38 countries in 2006. Country selection is driven by data availability in the World Input Output Database. Labor is partitioned into skilled and unskilled labor. Skilled workers are those with at least some tertiary education, unskilled workers are those without.

In the model, the relative wage between skilled and unskilled workers in each country is determined by the relative labor input and exogenous factor-augmenting productivity. To separately identify factor-augmenting productivity and the elasticity of substitution between labor types, Morrow and Trefler

(2017) augment their model's equilibrium conditions by a directed technical change equation similar to equation (45). Unfortunately, their approach requires to fix the technology substitution parameter δ exogenously. In their directed technical change equation, they (implicitly) assume $\delta = 1$. Given a value for δ , the directed technical change equation and the equation for relative wages at given technology identify the elasticity of substitution σ (without observing technology).

In their most elaborate estimation (the estimation of their equation 32), [Morrow and Trefler \(2017\)](#) find a value for σ of 1.89, which translates into a wage elasticity at exogenous technology (or, short-run wage elasticity) of $-1/1.89 = -0.53$. This is the first value from [Morrow and Trefler \(2017\)](#) I use in Table 1. Combining relative wage and directed technical change equations, the total wage elasticity, including directed technical change effects, is then obtained as $\sigma - 2 = -0.11$, the second estimate from [Morrow and Trefler \(2017\)](#) reported in Table 1.

Relative to the other studies listed in Table 1, a major shortcoming of [Morrow and Trefler \(2017\)](#) is that they do not have a strategy to isolate exogenous variation in factor inputs when estimating their directed technical change equation. Hence, part of the estimated relationship between technology and factor inputs may be driven by reverse causality, which leads to an overestimate of directed technical change effects. On the other hand, the fact that they estimate their model on cross-sectional data may imply underestimation of directed technical change effects, because the observed technology levels may not yet have fully adjusted to the most recent factor input changes.

B.2 Calibration of the Labor Supply Elasticity

In Section 7.1 of the main text, I claim that the model generates an elasticity of taxable income of 0.25 (as reported by [Saez et al., 2012](#), as their preferred estimate) if the labor supply elasticity e is set to 0.67. Here, I explain this claim.

The elasticity e measures the response of labor supply to an increase in the marginal retention rate if the tax function is locally linear. Starting from e , three adjustments are necessary to arrive at the elasticity of taxable income. First, the elasticity of taxable income must account for the potential non-linearity of the tax function. This adjustment is already incorporated in the labor supply elasticity ϵ_θ^R from Section 5.

Next, if labor supply of a given worker type θ changes, this will affect the wage of type θ in equilibrium, which in turn feeds back to labor supply. Hence, denoting the equilibrium elasticity of labor supply by $\bar{\epsilon}_\theta^R$, it must solve the equation

$$\bar{\epsilon}_\theta^R = \epsilon_\theta^R + \epsilon_\theta^w (\gamma_\theta^{own} + \rho_\theta^{own}) \bar{\epsilon}_\theta^R .$$

Finally, taxable income is the product of labor supply and wage. Thus, the elasticity of taxable income, denoted $\epsilon_\theta^{TI,R}$, is obtained from the equilibrium elasticity of labor supply $\bar{\epsilon}_\theta^R$ according to $\epsilon_\theta^{TI,R} = (1 + \gamma_\theta^{own} + \rho_\theta^{own}) \bar{\epsilon}_\theta^R$. Combining the previous equations, we thus obtain

$$\epsilon_\theta^{TI,R} = \frac{(1 + \gamma_\theta^{own} + \rho_\theta^{own}) \epsilon_\theta^R}{1 - \epsilon_\theta^w (\gamma_\theta^{own} + \rho_\theta^{own})} .$$

To compute an elasticity of taxable income close to what is typically estimated in the empirical literature, I make the following assumptions. First, I assume a locally linear tax function, because real-world tax schedules are piece-wise linear. This implies that the elasticities ϵ_θ^R and ϵ_θ^w are both equal to e (see equations (39) and (40) in Appendix A.2). Second, I only use the short-run response of wages to labor supply (represented by γ_θ^{own}) in the computation of the equilibrium elasticity of labor supply, because empirical estimates of the elasticity of taxable income typically measure the response of income to tax changes over relatively short horizons of about one or two years (Saez et al., 2012). Together, these assumptions lead to the following expression for the model-implied counterpart of empirical estimates of the elasticity of taxable income:

$$\tilde{\epsilon}_\theta^{TI,R} = \frac{(1 + \gamma_\theta^{own})e}{1 - e\gamma_\theta^{own}}.$$

Setting $\gamma_\theta^{own} = \gamma^{CES} = -0.5$ (as in the baseline calibration) and inverting the above equation shows that e takes a value of 0.67 if $\tilde{\epsilon}_\theta^{TI,R} = 0.25$.

B.3 Calibration of the Exogenous Technology Parameter

To calibrate the exogenous technology parameter κ , an estimate of the earnings distribution under the initial tax system $\bar{T}$ is needed. As explained in the main text, the initial tax system is set to approximate the US income tax in 2005. Hence, the income distribution under $\bar{T}$ should approximate the empirical earnings distribution of the US in 2005.

As is standard in the literature (e.g. Mankiw et al., 2009; Diamond and Saez, 2011), I approximate the empirical earnings distribution by merging a lognormal distribution (for the bulk of incomes) and a Pareto distribution (for the upper tail). I also assume that there is a mass point of workers with zero income (as, e.g., in Mankiw et al., 2009; Brewer et al., 2010), which I set to 2%.⁵⁰

Since the earnings distribution enters most of the formulae used in the simulations via its hazard ratio $ym_y/(1 - M_y)$, I directly target the empirical hazard ratio in 2005. In particular, I construct the hazard ratio as

$$\frac{ym_y}{1 - M_y} = \frac{ym_y^{\lognormal}}{1 - M_y^{\lognormal}} \left(1 - \Phi \left(\frac{y - 200000}{\sigma^{normal}} \right) \right) + \frac{ym_y^{Pareto}}{1 - M_y^{Pareto}} \Phi \left(\frac{y - 200000}{\sigma^{normal}} \right),$$

where the normal distribution used for smoothing has a mean \$200k, reflecting the region in the earnings distribution where the transition from lognormal to Pareto occurs. I then choose the parameters of the lognormal and the Pareto distribution to match key properties of the empirical hazard ratio in 2005. The Pareto shape parameter is set to 1.5, which is the hazard ratio of the empirical earnings distribution for high incomes (see, e.g., Figure 2 in Diamond and Saez, 2011). The lognormal mean and variance and the variance σ^{normal} of the smoothing function are set to 10.6, 0.85, and 75000, respectively. These values ensure that the average income matches its empirical counterpart of about \$63k and that the resulting hazard ratio peaks at about \$150k, decreases until about \$350k, and flattens out afterwards, as

⁵⁰The main results are robust to other values of the mass point between 0% and 10%.




Figure 3. The figure shows the hazard ratio of the income distribution under the initial tax $\bar{T}$ used to calibrate the exogenous technology parameter κ . The construction of the hazard ratio follows the description in the text. The hazard ratio approximates the empirical hazard ratio of the US earnings distribution in 2005, as depicted, for example, in Figure 2 in [Diamond and Saez \(2011\)](#).

depicted in Figure 3 (see again Figure 2 in [Diamond and Saez, 2011](#), for comparison with the empirical US hazard ratio in 2005).

Given the hazard ratio of earnings, I obtain the cumulative distribution function by solving the corresponding differential equation. Specifically, when k_y denotes the hazard ratio of the earnings distribution, the cumulative distribution function solves

$$\frac{dM_y}{dy} = \frac{k_y}{y} - \frac{k_y}{y} M_y .$$

Finally, the density function of incomes is obtained as the numerical derivative of M .

Since the distribution of types on the type space is uniform, the cumulative distribution function of incomes M returns for each income the type who earns this income under the initial tax $\bar{T}$. Hence, the income function $y : \theta \mapsto y_\theta$ is given by the inverse of M .

Given y_θ , it is straightforward to compute κ_θ from workers' first-order condition and the condition that wages equal marginal products of labor in aggregate production. First, multiplying the first-order condition (2) by l_θ and solving for it yields

$$l_\theta = (R'_T(y_\theta)y_\theta)^{\frac{e}{1+e}} ,$$

where I used that the disutility of labor is isoelastic in the quantitative analysis. With the estimate of $\bar{T}$ described in the main text, the previous equation allows to compute labor inputs under $\bar{T}$.

For the second step, start from equations (15) and (46), copied here for convenience:

$$w_\theta(l, \phi) = (\kappa_\theta \phi_\theta)^{\frac{\sigma-1}{\sigma}} (l_\theta h_\theta)^{-\frac{1}{\sigma}} F(l, \phi)^{\frac{1}{\sigma}}$$

$$\phi_\theta^* = \bar{C}^{\frac{1}{\delta}} (\kappa_\theta h_\theta l_\theta)^{\frac{\sigma-1}{(\delta-1)\sigma+1}} \left[\int_{\underline{\theta}}^{\bar{\theta}} (\kappa_\theta h_\theta l_\theta)^{\frac{(\sigma-1)\delta}{(\delta-1)\sigma+1}} d\theta \right]^{-\frac{1}{\delta}}.$$

The total amount of R&D resources $\bar{C}$ is not identified separately from κ , so I normalize it to satisfy

$$\bar{C} = \int_{\underline{\theta}}^{\bar{\theta}} (\kappa_\theta h_\theta l_\theta)^{\frac{(\sigma-1)\delta}{(\delta-1)\sigma+1}} d\theta$$

under the initial tax. Using this normalization in the above equation for ϕ_θ^* , plugging the equation into the expression for the wage w_θ , multiplying by l_θ , and solving for κ_θ yields:

$$\kappa_\theta = y_\theta^{\frac{1}{1+\gamma^{CES}+\rho^{CES}}} l_\theta^{-1} F^{\frac{\gamma^{CES}}{1+\gamma^{CES}+\rho^{CES}}}.$$

With $F = \int_{\underline{\theta}}^{\bar{\theta}} y_\theta m_y dy$ by Euler's theorem, this allows to compute κ .

Finally, the optimal tax formulae in Proposition 3 and Proposition 4 require the inverse hazard ratio of the exogenous parameters β and $\bar{\beta}$, respectively. In principle, β and $\bar{\beta}$ can be computed from κ and from the equilibrium technology under initial taxes via their definitions. Then, their pdf and cdf, and finally their hazard ratios can be computed. Here, to avoid unnecessary approximations, I choose a more direct way and compute the inverse hazard ratios of β and $\bar{\beta}$ directly from the hazard ratio of the income distribution, using equations (72) and (74). This ensures that the two hazard ratios inherit their shape directly from the shape of the initial hazard ratio of incomes (which is calibrated to match its empirical counterpart), without numerical differentiation or integration steps and the associated approximation errors in between.

C Supplementary Material

This appendix contains several results complementary to those presented in the main text. Section C.1 provides results for the welfare effects of tax reforms, extending the analysis of tax reforms in the main text. Sections C.2 and C.3 extend the quantitative analysis from the main text, computing welfare effects of the progressive tax reform analyzed in the main text and optimal marginal tax rates for a Rawlsian welfare function, respectively. Finally, Section C.4 contains all proofs of the results presented in this appendix.

C.1 Welfare Effects of Tax Reforms

Given that under certain conditions progressive tax reforms induce equalizing technical change, it is natural to suspect that taking into account directed technical change effects of tax reforms should raise the predicted welfare gains from progressive reforms. This conjecture is examined in the following.

To analyze the welfare effects of a tax reform τ , write welfare as a function of the tax system:

$$\tilde{W}(T) := V(\{u_\theta(c_\theta(T), l_\theta(T))\}_{\theta \in \Theta}) .$$

Given the welfare function $\tilde{W}(T)$, the welfare effect of a tax reform can now be decomposed as follows.

Proposition 5. *Fix an initial tax T . The welfare effect of a tax reform τ can be written as*

$$\begin{aligned} \left. \frac{d\tilde{W}(T + \mu\tau)}{d\mu} \right|_{\mu=0} &= \underbrace{\int_{\underline{\theta}}^{\bar{\theta}} (1 - g_\theta) \tau(y_\theta(T)) h_\theta d\theta}_{:=ME_\tau(T)} + \underbrace{\int_{\underline{\theta}}^{\bar{\theta}} T'(y_\theta(T)) y_\theta(T) (-\epsilon_\theta^R) \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} h_\theta d\theta}_{:=BE_\tau(T)} \\ &+ \underbrace{\int_{\underline{\theta}}^{\bar{\theta}} [g_\theta(1 - T'(y_\theta(T))) + T'(y_\theta(T))(1 + \epsilon_\theta^w)] y_\theta(T) \left. \frac{d \log(w_\theta(T, \phi^*(T + \mu\tau)))}{d\mu} \right|_{\mu=0} h_\theta d\theta}_{:=TE_\tau^W(T)} \\ &+ \underbrace{\int_{\underline{\theta}}^{\bar{\theta}} [g_\theta(1 - T'(y_\theta(T))) + T'(y_\theta(T))(1 + \epsilon_\theta^w)] y_\theta(T) \left. \frac{d \log(w_\theta(T + \mu\tau, \phi^*(T)))}{d\mu} \right|_{\mu=0} h_\theta d\theta}_{:=SE_\tau^W(T)} . \end{aligned}$$

Proof. See Appendix C.4. □

Proposition 5 shows that a tax reform has four distinct effects on welfare. The mechanical effect $ME_\tau(T)$ captures the effect from changing taxes and redistributing revenue in the absence of any behavioral responses. The behavioral effect $BE_\tau(T)$ captures the effect of the direct response of labor supply, holding wages constant. Both effects are well known in the literature.

The third term $TE_\tau^W(T)$ represents the welfare implications of the technical change induced by the

tax reform. The first part,

$$\int_{\underline{\theta}}^{\bar{\theta}} g_{\theta}(1 - T'(y_{\theta}(T)))y_{\theta}(T) \frac{d \log(w_{\theta}(T, \phi^*(T + \mu\tau)))}{d\mu} \Big|_{\mu=0} h_{\theta} d\theta ,$$

captures the direct effect of the technology-induced wage changes on workers' utility: from the change in pre-tax income, only the share $1 - T'(y_{\theta}(T))$ translates directly into a change of utility as the remaining share is taxed away. The second part,

$$\int_{\underline{\theta}}^{\bar{\theta}} T'(y_{\theta}(T))(1 + \epsilon_{\theta}^w)y_{\theta}(T) \frac{d \log(w_{\theta}(T, \phi^*(T + \mu\tau)))}{d\mu} \Big|_{\mu=0} h_{\theta} d\theta ,$$

is the welfare effect of the lump-sum redistribution of the revenue gain or loss induced by the wage adjustments to technical change. Here, the pre-tax income change is scaled by $1 + \epsilon_{\theta}^w$, as the wage change induces a labor supply adjustment of ϵ_{θ}^w .⁵¹

Importantly, even if the induced technical change reduces the skill premium (e.g., because τ is progressive and the conditions of Corollary 1 are satisfied), we cannot sign the directed technical change effects on welfare unambiguously. This is because, when starting, for example, from a progressive tax T , the reduction in high-skilled workers' wages passes through to the government budget to a larger extent than the simultaneous rise in low-skilled workers' wages, as marginal tax rates are higher for the high-skilled. Hence, directed technical change may reduce tax revenue following a progressive reform, which affects welfare negatively via reduced lump-sum transfers. This negative welfare effect potentially outweighs the positive effect coming from the reduction in pre-tax wage inequality through the induced technical change.⁵²

The final term in Proposition 5, $SE_{\tau}^W(T)$ captures the welfare effect of the within-technology substitution effects on wages caused by the tax reform. Its structure is analogous to that of $TE_{\tau}^W(T)$. Given that even the directed technical change component $TE_{\tau}^W(T)$ has an ambiguous effect on welfare, it is not surprising that also the substitution component $SE_{\tau}^W(T)$ cannot be signed in general.

Importantly, however, Proposition 5 can be combined with equation (26) from Proposition 1 for the relative wage effects of tax reforms. This yields a formula for the welfare effects of tax reforms in terms of empirically observable quantities and welfare weights. I use this formula to quantify the welfare effects of tax reforms and the contribution of directed technical change in Appendix C.2.

The implications of Proposition 5 may be somewhat unexpected in light of the previous results. After all, if a progressive reform induces equalizing technical change and the welfare function values equity, directed technical change effects should make progressive reforms in some way more attractive. To see precisely in which way this is indeed true, we must slightly adjust the question posed by Proposition 5.

Instead of asking how directed technical change alters the welfare effects of a given progressive tax reform, we now study how accounting for directed technical change affects the set of initial taxes under

⁵¹The labor supply adjustment does not enter the first part of $TE_{\tau}^W(T)$ because it does not affect workers' utility by the envelope theorem.

⁵²This is the inverse of the observation by Sachs et al. (2020) that within-technology substitution effects may increase the revenue gains from progressive tax reforms if the initial tax schedule is already progressive.

which welfare can be improved by some progressive reform. In particular, let

$$\mathcal{T} := \left\{ T \mid T \text{ has CRP, } \exists \tau \text{ progressive s.t. } D_\tau \tilde{W}(T) > 0 \right\}$$

denote the set of tax schedules with a constant rate of progressivity (CRP) that can be improved in a welfare sense by a progressive tax reform. The restriction to CRP taxes is imposed to invoke Corollary 1. Specifically, combining the CRP restriction with isoelastic disutility of labor and the CES production structure from Section 3.6 ensures, according to Corollary 1, that any progressive tax reform induces equalizing technical change.

As a benchmark for comparison that does not include directed technical change effects, let

$$D_\tau^{ex} \tilde{W}(T) := \left. \frac{d\tilde{W}(T + \mu\tau)}{d\mu} \right|_{\mu=0, \rho_\theta^{own} = \rho_{\theta, \tilde{\theta}} = 0 \forall \theta, \tilde{\theta}} \quad (75)$$

denote the welfare effect of a reform τ when counterfactually setting all directed technical change elasticities to zero (i.e., when holding technology fixed). Then, we can define

$$\mathcal{T}^{ex} := \left\{ T \mid T \text{ has CRP, } \exists \tau \text{ progressive s.t. } D_\tau^{ex} \tilde{W}(T) > 0 \right\}$$

as the set of CRP schedules that one would perceive to be improvable by progressive reforms if one were to ignore directed technical change.

Comparing the two thus defined sets I find that accounting for directed technical change expands the set of tax schedules under which welfare can be improved by a progressive reform.

Proposition 6. *Suppose F and Φ are CES as introduced in Section 3.6 and the disutility of labor is isoelastic. Then,*

$$\mathcal{T}^{ex} \subseteq \mathcal{T},$$

that is, the set of initial tax schedules that can be improved by a progressive reform becomes larger when accounting for directed technical change effects.

Proof. See Appendix C.4. □

Proposition 6 proposes a way in which directed technical change effects make progressive reforms more attractive. Specifically, accounting for directed technical change increases the scope for welfare improvements by progressive tax reforms.

The idea behind Proposition 6 relies on the mechanism design approach to income taxation. Consider a progressive tax reform that a tax planner who neglects directed technical change effects (the exogenous technology planner, he, henceforth) expects to raise welfare. For any such reform, a planner who correctly anticipates directed technical change effects (the endogenous technology planner, she, henceforth) can find another progressive reform that exactly replicates the labor allocation expected by the exogenous technology planner following his reform. But since progressive tax reforms induce




Figure 4. The figure displays changes in the lump-sum payment that are equivalent to the welfare effect of the progressive tax reform described in the text. These changes are measured in % of the pre-reform average income and displayed for different values of the relative inequality aversion parameter r . The baseline calibration described in the main text applies.

equalizing technical change, the endogenous technology planner anticipates a more equal wage distribution after her reform than the exogenous technology planner expects to find after his reform. Via incentive compatibility constraints, a more equal wage distribution allows to distribute consumption more equally as well. Hence, while the two planners expect the same labor allocation to materialize, the endogenous technology planner anticipates a more equal consumption distribution than the exogenous technology planner. Since this reasoning holds for any progressive reform of the exogenous technology planner, the endogenous technology planner can find a welfare-improving progressive reform whenever the exogenous technology planner can find one. Hence, the endogenous technology planner perceives the scope for welfare improvements through progressive tax reforms to be greater.

C.2 Quantitative Analysis: Welfare Effects of Tax Reforms

In this section, I complement the quantitative analysis of the wage effects of progressive tax reforms in Section 7 by assessing the welfare effects of these reforms. In particular, I consider the same progressive tax reform as described in Section 7 and compute its welfare effects based on Proposition 5 from Appendix C.1. The calibration is the same as described in the main text.

Figure 4 displays (lump-sum) consumption changes equivalent to the welfare change induced by the tax reform for different values of the relative inequality aversion parameter r . The lump-sum consumption changes are expressed in percent of initial average income.

As observed analytically (see Proposition 5 and the subsequent discussion), the influence of directed




Figure 5. The figure displays optimal marginal tax rates by income level for a relative inequality aversion parameter of $r = 50$. Otherwise, the baseline calibration described in the main text applies.

technical change effects on the welfare assessment of a given progressive reform is ambiguous. The figure shows that for lower degrees of inequality aversion, directed technical change raises the welfare gains from the progressive reform. Here, the reduction in pre-tax wage inequality induced by directed technical change outweighs the loss in tax revenue, which translates into a reduction in the lump-sum payment to all workers. For high values of inequality aversion, for which the welfare function approaches a Rawlsian objective, the negative revenue effect from directed technical change becomes dominant, reducing the welfare gains from the reform. Interestingly, for very high levels of inequality aversion, accounting for directed technical change even switches the sign of the welfare effect: when ignoring directed technical change, the reform appears to raise welfare, but when accounting for directed technical change, the reform becomes undesirable. Notwithstanding these results, Proposition 6 implies that even for very high values of inequality aversion, there must be a different progressive reform that raises welfare when accounting for directed technical change.

C.3 Quantitative Analysis: Rawlsian Optimal Taxes

To complement the quantitative analysis of optimal taxes in the main text, I compute optimal marginal tax rates for an inequality aversion parameter of $r = 50$, approximating a Rawlsian welfare function. For comparison, Figure 5 also displays the corresponding marginal tax rates preferred by the exogenous technology planner.

The qualitative insights from Section 7 remain valid with Rawlsian welfare. Directed technical change still reduces optimal marginal tax rates below the median income and increase them above. There are

differences in the magnitudes of these adjustments, however. The reduction in marginal tax rates below the median is somewhat smaller than in the baseline scenario (at the 10th percentile: 4 vs. 5 pp in the conservative case and 12 vs. 17 pp in the strong bias case), while the increase in marginal tax rates above the median becomes more pronounced (at the 90th percentile: 6 vs. 3 pp in the conservative case and 14 vs. 8 pp in the strong bias case). Optimal marginal tax rates are U-shaped now even in the strong bias case.

The reason for these differences is that, with a Rawlsian welfare function, incomes slightly below the median are not relevant from a social perspective except for their contribution to tax revenue. Hence, the incentive to redistribute pre-tax income from high earners to earners below the median via directed technical change effects is mitigated. This, however, was the driving force behind the low optimal marginal tax rates on below-median incomes in the baseline scenario.

The welfare gains from optimal taxes relative to those of the exogenous technology planner generally become larger with Rawlsian welfare. In the strong bias case, they are equivalent to an increase in the lump-sum payment of \$850 annually, which corresponds to 1.4% of average income under the exogenous technology planner's taxes.

C.4 Proofs for the Supplementary Material

This section contains all proofs for the results presented in Appendix C.

Proof of Proposition 5. Welfare can be written as

$$\tilde{W}(T) = V \left(\{ R_T(w_\theta(T, \phi^*(T))l_\theta(T, w_\theta)) + S(l(T, w), w(T, \phi^*(T)), T) - v(l_\theta(T, w_\theta)) \}_{\theta \in \Theta} \right) ,$$

where

$$S(l(T, w), w(T, \phi^*(T)), T) = \int_{\underline{\theta}}^{\bar{\theta}} T(w_\theta(T, \phi^*(T))l_\theta(T, w_\theta)) h_\theta d\theta .$$

For notational convenience, let $D_\tau w(T, \phi^*(T))$ denote the within-technology substitution derivative $dw(T + \mu\tau, \phi^*(T))/d\mu|_{\mu=0}$ and $D_{\phi, \tau} w(T, \phi^*(T))$ the directed technical change derivative $dw(T, \phi^*(T + \mu\tau))/d\mu|_{\mu=0}$ of the wage in direction of the reform τ . With this notation, perturbing the tax function in

direction of τ yields:

$$\begin{aligned}
D_\tau \tilde{W}(T) = & \int_{\underline{\theta}}^{\bar{\theta}} \left(-g_\theta h_\theta \tau(w_\theta l_\theta) + g_\theta h_\theta \int_{\underline{\theta}}^{\bar{\theta}} h_{\tilde{\theta}} \tau(w_{\tilde{\theta}} l_{\tilde{\theta}}) d\tilde{\theta} \right) d\theta \\
& + \int_{\underline{\theta}}^{\bar{\theta}} g_\theta h_\theta \int_{\underline{\theta}}^{\bar{\theta}} h_{\tilde{\theta}} T'(w_{\tilde{\theta}} l_{\tilde{\theta}}) w_{\tilde{\theta}} (-\epsilon_\theta^R) \frac{l_{\tilde{\theta}}}{1 - T'(w_{\tilde{\theta}} l_{\tilde{\theta}})} \tau'(w_{\tilde{\theta}} l_{\tilde{\theta}}) d\tilde{\theta} d\theta \\
& + \int_{\underline{\theta}}^{\bar{\theta}} g_\theta h_\theta (1 - T'(w_\theta l_\theta)) l_\theta D_{\phi, \tau} w_\theta(T, \phi^*(T)) d\theta \\
& + \int_{\underline{\theta}}^{\bar{\theta}} g_\theta h_\theta \int_{\underline{\theta}}^{\bar{\theta}} h_{\tilde{\theta}} T'(w_{\tilde{\theta}} l_{\tilde{\theta}}) \left(l_{\tilde{\theta}} D_{\phi, \tau} w_{\tilde{\theta}}(T, \phi^*(T)) + w_{\tilde{\theta}} \epsilon_{\tilde{\theta}}^w \frac{l_{\tilde{\theta}}}{w_{\tilde{\theta}}} D_{\phi, \tau} w_{\tilde{\theta}}(T, \phi^*(T)) \right) d\tilde{\theta} d\theta \\
& + \int_{\underline{\theta}}^{\bar{\theta}} g_\theta h_\theta (1 - T'(w_\theta l_\theta)) l_\theta D_\tau w_\theta(T, \phi^*(T)) d\theta \\
& + \int_{\underline{\theta}}^{\bar{\theta}} g_\theta h_\theta \int_{\underline{\theta}}^{\bar{\theta}} h_{\tilde{\theta}} T'(w_{\tilde{\theta}} l_{\tilde{\theta}}) \left(l_{\tilde{\theta}} D_\tau w_{\tilde{\theta}}(T, \phi^*(T)) + w_{\tilde{\theta}} \epsilon_{\tilde{\theta}}^w \frac{l_{\tilde{\theta}}}{w_{\tilde{\theta}}} D_\tau w_{\tilde{\theta}}(T, \phi^*(T)) \right) d\tilde{\theta} d\theta .
\end{aligned}$$

Using $\int_{\underline{\theta}}^{\bar{\theta}} h_\theta g_\theta d\theta = 1$, we can rearrange this expression to obtain

$$\begin{aligned}
D_\tau \tilde{W}(T) = & \int_{\underline{\theta}}^{\bar{\theta}} (1 - g_\theta) \tau(w_\theta l_\theta) h_\theta d\theta + \int_{\underline{\theta}}^{\bar{\theta}} T'(w_\theta l_\theta) w_\theta l_\theta (-\epsilon_\theta^R) \frac{\tau'(w_\theta l_\theta)}{1 - T'(w_\theta l_\theta)} h_\theta d\theta \\
& + \int_{\underline{\theta}}^{\bar{\theta}} [g_\theta (1 - T'(w_\theta l_\theta)) + T'(w_\theta l_\theta) (1 + \epsilon_\theta^w)] l_\theta D_{\phi, \tau} w_\theta(T, \phi^*(T)) h_\theta d\theta \\
& + \int_{\underline{\theta}}^{\bar{\theta}} [g_\theta (1 - T'(w_\theta l_\theta)) + T'(w_\theta l_\theta) (1 + \epsilon_\theta^w)] l_\theta D_\tau w_\theta(T, \phi^*(T)) h_\theta d\theta .
\end{aligned}$$

□

Proof of Proposition 6. Take any initial tax $T \in \mathcal{T}^{ex}$ and any reform τ^{ex} that is progressive and raises welfare when neglecting directed technical change effects, that is, $D_{\tau^{ex}} \tilde{W}(T) > 0$ (see equation (75)). The strategy of the proof is to construct another progressive reform τ^{en} that raises welfare when accounting for directed technical change effects, that is, $D_{\tau^{en}} \tilde{W}(T) > 0$. Constructing such a reform will prove that $T \in \mathcal{T}$ and hence $\mathcal{T}^{ex} \subseteq \mathcal{T}$.

I construct the reform τ^{en} such that it exactly replicates the labor input changes that τ^{ex} would induce if there were no directed technical change effects. To move back and forth between induced labor input changes and progressive reforms, I use Lemma 6. In particular, note that Proposition 6 considers reforms of CRP tax schedules when the disutility of labor is isoelastic. Under these conditions the elasticities ϵ_θ^R and ϵ_θ^w are constant in θ . Lemma 6 then says that any progressive reform τ induces labor input responses $\hat{l}_{\theta, \tau}$ that decrease in θ and, conversely, any reform that induces labor input changes that decrease in θ is progressive.

After these preparations, write welfare as a function of consumption and labor inputs only, that is,

$$W(c, l) := V(\{u_\theta(c_\theta, l_\theta)\}_{\theta \in \Theta}) .$$

Then, the effect of reform τ^{ex} on welfare, ignoring directed technical change effects, is fully determined by the responses of consumption and labor supply to τ^{ex} that we would obtain if technology were fixed. I analyze these responses in the following.

Step 1. Denote the labor input response to τ^{ex} that ignores directed technical change effects by

$$D_{\tau^{ex}}^{ex} l_{\theta}(T) := \left. \frac{dl_{\theta}(T + \mu\tau)}{d\mu} \right|_{\mu=0, \rho_{\theta}^{own} = \rho_{\theta, \tilde{\theta}} = 0 \forall \theta, \tilde{\theta}}$$

and similarly the consumption response that ignores directed technical change effects by $D_{\tau^{ex}}^{ex} c_{\theta}(T)$.

I now characterize the consumption response contingent on the labor input response using incentive compatibility constraints. In particular, at any tax $\tilde{T}$, consumption and labor allocations must satisfy

$$c_{\tilde{\theta}}(\tilde{T}) - v \left(\frac{w_{\tilde{\theta}}(\tilde{T}, \phi^*(\tilde{T})) l_{\tilde{\theta}}(\tilde{T})}{w_{\theta}(\tilde{T}, \phi^*(\tilde{T}))} \right) \leq c_{\theta}(\tilde{T}) - v(l_{\theta}(\tilde{T})) \quad \text{for all } \theta, \tilde{\theta}.$$

Via an envelope argument this implies

$$c'_{\theta}(\tilde{T}) = v'(l_{\theta}(\tilde{T})) \left[l'_{\theta}(\tilde{T}) + \hat{w}_{\theta}(\tilde{T}, \phi^*(\tilde{T})) l_{\theta}(\tilde{T}) \right] \quad \text{for all } \theta.$$

Here, $\hat{w}_{\theta} = w'_{\theta}/w_{\theta}$ and the notation $x'_{\theta}(T)$ is exclusively used to denote differentiation with respect to the type index θ . So, $l'_{\theta}(T)$ is the derivative of $l_{\theta}(T)$ with respect to θ (and at θ), holding T constant. Integrating over θ , the envelope condition yields:

$$c_{\theta}(\tilde{T}) = c_{\underline{\theta}}(\tilde{T}) + \int_{\underline{\theta}}^{\theta} v'(l_{\tilde{\theta}}(\tilde{T})) \left[l'_{\tilde{\theta}}(\tilde{T}) + \hat{w}_{\tilde{\theta}}(\tilde{T}, \phi^*(\tilde{T})) l_{\tilde{\theta}}(\tilde{T}) \right] d\tilde{\theta} \quad \text{for all } \theta. \quad (76)$$

The level $c_{\underline{\theta}}$ is determined via the resource constraint:

$$\int_{\underline{\theta}}^{\bar{\theta}} c_{\theta}(\tilde{T}) h_{\theta} d\theta = F(l(\tilde{T}), \phi^*(l(\tilde{T}))). \quad (77)$$

Using equation (76), the response of consumption to tax reform τ^{ex} , ignoring directed technical change effects, can be expressed as

$$\begin{aligned} D_{\tau^{ex}}^{ex} c_{\theta}(T) &= D_{\tau^{ex}}^{ex} c_{\underline{\theta}}(T) + \int_{\underline{\theta}}^{\theta} v''(l_{\tilde{\theta}}(T)) (D_{\tau^{ex}}^{ex} l_{\tilde{\theta}}(T)) \left[l'_{\tilde{\theta}}(T) + \hat{w}_{\tilde{\theta}}(T, \phi^*(T)) l_{\tilde{\theta}}(T) \right] d\tilde{\theta} \\ &\quad + \int_{\underline{\theta}}^{\theta} v'(l_{\tilde{\theta}}(T)) \left[D_{\tau^{ex}}^{ex} l'_{\tilde{\theta}}(T) + \hat{w}_{\tilde{\theta}}(T, \phi^*(T)) D_{\tau^{ex}}^{ex} l_{\tilde{\theta}}(T) \right] d\tilde{\theta} \\ &\quad + \int_{\underline{\theta}}^{\theta} v'(l_{\tilde{\theta}}(T)) l_{\tilde{\theta}}(T) D_{\tau^{ex}}^{ex} \hat{w}_{\tilde{\theta}}(T, \phi^*(T)) d\tilde{\theta}. \quad (78) \end{aligned}$$

The resource constraint (77) implies

$$\int_{\underline{\theta}}^{\bar{\theta}} h_{\theta} D_{\tau^{ex}}^{\tau^{ex}} c_{\theta}(T) d\theta = \frac{d}{d\mu} F(l(T + \mu\tau^{ex}), \phi^*(l(T))) \Big|_{\mu=0}. \quad (79)$$

Step 2. Suppose now that we can find a reform τ^{en} that replicates the labor input change $D_{\tau^{ex}}^{\tau^{ex}} l_{\theta}(T)$ while accounting for directed technical change effects (I verify that such a reform exists below). That is, take τ^{en} such that

$$D_{\tau^{en}} l_{\theta}(T) := \frac{dl_{\theta}(T + \mu\tau^{en})}{d\mu} \Big|_{\mu=0} = D_{\tau^{ex}}^{\tau^{ex}} l_{\theta}(T) \quad \text{for all } \theta.$$

Here, I again use the notation $D_{\tau}x(T)$ to denote Gateaux differentiation of a function x in direction of τ . If the function x depends on T both directly and via the endogenous technology $\phi^*(T)$, I write $x(T, \phi^*(T))$ and use D_{τ} to denote the effect on x of perturbing T while holding technology constant at $\phi^*(T)$. The directed technical change effect of the perturbation, that is, $dx(T, \phi^*(T + \mu\tau))/d\mu|_{\mu=0}$, is then denoted by $D_{\phi, \tau}x(T, \phi^*(T))$, as in the proof of Proposition 5 above.

Using equation (76), the consumption response to the reform τ^{en} , accounting for directed technical change effects, can be expressed as

$$\begin{aligned} D_{\tau^{en}} c_{\theta}(T) &= D_{\tau^{en}} c_{\underline{\theta}}(T) + \int_{\underline{\theta}}^{\theta} v''(l_{\tilde{\theta}}(T)) (D_{\tau^{en}} l_{\tilde{\theta}}(T)) \left[l'_{\tilde{\theta}}(T) + \widehat{w}_{\tilde{\theta}}(T, \phi^*(T)) l_{\tilde{\theta}}(T) \right] d\tilde{\theta} \\ &\quad + \int_{\underline{\theta}}^{\theta} v'(l_{\tilde{\theta}}(T)) \left[D_{\tau^{en}} l'_{\tilde{\theta}}(T) + \widehat{w}_{\tilde{\theta}}(T, \phi^*(T)) D_{\tau^{en}} l_{\tilde{\theta}}(T) \right] d\tilde{\theta} \\ &\quad + \int_{\underline{\theta}}^{\theta} v'(l_{\tilde{\theta}}(T)) l_{\tilde{\theta}}(T) \left[D_{\tau^{en}} \widehat{w}_{\tilde{\theta}}(T, \phi^*(T)) + D_{\phi, \tau^{en}} \widehat{w}_{\tilde{\theta}}(T, \phi^*(T)) \right] d\tilde{\theta}. \end{aligned} \quad (80)$$

Note that here the last line contains the total effect of τ^{en} on the wage growth rate $\widehat{w}_{\tilde{\theta}}$, that is, the sum of the direct effect $D_{\tau^{en}} \widehat{w}$ and the directed technical change effect $D_{\phi, \tau^{en}} \widehat{w}$. The resource constraint (77) now implies

$$\int_{\underline{\theta}}^{\bar{\theta}} h_{\theta} D_{\tau^{en}} c_{\theta}(T) d\theta = \frac{d}{d\mu} F(l(T + \mu\tau^{en}), \phi^*(l(T + \mu\tau^{en}))) \Big|_{\mu=0}. \quad (81)$$

The principle of taxation says that every incentive compatible and resource feasible consumption-labor allocation can be implemented by some tax $\tilde{T}$. By implication, the allocation change $\{D_{\tau^{en}} l_{\theta}(T), D_{\tau^{en}} c_{\theta}(T)\}_{\theta \in \Theta}$ can be implemented by some reform $\tilde{\tau}$. Hence, a reform τ^{en} as analyzed above indeed exists.

Step 3. Having characterized the relevant consumption and labor input changes, we can now compare the welfare effect of reform τ^{ex} while ignoring directed technical change effects with the welfare effect of reform τ^{en} while accounting for directed technical change. Since the labor input changes are identical in both scenarios, the only difference in the two welfare effects stems from the different

consumption responses:

$$D_{\tau^{en}} \tilde{W}(T) - D_{\tau^{ex}} \tilde{W}(T) = \int_{\underline{\theta}}^{\bar{\theta}} g_{\theta} h_{\theta} (D_{\tau^{en}} c_{\theta}(T) - D_{\tau^{ex}} c_{\theta}(T)) d\theta . \quad (82)$$

From equations (78) and (80), the difference in consumption responses is, for every θ ,

$$D_{\tau^{en}} c_{\theta}(T) - D_{\tau^{ex}} c_{\theta}(T) = \int_{\underline{\theta}}^{\theta} v'(l_{\tilde{\theta}}(T)) l_{\tilde{\theta}}(T) D_{\phi, \tau^{en}} \hat{w}_{\tilde{\theta}}(T, \phi^*(T)) d\tilde{\theta} . \quad (83)$$

Here I used that the labor response is the same in both scenarios, such that the within-technology substitution effect on the return to skill is the same as well, that is,

$$D_{\tau^{ex}} \hat{w}_{\tilde{\theta}}(T, \phi^*(T)) = D_{\tau^{en}} \hat{w}_{\tilde{\theta}}(T, \phi^*(T)) .$$

Next, consider the directed technical change effect on the return to skill:

$$\begin{aligned} D_{\phi, \tau^{en}} \hat{w}_{\tilde{\theta}}(T, \phi^*(T)) &= D_{\phi, \tau^{en}} \left[\frac{d}{d\theta} \log(w_{\tilde{\theta}}(T, \phi^*(T))) \right] \\ &= \frac{d}{d\theta} [D_{\phi, \tau^{en}} \log(w_{\tilde{\theta}}(T, \phi^*(T)))] \\ &= \frac{d}{d\theta} \left[\frac{1}{w_{\tilde{\theta}}(T, \phi^*(T))} D_{\phi, \tau^{en}} w_{\tilde{\theta}}(T, \phi^*(T)) \right] . \end{aligned}$$

Since τ^{ex} is progressive, the labor response $(1/l_{\tilde{\theta}}) D_{\tau^{ex}} l_{\tilde{\theta}}(T)$ is decreasing in $\tilde{\theta}$ by Lemma 6. Hence, the identical response $(1/l_{\tilde{\theta}}) D_{\tau^{en}} l_{\tilde{\theta}}(T)$ decreases in $\tilde{\theta}$ as well. Then by Lemma 1, the directed technical change effect $(1/w_{\tilde{\theta}}) D_{\phi, \tau^{en}} w_{\tilde{\theta}}(T, \phi^*(T))$ must also decrease in $\tilde{\theta}$. We therefore obtain

$$\begin{aligned} 0 &\geq \frac{d}{d\theta} \left[\frac{1}{w_{\tilde{\theta}}(T, \phi^*(T))} D_{\phi, \tau^{en}} w_{\tilde{\theta}}(T, \phi^*(T)) \right] \\ &= D_{\phi, \tau^{en}} \hat{w}_{\tilde{\theta}}(T, \phi^*(T)) . \end{aligned}$$

By equation (83), this implies that the consumption difference $D_{\tau^{en}} c_{\theta}(T) - D_{\tau^{ex}} c_{\theta}(T)$ is decreasing in θ . Moreover, from equations (79) and (81), we have

$$\begin{aligned} \int_{\underline{\theta}}^{\bar{\theta}} h_{\theta} D_{\tau^{en}} c_{\theta}(T) d\theta &= \frac{d}{d\mu} F(l(T + \mu\tau^{en}), \phi^*(l(T + \mu\tau^{en}))) \Big|_{\mu=0} \\ &= \frac{d}{d\mu} F(l(T + \mu\tau^{ex}), \phi^*(l(T))) \Big|_{\mu=0} \\ &= \int_{\underline{\theta}}^{\bar{\theta}} h_{\theta} D_{\tau^{ex}} c_{\theta}(T) d\theta , \end{aligned}$$

where the second equality uses an envelope argument. By implication:

$$\int_{\underline{\theta}}^{\bar{\theta}} h_{\theta} (D_{\tau^{en}} c_{\theta}(T) - D_{\tau^{ex}}^{ex} c_{\theta}(T)) d\theta .$$

So, inspecting equation (82) reveals that if g_{θ} were constant in θ , we would have

$$D_{\tau^{en}} \tilde{W}(T) - D_{\tau^{ex}}^{ex} \tilde{W}(T) = 0 .$$

But since both g_{θ} and $D_{\tau^{en}} c_{\theta}(T) - D_{\tau^{ex}}^{ex} c_{\theta}(T)$ are decreasing in θ , we must have

$$\int_{\underline{\theta}}^{\bar{\theta}} g_{\theta} h_{\theta} (D_{\tau^{en}} c_{\theta}(T) - D_{\tau^{ex}}^{ex} c_{\theta}(T)) d\theta \geq 0$$

and thereby

$$D_{\tau^{en}} \tilde{W}(T) - D_{\tau^{ex}}^{ex} \tilde{W}(T) \geq 0 .$$

So,

$$D_{\tau^{en}} \tilde{W}(T) > 0 .$$

Step 4. Finally, we know that the labor response $(1/l_{\theta})D_{\tau^{en}} l_{\theta}(T)$ decreases in θ . Thus, we can again invoke Lemma 6 to obtain that τ^{en} must be progressive. We have thereby shown that

$$D_{\tau^{en}} \tilde{W}(T) > 0$$

for a progressive reform τ^{en} . So, $T \in \mathcal{T}$.

Since the preceding reasoning applies to any $T \in \mathcal{T}^{ex}$, we have shown that $\mathcal{T}^{ex} \subseteq \mathcal{T}$. □