Essays on Economic Inequality

Inauguraldissertation
zur
Erlangung des Doktorgrades
der
Wirtschafts- und Sozialwissenschaftlichen Fakultät
der
Universität zu Köln

vorgelegt von

M.Sc. Jonas Löbbing

aus Düsseldorf

im Jahr 2020

Referent: Prof. Dr. Peter Funk

Korreferent: Prof. Dr. Felix Bierbrauer

Tag der Promotion: 14. Juli 2020

Acknowledgments

I am indebted to my advisors Peter Funk and Felix Bierbrauer. Peter Funk has been instrumental in virtually all stages of the first chapter of this thesis. Also for the composition of the second chapter, his comments have been a great help. While not always convenient, his uncompromising views have more than once pushed my research in the right direction.

Felix Bierbrauer provided decisive input for the second chapter, especially in the later stages of the project. I am also grateful to him for two inspiring courses I attended in the early stage of my PhD studies, which sparked my interest in mechanism design and its applications to public policy.

I would also like to thank Daron Acemoglu, who hosted me for three months at the Massachussetts Institute of Technology in 2017. I am grateful to him not only for his hospitality but also for his reading of and insightful comments on the first two chapters of my thesis.

My fellow students and friends – and in one case (from the names there should be no confusion as to whom this refers) wife – Razi Farukh, Tobias Föll, Anna Kerkhof, and Sebastian Mertesacker provided invaluable support in various ways over the course of my studies. I am deeply thankful.

Special thanks also to Oliver Gürtler, who recognized my talent for research early and supported me on my way; to Emanuel Hansen, who was always willing to listen to my questions and ideas related to public economics; and to Raphael Flore, Philipp Giesa, Anna Hartmann, Eduardo Hidalgo, Matthias Kaldorf, Lucas Radke, Marius Vogel, and Florian Wicknig, who made my time at the Center for Macroeconomic Research a great and inspiring experience.

Moreover the following persons provided helpful comments on particular parts of my thesis: Spencer Bastani, Emmanuele Bobbio, Christian Bredemeier, Pavel Brendler, Sebastian Findeisen, Clemens Fuest, Christian Holzner, Andreas Irmen, Bas Jacobs, Eckhard Janeba, Sebastian Koehne, Keith Kuester, Claus Thustrup Kreiner, Etienne Lehmann, Ilse Lindenlaub, Benjamin Lockwood, Pascal Michaillat, Jakob Miethe, Johannes Münster, Andreas Peichl, Johannes Pfeifer, Dominik Sachs, Paul Schempp, Florian Scheuer, Dirk Sliwka, Stefanie Stantcheva, Michèle Tertilt, Uwe Thuemmel, Christian Traxler, Hitoshi Tsujiyama, Daniel Waldenström, Nicolas Werquin, Daniel Wiesen, Han Ye, and Xianjia Ye.

Finally, I would like to thank my family for their support, which had nothing to do with research but was every bit as important as what is mentioned above.

Contents

List of Tables

List of Figures

1 Introduction

Economic Inequality has been rising according to various measures and in many countries over the last four to five decades (e.g. Alvaredo, Chancel, Piketty, Saez and Zucman, 2018). Rising inequality is widely regarded as a major social problem. It is hence an important task for economic research to explore the ways in which economic policy can reduce inequality.

To solve a problem one should understand its roots. In the same way, to design policies that effectively address the rise in inequality, one should understand why inequality is rising.

The present thesis contributes to this understanding. Specifically, it advances our understanding of the forces that shape income inequality in a market-based economy and studies implications for the design of redistributive policy. Thereby, it extends the economic theory of the income distribution. The focus on economic theory is natural in so far as policy design requires counterfactual analysis; counterfactual analysis in turn requires theory.

1.1 Overview of the thesis

The thesis is composed of three, largely independent, research papers. The first two papers take a macroeconomic perspective, while the third paper zooms in on a particular aspect of income inequality, pursuing a more microeconomic approach.

Directed Technical Change In a market-based economy, production technologies operated by firms are a key determinant of the distribution of income. Hence, to understand changes in income inequality, an important step is to understand what drives changes in technologies.

This question is studied by the theory of directed technical change. According to the basic premise of this theory, which type of new technologies are introduced in an economy depends on the choices of profit-seeking agents (e.g., the management of firms, commercial researchers) and hence responds to the economic environment.

As a central result, the theory predicts that new technological developments are complementary to those input factors that are in abundant supply. Related to wage inequality, this means that the degree to which new technologies favor high- relative low-income earners depends on the relative supply of workers with qualifications that allow them to earn high incomes. If these workers are abundant relative to workers with less well rewarded qualifications, technologies will be developed that further raise their relative wages and thereby income inequality. Previous work has established this result for models that impose tight restrictions on the set of feasible technologies (allowing only for so-called labor-augmenting technologies, see Acemoglu 2007).

Chapter 2, based on the research paper "An Elementary Theory of Directed Technical Change and Wage Inequality", shows that the result is true in much more general environments. This is an important insight because digital technologies have enabled the development of machines that directly replace workers in a growing range of tasks. These labor-replacing technologies, however, are not well captured by models with labor-augmenting technology (e.g. Acemoglu and Restrepo, 2018a). Thus, my work makes central insights from directed technical change theory applicable to the growing set of models developed for the explicit study of such automation technologies.

Redistributive Income Taxes Given the strong impact of technology on the income distribution, an important question is how public policy can make firms use technologies that give rise to a more equal distribution of incomes. A natural approach is to consider taxes imposed directly on the use of certain technologies. Yet, recent research has shown that the scope for such direct taxes to reduce inequality is severely limited (see, e.g., the recent study of robot taxes by Thuemmel 2018). The main reason is that we lack reliable information about exactly which technologies, and applications thereof, are responsible for high inequality. Moreover, even if this information were available, monitoring the actual use of such technologies in firms may well be prohibitively costly.

Whenever policy cannot be directly conditioned on an undesirable activity – the use of disequalizing technology in the present case – economic theory prescribes to target complementary factors. Complementary to disequalizing technology are those factors that benefit from it, that is, high income workers.

This suggests to target the use of disequalizing production technologies indirectly via the income tax. Chapter 3, based on my research paper "Redistributive Income Taxation with Directed Technical Change", studies this policy option in detail. It first augments the canonical Mirrlees (1971) model of income taxation to include endogenous technology development and adoption choices of firms. Then, it analyzes how the presence of endogenous production technology alters the model's prescriptions about the optimal design of redistributive income taxes.

My main results show that directed technical change effects make the optimal tax more progressive, raising marginal tax rates at the top and lowering them at the bottom of the income distribution. The key mechanism is indeed that, by imposing higher tax rates on high income earners, labor supply of these workers is discouraged relative to the supply of low income earners. This induces firms to adopt and develop technologies that are more complementary to low-income occupations and hence reduces wage inequality before taxes. A more progressive tax therefore not only reduces post-tax but also pre-tax wage inequality. The optimal tax capitalizes on this pre-distributive effect. For reasonable calibrations of the model, the new effects are quantitatively significant: directed technical change reduces marginal tax rates on below-median incomes by up to 15 percentage points and increases them for above-median incomes by about half of this amount.

Credence Goods While income taxes are designed with the purpose to reduce inequality, other policies affect the income distribution less intentionally, but rather as a by-product of their intended effects. One such policy is the regulation of markets for expert services. Such markets, also called markets for credence goods, are subject to heavy regulation of prices and entry in virtually all industrialized countries. At the same time, occupations active on credence goods markets consistently rank among the top-earning occupations in most countries, a prime example being physicians and related occupations. Arguably, these high incomes reflect at least in parts the regulations imposed on the corresponding markets, ensuring high prices and low competitive pressure. From an equity perspective, deregulation of such markets may be a useful policy.

Chapter 4, based on the paper "Inefficiency and Regulation of Credence Goods Markets with Altruistic Experts", develops an argument against such deregulation. It proposes an efficiency-based rationale for regulation, specifically tailored to markets for credence goods. The defining characteristic of a credence good is that consumers cannot reliably assess its quality. Hence, they must trust the expert to provide an appropriate service. In our theory, there is indeed reason for such trust: experts are interested not only in their monetary payoff but also in the utility of their customers.

The key innovation is that experts' social motivation is income-dependent: the marginal rate of substitution between income and customer utility declines in income, such that the expert is more willing to forgo additional income to the benefit of customers when the amount already earned is high. In a common agency setting, where many consumers are served by the same expert, this creates an externality across consumers: one consumer's payment increases the expert's income, which in turn makes the expert more willing to forgo additional income in order to provide a higher quality service to all consumers. Under certain conditions, this externality implies that regulation that fixes prices above their competitive level achieves Pareto improvements. When market entry of experts is elastic with respect to profits, price regulation must be accompanied by entry restrictions to seize Pareto gains.

The externality we discover is fundamental in the sense that it arises in any moral hazard problem, provided that agents have non-linear social preferences and there is common agency.

Yet, its relevance depends on the availability of mechanisms for the provision of explicit monetary incentives. By considering the pure credence goods case, we focus on the extreme case where no explicit monetary incentives can be provided. This arguably makes the externality and its implications for regulation policy most relevant.

1.2 Contribution to Chapter 4

While Chapters 2 and 3 are based on research papers produced entirely by myself, Chapter 4 is based on joint work with Razi Farukh and Anna Kerkhof.

The research idea was developed in discussions between Razi Farukh and myself. I developed its formal representation and the proofs of our formal results. Razi Farukh and Anna Kerkhof wrote the first draft of the paper, which I revised.

2 An Elementary Theory of Directed Technical Change and Wage Inequality

Author: Jonas Löbbing

2.1 Introduction

Since the 1980s, many advanced economies have witnessed substantial increases in wage inequality between groups of workers with different levels of educational attainment. A broad empirical literature attributes parts of this increase to skill-biased technical change.[1] Appealing to skill-biased technical change as an exogenous explanation for the observed changes in the wage structure, however, is not entirely satisfactory. After all, the technologies that are used in an economy are eventually chosen by economic agents, about whose decisions economics should have something to say. This is the starting point for the theory of endogenously directed technical change (see Acemoglu, 1998; Kiley, 1999).[2] Central results of the theory predict how the skill bias of technical change depends on the supply of skills firms face in the labor market. In particular, they provide conditions under which (i) there is *weak relative equilibrium bias of technology* (weak bias, henceforth), meaning that any increase in the relative supply of skill induces skill-biased technical change, and (ii) there is *strong relative equilibrium bias of technology* (strong bias, henceforth), meaning that the positive effect of the induced technical change on the skill premium dominates the (typically negative) direct effect, such that the skill premium increases in relative skill supply (e.g. Acemoglu, 2002, 2007).[3] With the notable exception of Acemoglu (2007) (discussed below),

[1] See Bound and Johnson (1992), Katz and Murphy (1992), and Goldin and Katz (2008) on skill-biased technical change in general, and Graetz and Michaels (2018), Acemoglu and Restrepo (2019), and Dauth, Findeisen, Suedekum and Woessner (2017) on the effects of automation technology in particular.

[2] Henceforth, I use the terms "endogenously directed technical change", "directed technical change", and "endogenous technical change" equivalently.

[3] In a market economy, firms' technology adoption and development choices are based on the supply or demand curves they face in the markets they operate in. The supply of skills in the labor market is therefore a transmitter for the effects of many other variables on the skill bias of technical change. The analysis of such variables hence often relies on results that relate the skill bias of technical change to the supply of skills. An important example is given by the analysis of the effects of international trade on automation in Section 2.5.4.

these conditions are limited to settings in which aggregate production takes the specific form $F(\theta_1 L_1, \theta_2 L_2)$, where L_1 and L_2 denote the supply of skilled and unskilled labor, and θ_1 and θ_2 represent the endogenous, differentially labor-augmenting technology.

At the same time, the most recent literature on the effects of technical change on wage inequality analyzes labor-replacing (that is, automation) technology, typically in assignment models with labor and capital where capital takes the form of machines that perfectly substitute for labor in the production of tasks (e.g. Acemoglu and Autor, 2011; Autor and Dorn, 2013; Acemoglu and Restrepo, 2018a; Feng and Graetz, 2018; Aghion, Jones and Jones, 2017). In these models, the relevant technology variables can in general not be represented as labor-augmenting technology, such that they are outside the scope of the main results on directed technical change described above.

This paper generalizes the central results from directed technical change theory on weak and strong bias beyond the special case of differentially labor-augmenting technology and thereby makes them applicable to automation technology in Roy-like assignment models.[4] The first part of the paper derives general conditions for the phenomena of weak and strong bias that are independent of any functional form restriction, drawing on techniques from the theory of monotone comparative statics (Milgrom and Shannon, 1994). Besides making directed technical change theory applicable to automation technology, the results clarify the general mechanisms, based on simple notions of complementarity, that underlie the phenomena of weak and strong bias. The second part applies these results to obtain novel insights about the endogenous determination of automation technology in a Roy-like assignment model, with potential implications for redistributive labor market and trade policy.

The first part starts from a reduced form characterization of wages and equilibrium technology that is shown to arise from a range of different microfoundations of endogenous technical change, including standard approaches from endogenous growth theory. Building on this reduced form characterization, conditions are identified under which there is weak bias of technology, meaning that any increase in the relative supply of skill induces skill-biased technical change. The only essential condition is that the skill bias of technology is scale invariant, in the sense that a proportional change in the supply of all skill levels does not induce biased technical change. This is guaranteed by a restriction close to homotheticity of aggregate production in all labor inputs,

[4]At first glance, it may seem that Uzawa's theorem provides a justification for the restriction to labor-augmenting technology. But Uzawa's theorem only applies to the component of technology that grows over time on a balanced growth path, whereas the literature on endogenously directed technical change has mainly been concerned with the component of technology that is stationary on a balanced growth path, inducing changes in the stationary long-run distribution of (relative) wages. Moreover, with the labor share and the (risk-free) real interest rate declining over several decades (e.g. Karabarbounis and Neiman, 2014; Caballero, Farhi and Gourinchas, 2017), the general desirability for a model to generate balanced growth is no longer obvious.

which is remarkably weak compared to existing results (e.g. Acemoglu, 2007, Theorem 1). Most importantly, the restriction to differentially labor-augmenting technology from previous work can be deleted without replacement.[5]

While an increase in the relative supply of skill tends to induce skill-biased technical change, it also has a direct effect on the wage distribution, which typically depresses skill premia. The second set of results provides necessary and sufficient conditions for the occurrence of strong bias, meaning that the effect of the induced technical change dominates the direct effect such that skill premia increase with relative skill supply. It is shown that the induced technical change effect dominates everywhere if and only if the aggregate production function is quasiconvex. Reversely, if and only if aggregate production is quasiconcave, the direct effect dominates everywhere. These conditions provide an interesting analogy to endogenous growth theory, where convexity of aggregate production along rays through the origin (that is, increasing returns to scale) is required to generate persistent growth in a wide class of models (cf. Romer, 1986). As in these models, the aggregate (quasi-)convexity requirement discovered here has implications for the market structures needed in a model to analyze the case where skill premia increase in relative skill supply. In particular, either deviations from perfect competition or spillover effects across firms' technologies are needed.

While perhaps most natural in settings with two different levels of skill, all results in the first part of the paper also hold in settings with an arbitrary number of skill levels. Such settings allow to analyze technical change that is not monotonically skill-biased but causes the returns to skill to become, for example, more convex (a phenomenon often referred to as wage polarization in the literature). It turns out that, in principle, the techniques used to derive the monotone skill bias results can also be used to derive analogous results for non-monotone changes in the returns to skill.

The second part of the paper uses the techniques developed in the first part to derive novel predictions about the endogenous evolution of automation technology in the Roy-like assignment model proposed by Teulings (1995) (see Costinot and Vogel, 2010 for decisive progress in comparative statics for this model), augmented to incorporate capital as an additional production factor as in Acemoglu and Autor (2011) or Feng and Graetz (2018). In the model, a continuum of differentially skilled workers and capital, taking the form of machines that perfectly substitute for labor in the production of tasks, are assigned to a continuum of tasks, which in turn are combined to produce a single final good. In line with recent forecasts on the future automation potential

[5] The results in this part of the paper imply a LeChatelier Principle for relative demand curves, analogous to the conventional LeChatelier Principle that applies to absolute demand curves (e.g. Milgrom and Roberts, 1996). For an explicit formulation of the implied LeChatelier Principle for relative demand see Loebbing (2016), an earlier version of the present paper.

for different tasks (e.g. Frey and Osborne, 2017; Arntz, Gregory and Zierahn, 2016), machines are assumed to have comparative advantage in less complex tasks than labor, such that any increase in the set of tasks performed by machines (automation) displaces low-skilled workers from some of their previous tasks.[6]

The first result pertains to automation itself, as measured by the size of the set of tasks performed by machines. It says that any increase in the relative supply of skills induces automation, representing a skill-biased technical change. The induced automation, however, will never be strong enough to outweigh the initial direct effect of the increase in relative skill supply on the wage distribution, because aggregate production is quasiconcave. In consequence, low-skilled workers will always benefit in total from an increase in relative skill supply.

The second result endogenizes the productivity of machines. It shows that any increase in relative skill supply does not only stimulate automation but also investment into improving the productivity of machines, which in turn reinforces automation. The reinforcement between automation and machine productivity potentially reverses the result from the case with exogenous machine productivity: now, low-skilled workers' wages may decline, both relative to high-skilled workers' wages and in absolute terms, in response to an increase in the relative supply of skills. The reason is that the endogenous response of machine productivity "convexifies" the aggregate production function and may thus offset its quasiconcavity. These results provide a promising starting point for analyzing the interaction between labor market policies and automation, as many such policies (for example labor income taxation or unemployment insurance) affect firms primarily by changing the supply of workers they face.

The final and perhaps most important applied result considers the effect of trade in tasks between a skill-abundant, technologically advanced and a skill-scarce, technologically backward country. The trade analysis is a natural step within the assignment framework, because trade and changes in labor supply are in some sense equivalent here (see Costinot and Vogel, 2010). It turns out that trade with a skill-scarce country acts like a decrease in the relative supply of skills and hence reduces incentives to invest into automation technology in the skill-abundant country. Intuitively, trade makes the performance of low-skilled labor from abroad accessible to firms in the skill-abundant country. This reduces the incentives to automate tasks performed by low-skilled labor and hence, via the reinforcement mechanism, also reduces incentives to improve machine productivity. Analogously to the closed economy setting, this discouragement effect of trade on automation may be strong enough to overturn the standard Heckscher-Ohlin effect, according to which trade with a skill-scarce country raises skill premia at home. In consequence,

[6]This assumption is also broadly supported by recent estimates of the impact of industrial robots (Graetz and Michaels, 2018; Acemoglu and Restrepo, 2019) and a wider set of automation technologies in US manufacturing (Lewis, 2011) on the structure of employment and wages.

the overall effect of trade on low-skilled workers' wages in the skill-abundant country may turn out to be positive, both in absolute terms and relative to more skilled workers' wages.

From the perspective of the skill-scarce country, the standard Heckscher-Ohlin effect implies a reduction in skill premia. But there is a countervailing effect in this setting, because trade exposes low-skilled workers in the skill-scarce and technologically backward country to competition from the advanced machines of the skill-abundant country. Again, this automation-related effect may be sufficiently strong to overturn the Heckscher-Ohlin effect, such that skill premia in the skill-scarce country increase in response to trade opening.

Both findings are potentially relevant for trade policy. The negative effect of trade on automation casts doubt on policies that restrict trade with developing or emerging economies to protect low-skilled workers in developed countries. By stimulating automation, the desired effects of such policies may be severely mitigated or even reversed. The exposure of low-skilled workers in developing countries to competition from advanced foreign machines may provide a rationale for import restrictions on certain goods or comprehensive trade adjustment programs in developing countries. Real-world examples of such policies are the frequent exemptions from commitments to cut tariffs on agricultural imports granted to developing countries in various WTO negotiations on agricultural trade.[7]

The remainder of the paper is structured as follows. Section 2.2 introduces the reduced form characterization of wages and equilibrium technology that provides the basis for the general results on directed technical change in the following sections. Section 2.3 presents these results for the case with only two different levels of skill. Section 2.4 generalizes them to skill supply of arbitrary dimension. Section 2.5 applies the results to endogenous automation technology in assignment models, and Section 2.6 concludes.

Related Literature The paper has links to several strands in the existing literature. The first part of the paper extends the literature on directed technical change and wage inequality (e.g. Acemoglu, 1998, 2002 and Kiley, 1999), generalizing the key theoretical results of that literature. Most closely related to this analysis is Acemoglu (2007), who provides an endogenous technical change analysis on a similar level of generality. In contrast to the present paper, Acemoglu (2007) analyzes the effects of technical change induced by changes in the supply of a given skill level on the absolute wage of that skill, rather than on relative wages between different skills. From a purely theoretical perspective, the first part of the present paper can thus be viewed as the completion of a general theory of the effects of skill supply on the direction of technical change,

[7] Agricultural trade is a particularly fitting example, since agricultural production is highly automated in developed economies, but still very intensive in low-skilled labor in many developing countries (see e.g. de Vries, Timmer and de Vries, 2015).

with the first part on absolute wages given by Acemoglu (2007) and the second part on relative wages presented here. The analysis of relative wages is indispensable when the goal is to study implications of endogenous technical change for wage inequality.

The second part of the paper bridges the gap between the literature on directed technical change and the more recent strand of work on (exogenous) technical change and wage inequality in Roy-like assignment models (e.g. Costinot and Vogel, 2010 and Acemoglu and Autor, 2011).[8] The analysis of the effects of international trade on automation technology is related to existing work on international trade in assignment models (see Costinot and Vogel, 2015 for a survey of the use of assignment models in international economics), but also to Acemoglu (2003) who analyzes the effects of trade on directed technical change in a setting with differentially labor-augmenting technology. Most closely related to the second part of the paper are recent papers by Acemoglu and Restrepo (2018a), Hémous and Olsen (2018), Feng and Graetz (2018), Acemoglu and Restrepo (2018b), and Krenz, Prettner and Strulik (2018). Acemoglu and Restrepo (2018a) and Hémous and Olsen (2018) analyze the dynamic evolution of automation technology and its response to exogenous technology shocks rather than its response to changes in the structure of labor supply and international trade. Feng and Graetz (2018) provide an analysis similar to the first result on endogenous automation in the present paper, but they neither study endogenous investment into machine productivity nor the effects of international trade, both of which are crucial for the most important results on endogenous automation in this paper. Acemoglu and Restrepo (2018b) analyze the effects of the demographic structure on endogenous automation, with the focus on the effects of automation on productivity and the labor share. Finally, in parallel work Krenz et al. (2018) provide a joint theoretical analysis of offshoring and automation, but, unlike the model presented in Section 2.5.4, their model does not feature endogenous investment into the productivity of automation technology and hence (by the results of Section 2.4.2) cannot generate a reversal of the standard Heckscher-Ohlin effects, which is the main result of the analysis of the interplay between trade and automation in this paper.

2.2 A Simple Framework for Directed Technical Change

Consider a general equilibrium model with a continuum of firms and a continuum of workers.[9] Workers inelastically supply labor L and consume a single final good. They make no meaningful decisions. Firms are identical and produce the final good from labor according to a production

[8]See Acemoglu and Restrepo (2018c) for a list of advantages of the assignment approach with labor-replacing technology over the labor-augmenting technology approach in studying the effects of technical change on wage and income inequality.

[9]The model is identical to economy D from Acemoglu (2007).

function $F(L_i, \theta_i)$, where L_i is firm i's labor input and θ_i is a variable denoting firm i's production technology. The mass of firms is one.

Labor supply is differentiated according to skill levels s, that is, $L = \{L_s\}_{s \in S}$. Every L_s is a positive real number. The skill set can be of arbitrary size, that is, $S \subset \mathbb{R}$ is either a finite set or an interval. The technology variables θ_i are restricted to some set Θ. F and Θ satisfy the following assumption.

Assumption 1. *The set of feasible technologies Θ is compact. The production function $F(L, \theta)$ is continuous in θ, continuously differentiable in L, and the derivative $\nabla_L F(L, \theta)$ is strictly positive everywhere.*

Compactness of Θ requires that a topology is specified, which is presupposed. If S is finite, the derivative $\nabla_L F(L, \theta)$ is simply the gradient of F with respect to L, and every partial derivative is assumed to be strictly positive. If S is a continuum, $\nabla_L F(L, \theta)$ is the Gateaux derivative of F with respect to L, which can be represented as a real-valued function on S. This function is then assumed to be strictly positive at every s.

Under Assumption 3 it is straightforward to characterize an equilibrium in the model described above. Since changes in technology will be characterized by their effects on the wage distribution, it useful to define an *exogenous technology equilibrium* at first, where all firms' technologies are fixed at some $\theta \in \Theta$. In an exogenous technology equilibrium, firms choose their labor inputs L_i to maximize profits, taking wages $w = \{w_s\}_{s \in S}$ and their technologies $\theta_i = \theta$ as given. In a symmetric exogenous technology equilibrium, wages must satisfy

$$w(L, \theta) = \nabla_L F(L, \theta), \tag{2.1}$$

where the final good is used as the numéraire. Next, consider an *endogenous technology equilibrium*, where firms do not only choose their labor inputs to maximize profits, but also their technologies θ_i. Again, they take wages as given. In a symmetric endogenous technology equilibrium, the symmetric technology choice of firms, denoted θ^*, satisfies

$$\theta^*(L) \in \operatorname*{argmax}_{\theta \in \Theta} F(L, \theta). \tag{2.2}$$

Moreover, wages are given by $\nabla_L F(L, \theta^*(L))$, so equation (2.1) continues to hold: wages in an endogenous technology equilibrium are identical to wages in an exogenous technology equilibrium when technology is fixed at the equilibrium technology $\theta^*(L)$.

The model described above is special in that firms choose their production technologies independently of each other from an exogenous set. More elaborate models allow firms' technology choices to depend on each other and technologies to be developed and supplied to firms by a different set of agents, thus introducing a market for technology. Appendix A.2.1 shows that equations (2.1) and (2.2) continue to hold in such more sophisticated models. In particular, the appendix considers two models with endogenous production technology that follow standard modeling approaches from endogenous growth theory. The first model allows for spillovers across firms' technology choices, as in learning by doing models of endogenous growth (e.g. Romer, 1986; Lucas, 1988). The second model introduces a technology sector, where monopolistically competitive technology firms invest into the development of technologies and supply intermediate goods embodying their technologies to final good firms. This specification follows monopolistic competition based models of endogenous growth such as Romer (1990) and Aghion and Howitt (1992). In both models, wages and technologies are determined as by equations (2.1) and (2.2) in symmetric exogenous or, respectively, endogenous technology equilibria.[10]

Appendix A.2.1 also provides conditions for existence and uniqueness of symmetric equilibria. Remarkably, none of the models requires to impose any specific functional form restrictions on the production function F. The important difference between the simple baseline model described above and the more elaborate models in the appendix is that the former requires the endogenous technology production function $\overline{F}(L) := F(L, \theta^*(L))$ to be concave for a symmetric endogenous technology equilibrium to exist at all L. The more elaborate models only require concavity of $F(L, \theta)$ in L alone. The reason is that in the baseline model, equilibrium technologies and labor inputs are the joint outcome of individual firms' independent profit maximization problems, whereas in the more elaborate models they are an equilibrium outcome that arises from interdependent choices of multiple different agents. This distinction becomes relevant in Section 2.3.2 below.[11]

Since equations (2.1) and (2.2) provide a characterization of wages and equilibrium technology in a reasonably general class of models, the analysis in the first part of the paper builds on these equations, imposing Assumption 3. The goal is to answer the following questions.

[10] Acemoglu (2007) presents three more models (his economies C, M, and O) of endogenous technical change that satisfy equations (2.1) and (2.2). His models are related to but distinct from those presented in Appendix A.2.1.

[11] While all specific models presented in this paper are static, the models in Appendix A.2.1 can naturally be extended to dynamic versions, which generate constant growth paths with stationary relative wages between skill groups. These relative wages are then identical to the relative wages that prevail in equilibrium of the static model. The comparative statics results derived for the static class of models in the following sections can thus be interpreted as comparative statics on the constant growth path for a corresponding class of dynamic models. For an explicit treatment of dynamic models see Section 3.2 and Appendix B in Loebbing (2016).

Question 1 How do increases in the relative supply of skills affect the skill bias of technology?

Question 2 How do increases in the relative supply of skills affect skill premia (after adjustment of technology)?

According to equations (2.1) and (2.2), changes in the supply of skills L affect wages via two channels. First, when holding technology fixed, there is a direct effect as by equation (2.1). Second, the equilibrium technology $\theta^*(L)$ responds according to equation (2.2), which in turn affects wages as well. Question 1 asks for the second effect (the *induced technical change effect*, henceforth), while Question 2 asks for the combined impact of both effects on wages (the *total effect*). This distinction follows Acemoglu (2002, 2007) who also organizes his results around these two questions.

To pose the questions formally, precise definitions of an increase in relative skill supply and skill-biased technical change in environments with more than two skill levels are needed. Let an increase in relative skill supply be defined as an increase in skill supply ratios along the entire skill set.

Definition 1. An *increase in relative skill supply* is a change in labor supply from L to L' such that

$$\frac{L_{s'}}{L_s} \leq \frac{L'_{s'}}{L'_s}$$

for all $s \leq s'$.

We say that L has smaller relative skill supply than L' and write $L \preceq^s L'$.

Similarly, let a skill-biased technical change be a change in technology θ that raises skill premia along the entire skill set.

Definition 2. A *skill-biased technical change* is a change in technology from θ to θ' such that

$$\frac{w_{s'}(L,\theta)}{w_s(L,\theta)} \leq \frac{w_{s'}(L,\theta')}{w_s(L,\theta')}$$

for all $s \leq s'$ and all L.

We say that θ is less skill-biased than θ' and write $\theta \preceq^b \theta'$.

Moreover, if a wage vector w has lower skill premia along the entire skill set than another wage vector w' (such as $w(L,\theta)$ relative to $w(L,\theta')$ in Definition 2), it will sometimes be convenient to write $w \preceq^p w'$ for brevity. For the relations $\preceq^s$, $\preceq^b$, and $\preceq^p$, the corresponding strict relations $\prec^s$, $\prec^b$, and $\prec^p$ are defined as usual.

Finally note that, without further assumptions, the equilibrium technology $\theta^*(L)$ may not be uniquely determined by equation (2.2). While all results below could in principle be formulated in terms of sets of technologies or wages, this would substantially complicate the notation. It is therefore convenient to restrict attention to equilibria in which θ^* is the supremum of the set $\operatorname{argmax}_\theta F(L, \theta)$, where the supremum is taken with respect to the skill bias order $\preceq^b$. In all models of this paper that impose more structure on the technology set Θ (either in Appendix A.2.1 or in the applied Section 2.5), weak conditions guarantee that $\operatorname{argmax}_\theta F(L, \theta)$ is a singleton, so the selection of a unique $\theta^*(L)$ does not seem very restrictive.

2.3 Directed Technical Change with Two Skill Levels

Both for general expository reasons and for better comparability with existing results, it is convenient to start with an analysis of settings with only two different levels of skill. Suppose therefore that labor supply takes the form $L = (L_1, L_2) \in \mathbb{R}^2_{++}$, and let L_1 denote unskilled and L_2 skilled labor.

2.3.1 The Induced Technical Change Effect

First consider the induced technical change effect addressed in Question 1 above. The following result identifies sufficient conditions for any increase in relative skill supply to induce skill-biased technical change. This phenomenon is called weak relative bias of technology and proved for differentially labor-augmenting technology in Acemoglu (2007) (see Corollary 2 below).

Proposition 1 (Special case of Theorem 1). *Let $L \in \mathbb{R}^2_{++}$. Moreover, suppose that the equilibrium technology $\theta^*(L)$ is homogeneous of degree zero in L, and that any two technologies $\theta, \theta' \in \Theta$ can be ordered according to their skill bias, that is, either $\theta \preceq^b \theta'$ or $\theta' \preceq^b \theta$.*

Then, any increase in relative skill supply induces skill-biased technical change:

$$L \preceq^s L' \Rightarrow \theta^*(L) \preceq^b \theta^*(L').$$

Proof. The first step is to note that, starting with any change from L to L' that raises relative skill supply, the scale invariance (or zero homogeneity) of $\theta^*(L)$ allows to scale L' up or down without changing $\theta^*(L')$. We can therefore restrict attention to labor supply changes that keep output constant while holding technology fixed, that is, to changes from L to L' such that $F(L, \theta^*(L)) = F(L', \theta^*(L))$. In other words, we can assume without loss of generality that L'

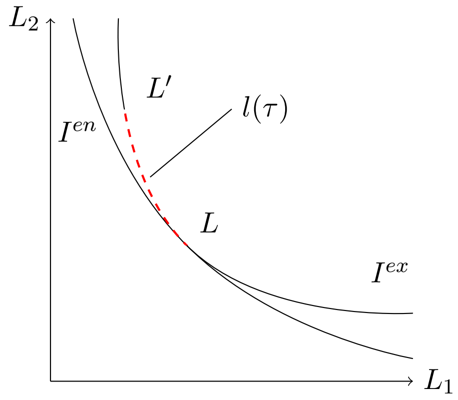

Figure 2.1. I^{ex} and I^{en} are the exogenous and the endogenous technology isoquants through L. The move from L to L' is parameterized by $l(\tau)$ such that $l(0) = L$ and $l(1) = L'$. Moving along $l(\tau)$ leaves output constant when $\theta^*(L)$ is fixed, but must raise output when technology is allowed to adjust (I^{en} is above I^{ex} everywhere). This means that technology adjusts in a way that is complementary to the rise in relative skill supply L_2/L_1.

is on the exogenous technology isoquant through L, which consists of all points L'' satisfying $F(L'', \theta^*(L)) = F(L, \theta^*(L))$.

Let $l(\tau) = (l_1(\tau), l_2(\tau))$ parameterize the path from L to L' along the exogenous technology isoquant of F. In particular, let $l(0) = L$, $l(1) = L'$, and $F(l(\tau), \theta^*(L)) = F(L, \theta^*(L))$ for all $\tau \in [0, 1]$. Since relative skill supply increases from L to L', the first entry of $l(\tau)$, $l_1(\tau)$, is decreasing, while the second entry, $l_2(\tau)$, is increasing in τ.

Figure 2.1 illustrates such a change along the exogenous technology isoquant I^{ex}. The dashed red segment of the exogenous technology isoquant is the image of the path $l(\tau)$.

For the second step note that, since $\theta^*(l(\tau))$ maximizes F at $l(\tau)$, we must have

$$F(l(0), \theta^*(L')) \leq F(l(0), \theta^*(L)) = F(l(1), \theta^*(L)) \leq F(l(1), \theta^*(L')). \tag{2.3}$$

In Figure 2.1, this corresponds to the exogenous technology isoquant I^{ex} being located above the endogenous technology isoquant I^{en}, which consists of all points L'' satisfying $F(L'', \theta^*(L'')) = F(L, \theta^*(L))$. If now both technologies are equally skill-biased, $\theta^*(L) \sim^b \theta^*(L')$, the statement of the theorem is true.[12] So we can restrict attention to cases with $\theta^*(L) \nsim^b \theta^*(L')$. In these cases, at least one of the two inequalities in (2.3) must be strict, because θ^* is selected as the supremum of the maximizer set in equation (2.2). (If both inequalities were equalities, we would either select

[12] The notation $\theta \sim^b \theta'$ means that both $\theta \preceq^b \theta'$ and $\theta' \preceq^b \theta$.

$\theta^*(L)$ at both $l(0)$ and $l(1)$, or $\theta^*(L')$.) This implies

$$F(l(0), \theta^*(L')) < F(l(1), \theta^*(L')),$$

and, by the mean value theorem, there exists a $\tau' \in (0, 1)$ such that

$$\begin{aligned} 0 &< \nabla_L F(l(\tau'), \theta^*(L')) \frac{\mathrm{d}\, l(\tau')}{\mathrm{d}\, \tau} \\ &= w_1(l(\tau'), \theta^*(L')) \frac{\mathrm{d}\, l_1(\tau')}{\mathrm{d}\, \tau} + w_2(l(\tau'), \theta^*(L')) \frac{\mathrm{d}\, l_2(\tau')}{\mathrm{d}\, \tau}. \end{aligned} \tag{2.4}$$

At the same time, by construction of $l(\tau)$, $F(l(\tau), \theta^*(L))$ is constant in τ, such that

$$w_1(l(\tau'), \theta^*(L)) \frac{\mathrm{d}\, l_1(\tau')}{\mathrm{d}\, \tau} + w_2(l(\tau'), \theta^*(L)) \frac{\mathrm{d}\, l_2(\tau')}{\mathrm{d}\, \tau} = 0. \tag{2.5}$$

Finally, rearranging and combining equations (2.4) and (2.5) yields

$$\frac{w_2(l(\tau'), \theta^*(L'))}{w_1(l(\tau'), \theta^*(L'))} > \frac{w_2(l(\tau'), \theta^*(L))}{w_1(l(\tau'), \theta^*(L))}. \tag{2.6}$$

Intuitively, if an increase in τ raises output at $\theta^*(L')$ by more than at $\theta^*(L)$, then the relative return to skilled labor must be greater under $\theta^*(L')$ as well. Since, by hypothesis, $\theta^*(L)$ and $\theta^*(L')$ can be ordered according to their skill bias, this implies that $\theta^*(L')$ is more skill-biased than $\theta^*(L)$, that is, $\theta^*(L) \preceq^b \theta^*(L')$. □

While the proposition applies to a wide range of specific models (see Appendix A.2.1 and the discussion in the previous section), it reveals a common thread across all of them: when relative skill supply increases, firms switch to technologies that are best suited to translate the increased availability of skilled (relative to unskilled) workers into output gains. Such technologies in turn are those under which the relative returns to skilled labor are high.

The conditions of the proposition are remarkably weak compared to existing results (see Corollary 2 below). The scale invariance condition ($\theta^*(L)$ is homogeneous of degree zero) is always satisfied when F is homogeneous in labor. Indeed, a condition slightly weaker than homogeneity is sufficient to guarantee scale invariance of θ^*.

Remark 1. Suppose $F(L, \theta)$ can be written as the composition of an inner function $f(L, \theta)$ that is linear homogeneous in L and an outer function $g(f, L)$ that is strictly increasing in f. Then, the set $\operatorname{argmax}_\theta F(L, \theta)$ and hence θ^* are homogeneous of degree zero in L.

The completeness condition (any two technologies can be ordered according to their skill bias) is only required because the theorem allows for changes in labor supply of arbitrary size. Once attention is restricted to local changes, it can be dropped without replacement.

Corollary 1. *Let $L \in \mathbb{R}^2_{++}$ and $\Theta \subset \mathbb{R}^N$ for arbitrary N. Suppose that $\theta^*(L)$ is homogeneous of degree zero and differentiable in L, and $w(L,\theta)$ is differentiable in θ.*

Then any local increase in relative skill supply induces skill-biased technical change:

$$\nabla_\theta \frac{w_2(L,\theta^*(L))}{w_1(L,\theta^*(L))} \nabla_L \theta^*(L) \, \mathrm{d}\, L \geq 0$$

for any L and $\mathrm{d}\, L$ such that $\mathrm{d}\, L_1/L_1 \leq \mathrm{d}\, L_2/L_2$.

Proof. The proof of the corollary replicates the proof of Proposition 1 with the tools of differential calculus and is provided in Appendix A.1.1 for completeness. □

Corollary 1 states that, under scale invariance of $\theta^*(L)$, the technical change $\nabla_L \theta^*(L) \, \mathrm{d}\, L$, induced by a local increase in relative skill supply $\mathrm{d}\, L$, raises the skill premium.

Results in the existing literature, in contrast, are restricted to differentially labor-augmenting technology, that is, to settings where F takes the form $F(\theta_1 L_1, \theta_2 L_2)$. The most general of these existing results can be obtained as a further corollary to Corollary 1.

Corollary 2 (cf. Theorem 1, Acemoglu, 2007)**.** *Let $L \in \mathbb{R}^2_{++}$ and $\Theta = \{\theta \in \mathbb{R}^2_{++} \mid C(\theta) \leq c\}$ for some constant $c > 0$ and a twice continuously differentiable, strictly convex, and homothetic function C with finite (but not necessarily constant) elasticity of substitution. Suppose that F can be written as $F(\theta_1 L_1, \theta_2 L_2)$, F is twice continuously differentiable, concave, and homothetic with finite (but not necessarily constant) elasticity of substitution.*

Then any local increase in relative skill supply induces skill-biased technical change.

Proof. Homotheticity of F, together with the labor-augmenting form of θ, guarantees scale invariance of $\theta^*(L)$. Moreover, the curvature and differentiability assumptions on F and C, together with finiteness of the elasticities of substitution, ensure differentiability of $\theta^*(L)$ and of wages $w(L,\theta)$. Therefore, all conditions of Corollary 1 are satisfied and its conclusion applies. □

The major restriction in Corollary 2 compared to Corollary 1 is that θ takes the labor-augmenting form $F(\theta_1 L_1, \theta_2 L_2)$. Corollary 1 shows that this restriction can be deleted without replacement. This is partly reassuring for existing work, as it shows that the restriction to labor-augmenting technologies is not essential for the most basic results on endogenous technical change and wage inequality. But more importantly, it allows to take these results to new types of models, especially to models with labor-replacing technologies as those discussed in Section 2.5.

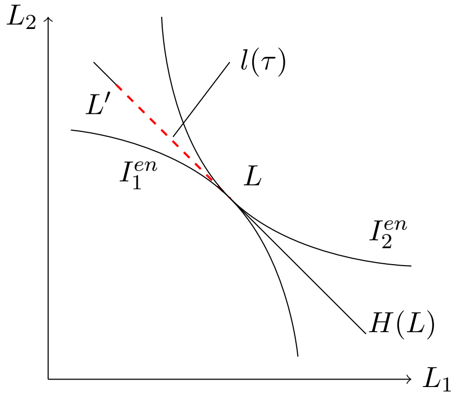

Figure 2.2. The figure shows two alternative endogenous technology isoquants, I_1^{en} for the case of a quasiconvex endogenous technology production function $\overline{F}$ and I_2^{en} for the quasiconcave case. In the quasiconcave case, moving from L to L' along the line $H(L)$ reduces $\overline{F}$ monotonically. Therefore, the ratio of marginal products w_2/w_1 must be below its initial value at L on the entire way to L'. The opposite conclusion applies to the quasiconvex case.

2.3.2 The Total Effect

The preceding analysis shows that under fairly general conditions any increase in relative skill supply induces skill-biased technical change. But the direct effect of an increase in relative skill supply on the skill premium, holding technology constant, is typically negative, so the two effects counteract each other. The following result provides exact conditions under which the induced technical change effect dominates the direct effect or, in the words of Question 2 above, under which an increase in relative skill supply raises the skill premium after adjustment of technology. This phenomenon is called strong relative bias of technology in previous work, and again results only exist for differentially labor-augmenting technology (cf. Acemoglu, 2007).

Proposition 2 (Special case of Theorem 3). *Let $L \in \mathbb{R}^2_{++}$ and suppose that $w(L, \theta^*(L))$ is homogeneous of degree zero in L. Then, there exists an increase in relative skill supply that strictly raises the skill premium, formally: $\exists L \preceq^s L'$ such that $w(L, \theta^*(L)) \prec^p w(L', \theta^*(L'))$, if and only if the endogenous technology production function $\overline{F}(L) := F(L, \theta^*(L))$ is not quasiconcave.*

Moreover, any increase in relative skill supply raises the skill premium, $L \preceq^s L' \Rightarrow w(L, \theta^*(L)) \preceq^p w(L', \theta^*(L'))$, if and only if $\overline{F}$ is quasiconvex.

Sketch of proof. The full proof is given in Appendix A.1.1. A sketch of it is provided here to convey its main idea.

The first step is to note that, starting from any labor supply L, the scale invariance (or homogeneity of degree zero) of $w(L, \theta^*(L))$ allows to restrict attention to changes along the line $H(L)$ that is tangent to the endogenous technology isoquant through L, that is, to the isoquant of the endogenous technology production function $\overline{F}$ (defined in the proposition). Figure 2.2 shows two alternative shapes of the endogenous technology isoquant through L. In one case (I_1^{en}) $\overline{F}$ is quasiconvex, in the other case (I_2^{en}) $\overline{F}$ is quasiconcave. In the quasiconcave case, $\overline{F}$ is decreasing along the path $l(\tau)$ that runs along $H(L)$ from L to L', in the quasiconvex case it is increasing along this path.

Second, at any point on $l(\tau)$, an infinitesimal move in direction of L' along the line $H(L)$ will decrease (increase) $\overline{F}$ if and only if the marginal gain from increasing L_2 exceeds (falls short of) the marginal loss from decreasing L_1. At L these two effects cancel each other exactly by construction of $H(L)$. Therefore, if $\overline{F}$ decreases (increases) along $l(\tau)$, the ratio of marginal products of L_2 and L_1, w_2/w_1, must be smaller (greater) at any point on $l(\tau)$ than at L. It follows that the endogenous technology skill premium $w_2(l(\tau), \theta^*(l(\tau)))/w_1(l(\tau), \theta^*(l(\tau)))$ falls (rises) in relative skill supply if $\overline{F}$ is quasiconcave (quasiconvex). This provides for the "only if" statement in the first part of the proposition and for the "if" statement in the second part.

The converse statements are obtained by noting that any failure of quasiconcavity (quasiconvexity) allows to find L, $H(L)$, and an L' on $H(L)$ such that $\overline{F}$ must increase (decrease) in direction of L' at some point on the line segment between L and L'. □

The requirement that $w(L, \theta^*(L))$ is homogeneous of degree zero in L can be ensured by slightly strengthening the condition in Remark 1.

Remark 2. Suppose $F(L, \theta)$ can be written as the composition of an inner function $f(L, \theta)$ that is linear homogeneous in L and an outer function $g(f)$ that is strictly increasing in f. Then, the endogenous technology wages $w(L, \theta^*(L))$ are homogeneous of degree zero in L.

Proposition 2 provides an exact link between skill premia that increase in relative skill supply (strong relative bias) and curvature properties of the aggregate production function $\overline{F}$. This reveals an interesting theoretical analogy to endogenous growth theory. There, increasing returns to scale in aggregate production are necessary for persistent growth in a large class of models (cf. Romer, 1994; Acemoglu, 2009). While increasing returns to scale constitute a failure of concavity along lines through the origin, the failure of concavity required in Proposition 2 concerns the contour sets of $\overline{F}$ and is in this sense orthogonal to returns to scale.

From an applied modeling perspective, Proposition 2 is informative about how to set up a model to analyze the case where skill premia increase in relative skill supply. In particular, it says that one must depart from the baseline model of endogenous technology choices presented in

Section 2.2, where identical firms choose their technologies independently of each other and hence production functions must be (at least locally) concave in labor and technology. Appendix A.2.1 and A.2.1 discuss two such ways of departure, both of which introduce some form of interdependence between firms' technology choices. The first does so in an ad hoc way, assuming spillovers between firms' technologies without further specifying them. In the second model, interdependence occurs via the market for technologies or innovations, where technology firms supply their innovations to final good firms, and non-rivalry of innovations implies that technology firms sell their ideas to all active final good firms at once (see Appendix A.2.1 for details on these models). In both cases, interdependence between firms' technologies breaks the requirement that production functions are jointly concave in labor and technology, and hence allows for the failure of concavity required by Proposition 2.

One can also interpret the baseline model presented in Section 2.2 as describing a process of pure technology adoption whereas the other models incorporate some features of true innovation (such as spillovers from imperfect protection of an individual firms' knowledge, or imperfect competition from the partial protection of intellectual property). Then, Proposition 2 admits the conclusion that technology adoption alone is not sufficient for strong relative bias of technology. Some portion of innovation is needed for this to occur.

Previous work has considered a local version of strong relative bias. In the setting with labor-augmenting technology, Acemoglu (2007) shows that this local version arises if and only if the elasticity of substitution between the two arguments of the function $F(\theta_1 L_1, \theta_2 L_2)$ exceeds some threshold value. Since the labor-augmenting technology setting is a special case of my analysis, Proposition 2 implies that the elasticity of substitution crosses this threshold exactly when the upper contour sets of $\overline{F}$ change their curvature from convex to concave. While the relation between skill premia that increase in relative skill supply and curvature properties of the aggregate production function could already be anticipated from the specific existing results, Proposition 2 formulates this relation precisely.

2.4 Directed Technical Change with Multiple Skill Levels

Consider now the general case with arbitrarily many skills. The next two subsections present the more general theorems behind Propositions 1 and 2, while the third subsection discusses how to extend these results to non-monotone changes in skill supply and skill premia. The main insight is that the results from the previous section are not specific to the two skills case but generalize quite naturally to settings with arbitrarily many skill levels.

2.4.1 The Induced Technical Change Effect

For the two skills case, Proposition 1 provides conditions under which any increase in the relative supply of skilled to unskilled labor induces skill-biased technical change. The following result shows that the statement of Proposition 1 holds under exactly the same conditions for many skills.

Theorem 1. *Suppose that the equilibrium technology $\theta^*(L)$ is homogeneous of degree zero in L, and that any two technologies $\theta, \theta' \in \Theta$ can be ordered according to their skill bias, that is, either $\theta \preceq^b \theta'$ or $\theta' \preceq^b \theta$.*

Then, any increase in relative skill supply induces skill-biased technical change:

$$L \preceq^s L' \Rightarrow \theta^*(L) \preceq^b \theta^*(L').$$

Proof. See Appendix A.1.1. □

Recalling the definitions of increases in relative skill supply and skill-biased technical change for multiple skill environments, Theorem 1 says that an increase in all supply ratios of more versus less skilled workers induces technical change that raises all skill premia in the model.

In the application Section 2.5, a slightly different version of Theorem 1 turns out to be useful. This second version builds on a somewhat less demanding definition of skill-biased technical change. Indeed, from an economic point of view, Definition 2 appears unnecessarily strong, as it requires technical change to raise all skill premia at any point of the labor supply space to qualify as skill-biased. What matters from an applied perspective, however, is that the change in technology raises skill premia at those labor supply levels where it can actually happen; that is, at those labor supplies where it increases aggregate production F. This leads to the following alternative definition of skill-biased technical change.

Definition 3. A skill-biased technical change is a change in technology from θ to θ' such that

$$F(L, \theta) \leq F(L, \theta') \Rightarrow w(L, \theta) \preceq^p w(L, \theta').$$

We write $\theta \preceq^{b'} \theta'$.

Using this definition, Theorem 1 can be restated as follows.

Theorem 2. *Suppose that the equilibrium technology $\theta^*(L)$ is homogeneous of degree zero in L, and that any two technologies $\theta, \theta' \in \Theta$ can be ordered according to their skill bias following the alternative Definition 3, that is, either $\theta \preceq^{b'} \theta'$ or $\theta' \preceq^{b'} \theta$.*

Then, any increase in relative skill supply induces skill-biased technical change according to Definition 3:

$$L \preceq^s L' \Rightarrow \theta^*(L) \preceq^{b'} \theta^*(L').$$

Proof. See Appendix A.1.1. □

Theorem 2 provides somewhat more flexibility in applications, which is important especially because the condition that any two technologies can be ordered according to their skill bias can be quite restrictive in models with multiple skill types. Section 2.5.3 demonstrates its usefulness in a case where Theorem 1 would not be applicable.

Finally, Appendix A.2.2 shows that there is still some slack in the conditions of Theorems 1 and 2. Indeed, both the scale invariance condition on θ^* and the condition that any two technologies can be ordered according to their skill bias can be slightly relaxed. Since this discussion is mainly technical and does not play a role in the application part of the paper, it is deferred to the appendix. The appendix also clarifies the relation of Theorem 1 to the main results of monotone comparative statics developed by Milgrom and Shannon (1994).

2.4.2 The Total Effect

Considering the total effect of an increase in relative skill supply, Proposition 2 says for the two skills case that there exists an increase in relative skill supply that raises the skill premium if and only if the endogenous technology production function $\overline{F}$ fails to be quasiconcave. Moreover, any increase in relative skill supply raises the skill premium if and only if $\overline{F}$ is quasiconvex. The following result extends these insights to the general case with arbitrarily many skills.

Theorem 3. *Suppose that $w(L, \theta^*(L))$ is homogeneous of degree zero in L. Then, the following statements hold.*

(1) *If there exists an increase in relative skill supply that strictly raises skill premia, formally: $\exists L \preceq^s L'$ such that $w(L, \theta^*(L)) \prec^p w(L', \theta^*(L'))$, then $\overline{F}$ is not quasiconcave.*

 Moreover, if $\overline{F}$ is not quasiconcave along some line in direction of $\preceq^s$, then there exists an increase in relative skill supply that does not lower all skill premia, formally: $\exists L \preceq^s L'$ such that $w(L', \theta^*(L')) \npreceq^p w(L, \theta^*(L))$.

(2) *If it holds that any increase in relative skill supply raises all skill premia, formally: $L \preceq^s L' \Rightarrow w(L, \theta^*(L)) \preceq^p w(L', \theta^*(L'))$, then $\overline{F}$ is quasiconvex along all lines in direction of $\preceq^s$.*

 Moreover, if $\overline{F}$ is quasiconvex, then no increase in relative skill supply will lower all skill premia, formally: $L \preceq^s L' \Rightarrow w(L', \theta^*(L')) \not\prec^p w(L, \theta^*(L))$.

Proof. See Appendix A.1.1. □

The first statement in Part 1 of the theorem replicates the only if part of the first part of Proposition 2: only if $\overline{F}$ is not quasiconcave, there can be an increase in relative skill supply that strictly raises skill premia. This result captures the most important insight from Section 2.3.2 and extends it to the many skills case. It implies that one has to use models in which aggregate production may fail to be quasiconcave to analyze cases where skill premia increase in relative skill supply. As discussed in detail in Section 2.3.2, the possibility of a failure of quasiconcavity in aggregate production is closely linked to the specific mechanisms that determine equilibrium technologies in the model.

The converse of this main result, however, does not extend one-to-one to the many skills environment. The reason for this is twofold. First, with high-dimensional skill supply, $\overline{F}$ may fail to be quasiconcave in directions orthogonal to changes in relative skill supply. Such failures of quasiconcavity do not have immediate consequences for the response of skill premia to increases in relative skill supply and hence do not admit the conclusion that skill premia do not decrease in relative skill supply. Second, in the two skills case, if the skill premium does not fall in relative skill supply, it must necessarily increase, as it is one-dimensional. With many skills, however, there may be instances where skill premia fall in relative skill supply in some ranges of skill but increase in other ranges. The partial converse offered by Theorem 3 thus (i) restricts attention to cases where $\overline{F}$ fails to be quasiconcave along lines in direction of changes in relative skill supply and (ii) says that in such cases not all skill premia fall when relative skill supply increases. When restricted to two skill levels, this statement becomes a full converse, so Part 1 of Theorem 3 covers the first part of Proposition 2 as a special case.

Analogous adjustments are required in Part 2 of Theorem 3 to extend the second part of Proposition 2 to the many skills environment. Again, once attention is restricted to two skill groups, Part 2 of Theorem 3 becomes an "if and only if" statement that replicates the second part of Proposition 2 exactly.

The main takeaway from Theorem 3 is that the principal insight from the two skills case regarding the type of models needed to analyze the case where skill premia increase in relative skill supply extends to environments with arbitrarily many skills. The same holds for the analogy to the non-concavities required for persistent growth in endogenous growth theory (see Section 2.3.2 for detailed discussion).

2.4.3 Non-Monotonically Biased Technical Change

The previous discussion was focused on skill supply and wage changes that are monotone in skill, in particular on increases in relative skill supply along the entire skill set and on increases in all skill premia. Yet the results are more versatile than it may seem at first glance. This is because none of them requires wages to increase in the skill index s. Hence, the interpretation of a higher s as denoting a more skilled type of labor is not implied by any of the formal results so far.

Consider for example a three skill setting with $S = \{1, 2, 3\}$. We can now interpret L_1 as the supply of middle-skill workers, L_2 as low-skill, and L_3 as high-skill workers. Then, under the conditions of Theorem 1, any change in labor supply such that the low versus middle-skill ratio and the high versus low-skill ratio increase will induce polarizing technical change; that is, technical change that raises low-skill workers' wages relative to middle-skill wages and high-skill wages relative to low-skill wages.

The common notion of wage polarization, however, does not contain any restriction as to whether high-skill wages increase relative to low-skill wages or vice versa. Accordingly, the following definition of polarizing technical change dispenses with such a restriction. The notation again follows the convention from the previous sections whereby a higher index denotes a more skilled type of labor, that is, L_1 denotes low-skilled, L_2 middle-skilled, and L_3 high-skilled labor.

Definition 4. Let $L = (L_1, L_2, L_3) \in \mathbb{R}^3_{++}$. A *polarizing technical change* is a change in technology from θ to θ' such that

$$\frac{w_2(L,\theta)}{w_1(L,\theta)} \geq \frac{w_2(L,\theta')}{w_1(L,\theta')} \quad \text{and} \quad \frac{w_3(L,\theta)}{w_2(L,\theta)} \leq \frac{w_3(L,\theta')}{w_2(L,\theta')}$$

for all L.

We say that θ is less polarizing than θ'.

When adopting such a broader definition of polarizing technical change, the loss of information about the change in the high to low-skill relative wage implies that the set of skill supply changes for which we can sign the effect on polarizing technical change becomes smaller. The following result identifies such a set for the three skills case.

Theorem 4. *Let $L \in \mathbb{R}^3_{++}$. Moreover, suppose that the equilibrium technology $\theta^*(L)$ is homogeneous of degree zero in L, and that for any two technologies $\theta, \theta' \in \Theta$, either θ is less polarizing than θ' according to Definition 4, or vice versa.*

Then, any change in labor supply from L to L' such that low- and high-skilled labor supply change proportionately to each other and increase relative to middle-skilled labor induces polarizing technical change.

Formally, for any L and L' with

$$\frac{L_2}{L_1} \geq \frac{L'_2}{L'_1} \quad \textit{and} \quad \frac{L_3}{L_1} = \frac{L'_3}{L'_1}$$

it holds that

$$\frac{w_2(L', \theta^*(L))}{w_1(L', \theta^*(L))} \geq \frac{w_2(L', \theta^*(L'))}{w_1(L', \theta^*(L'))} \quad \textit{and} \quad \frac{w_3(L', \theta^*(L))}{w_2(L', \theta^*(L))} \leq \frac{w_3(L', \theta^*(L'))}{w_2(L', \theta^*(L'))}.$$

Proof. See Appendix A.1.1. □

Theorem 4 applies to polarized changes in skill supply that are balanced, in the sense that the relative increases in the supply of low versus medium and high versus medium skills must be equal in size. It can be shown by examples that this restriction cannot be dispensed with. Yet, as discussed above, when adopting a more exclusive definition of polarizing technical change that signs all relative wage changes, the balancedness restriction can be dropped and the results of the previous sections are readily applicable.

Since the focus of the paper is on monotonically skill-biased technical change, a more general treatment of polarizing technical change beyond the three skills case seems inept here. Note at this point, however, that there are a number of reasons for the focus on monotonically skill-biased technical change. First, empirically, when identifying skill by education level, skill supply changes in most developed economies have taken the form of monotone increases in relative skill supply during the last decades. Second, there is evidence that technical change has been monotonically skill-biased at least in the United States over the last four decades (Sevinc, 2018).[13] Third, the application part of the paper focuses on automation technology and the existing evidence on the impact of automation technologies on the wage distribution supports the view that this impact is monotonically skill-biased (Lewis, 2011; Acemoglu and Restrepo, 2019; Dauth et al., 2017). In addition, recent attempts to forecast the future potential for automation across occupations find that the risk of automation decreases monotonically with average occupational education levels, suggesting a monotonic skill bias of anticipated future automation (e.g. Frey and Osborne, 2017; Arntz et al., 2016). Fourth, the analysis of changes in relative skill supply serves as a starting point for analyzing the effects of further potential determinants of the skill bias of technical change such

[13]This does not contradict the observation that, over some periods, wage growth has been polarized across occupations, with medium-paying occupations having experienced the smallest mean wage growth (as documented, for example, in Autor, Katz and Kearney, 2006 and Autor and Dorn, 2013). Both findings are reconciled through the fact that average skill and average wages are somewhat disconnected across occupations in the bottom part of the occupational wage distribution, potentially due to systematic differences in non-wage amenities of jobs (Sevinc, 2018).

as international trade. Section 2.5.4 shows that in certain environments trade with a skill-scarce country acts like a monotone decrease in relative skill supply, so the results of Sections 2.3.1 to 2.4.2 apply.

2.5 Endogenous Automation Technology

Differentially labor-augmenting technology, as analyzed by previous work on directed technical change, is a fairly abstract concept. Its relation to common intuitive notions of technical change is loose at best. Moreover, it cannot deliver results on technical change that are directly testable in empirical work, because labor-augmenting technology variables have no directly measurable empirical counterpart.[14]

A more concrete formalization of technical change is given by labor-replacing technology in models with a flexible assignment of production factors to tasks. Such models formalize the intuitive notion that technical progress allows the production of machines that take over tasks previously performed by human labor.

An endogenous technical change analysis in this type of model thus has a number of benefits over the labor-augmenting technology approach.[15] First, the results align well with intuitive notions of technical change. Second, they can be tested directly in empirical work, as labor-replacing technology variables can be identified with empirical measures of concrete automation technologies.[16] Third, they make statements about a form of technical change that is widely perceived to be among the most important determinants of future changes in the employment and wage structure. Finally, the literature on assignment models of the type analyzed here is growing rapidly, with applications in labor (e.g. Acemoglu and Autor, 2011), trade (e.g. Costinot and Vogel, 2010), growth (e.g. Acemoglu and Restrepo, 2018a), and public economics (e.g. Rothschild and Scheuer, 2013). Bringing results on directed technical change to the assignment environment keeps them connected to the newest strand of the theoretical literature on wage and income inequality.

The following sections conduct such a directed technical change analysis in assignment models,

[14] Consequently, empirical examinations of models of endogenous labor-augmenting technology are restricted to the reduced form relationship between labor inputs and wages, which captures both the direct and the induced technical change effect and hence does not allow for precise conclusions about either one of them (e.g. Blum, 2010; Morrow and Trefler, 2017; Carneiro, Liu and Salvanes, 2019).

[15] See Acemoglu and Restrepo (2018c) for a complementary list of advantages of the labor-replacing technology approach.

[16] See for example the use of counts of industrial robots as a measure for automation technology by Graetz and Michaels (2018); Acemoglu and Restrepo (2019); Dauth et al. (2017); Abeliansky and Prettner (2017); Acemoglu and Restrepo (2018b), the use of survey data on the adoption of various automation technologies in manufacturing by Lewis (2011), and the use of data on harvesting machines in agriculture by Clemens, Lewis and Postel (2018).

applying the results developed in the previous part of the paper.

2.5.1 Setup

The analysis builds on the assignment model by Teulings (1995), augmented to incorporate capital as an additional production factor. There is a continuum of tasks (or intermediate goods), indexed by $x \in X = [\underline{x}, \overline{x}]$, and a single final good. Final good producers produce the final good out of tasks according to

$$Y = \beta \left(\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \, \mathrm{d}\, x \right)^{\frac{\epsilon}{\epsilon-1}} \tag{2.7}$$

with $\beta > 0$, and $\epsilon > 0$ being the elasticity of substitution across tasks. Task producers produce tasks linearly from capital and labor,

$$Y_x = \alpha(x) K_x + \int_{\underline{s}}^{\overline{s}} \gamma(s, x) L_{s,x} \, \mathrm{d}\, s,$$

where K_x denotes the amount (or density) of capital assigned to task x, $L_{s,x}$ is the amount of labor of skill s assigned to task x (or the joint density of labor over skills and tasks), and $\alpha(x)$ and $\gamma(s, x)$ are task specific productivities of capital and the differentially skilled types of labor.

There is a continuum of skills, indexed by $s \in S = [\underline{s}, \overline{s}]$, and labor supply $\{L_s\}_{s \in S}$ (or the marginal density of labor over skills) is exogenous. The total amount of capital is denoted by $K = \int_{\underline{x}}^{\overline{x}} K_x \, \mathrm{d}\, x$. Capital is produced at marginal cost r from final good. This mimics the steady-state of dynamic models in which capital is accumulated over time and the long-run interest rate is fixed by preferences and depreciation.

The final good is the numéraire, task prices are denoted by p_x, wages by w_s, and the price of capital by p_c. All firms maximize profits and all markets are competitive.

An equilibrium consists of wages, task prices, a price for capital, a joint distribution of labor over tasks and skills, and distributions of capital and task output levels over tasks such that all markets clear given profit maximizing behavior by firms.[17] To simplify a more detailed characterization of equilibrium, some of its basic properties are derived first.

The pattern according to which capital and skills are assigned to tasks is determined by comparative advantage and hence by the shape of the productivity schedules $\alpha(x)$ and $\gamma(s, x)$. Let these schedules be strictly positive, twice differentiable, and satisfy the following comparative advantage assumption.

[17] Note that workers, who consume the final good, supply labor inelastically, and own the firms, do not have any meaningful choices, so they are omitted from the exposition.

Assumption 2. *More skilled workers have comparative advantage in higher x (henceforth, more complex) tasks, that is,*

$$\frac{\gamma(s,x')}{\gamma(s,x)} < \frac{\gamma(s',x')}{\gamma(s',x)}$$

for all $s < s'$ and $x < x'$.

Moreover, all workers have comparative advantage over capital in more complex tasks, that is,

$$\frac{\alpha(x')}{\alpha(x)} < \frac{\gamma(s,x')}{\gamma(s,x)}$$

for all s and $x < x'$.

The assumption about comparative advantage across skills gives a meaning to the task index x. A higher x now indicates a task in which more skilled workers have comparative advantage. In this sense x can be viewed as a measure of a task's complexity.

The assumption about the comparative advantage between capital and workers, in contrast, is more restrictive. It implies that capital will always perform a set of least complex tasks while workers sort into tasks of higher complexity. Low-skilled workers will thus always be the first to lose their tasks to machines when automation technology advances. Though restrictive, there are good reasons for this assumption in the present context. First, empirical studies suggest that the use of industrial robots, an important form of automation technology in the manufacturing sector, has negative effects on low-skilled workers' wages and employment shares, while results for medium-skilled workers are ambiguous and high-skilled workers may gain somewhat on both margins (see Graetz and Michaels, 2018; Acemoglu and Restrepo, 2019).[18] Second, recent forecasts of the future potential for automation across occupations predict invariably that the risk of automation decreases almost monotonically with average education levels of workers in a given occupation (Frey and Osborne, 2017; Arntz et al., 2016; Nedelkoska and Quintini, 2018). Third, Lewis (2011) shows empirically that investment into various automation technologies in US manufacturing in the 1980s and 1990s was a substitute for the least-skilled but a complement to medium-skilled workers.

As already noted, the comparative advantage assumption has clear implications for the sorting of capital and workers into tasks. In particular, it implies that there is a threshold task $\widetilde{x}$ such that capital performs all tasks below $\widetilde{x}$ while labor performs all tasks above. Moreover, more

[18] In more detail, Graetz and Michaels (2018) analyze a panel of industrialized countries and find negative (positive) effects of robot use on the share of hours worked and the wage bill share of low-skilled (high-skilled) workers, whereas results for medium-skilled workers are insignificant. Acemoglu and Restrepo (2019) find that across US commuting zones the effects of robots on wages and employment to population ratios are monotonic over five education groups, with the largest negative effects for the least educated group. Observation periods in both studies start in 1993 and end in 2005 (Graetz and Michaels) and 2007 (Acemoglu and Restrepo).

skilled workers perform more complex tasks, such that the assignment of skills to tasks can be summarized by a unique matching function $m(s)$, which assigns a task to each skill s.

Lemma 1. *In any equilibrium, there exists an automation threshold* $\widetilde{x} \in X$ *and a strictly increasing and continuous matching function* $m : S \to [\widetilde{x}, \overline{x}]$ *such that*

$$\begin{aligned} L_{s,x} > 0 \quad &\textit{if and only if} \quad x = m(s) \\ K_x > 0 \quad &\textit{if and only if} \quad x < \widetilde{x}. \end{aligned}$$

Proof. See Appendix A.1.1. □

This representation of the assignment of factors to tasks allows to give a detailed characterization of equilibrium in terms of the automation threshold $\widetilde{x}$ and the matching function m. Accordingly, an equilibrium consists of

- an automation threshold $\widetilde{x}$, a matching function $m : S \to [\widetilde{x}, \overline{x}]$, an assignment of capital to tasks $\{K_x\}_{x \in X}$, and task output $\{Y_x\}_{x \in X}$;
- task prices $\{p_x\}_{x \in X}$, wages $\{w_s\}_{s \in S}$, and a capital price p_c;

such that

(E1) $Y_x = \alpha(x) K_x$ if $x < \widetilde{x}$ and $Y_x = \gamma(m^{-1}(x), x) L_{m^{-1}(x)} \frac{\mathrm{d}\, m^{-1}(x)}{\mathrm{d}\, x}$ if $x \geq \widetilde{x}$; (market clearing)

(E2) $p_x = \frac{\partial Y}{\partial Y_x}$ for all x, where Y is given by (2.7); (final good firms)

(E3) $m(s) \in \operatorname{argmax}_{x \in X} \gamma(s, x) p_x$ for all s;

(E4) $w_s = \gamma(s, m(s)) p_{m(s)}$ for all s;

(E5) $p_c = r = \alpha_x p_x$ for all $x < \widetilde{x}$;

(E6) $\frac{w_{\underline{s}}}{\gamma(\underline{s}, \widetilde{x})} = \frac{r}{\alpha(\widetilde{x})}$.

(E3)–(E6): (task producers)

Condition (E1) establishes that the markets for tasks, capital, and labor clear. It derives the amount of labor used in a given task x (the marginal density of labor at x) via a change of variable from the exogenous supply of skills L_s (the marginal density of labor at s), using the assignment of skills to tasks $m(s)$ and labor market clearing. (E2) follows from final good firms' profit maximization. Task producers' profit maximization is reflected in the remaining conditions: each skill is assigned to the task where its marginal product is greatest (E3); this marginal product determines the wage (E4); capital is assigned where its marginal product is greatest and this marginal product determines the price of capital, which in turn must be equal to capital's marginal cost (E5); and the threshold task $\widetilde{x}$ is determined such that task producers are indifferent between using capital and skill $\underline{s}$ in this task (E6).

An immediate consequence of task producers' profit maximization is that relative wages are fully determined by the matching function. In particular, applying the envelope theorem to conditions (E3) and (E4) yields[19]

$$\frac{\mathrm{d}\log w_s}{\mathrm{d}\, s} = \frac{\partial \log \gamma(s, m(s))}{\partial s}. \tag{2.8}$$

As a final remark, the marginal cost of capital must respect a lower bound to guarantee equilibrium existence:

$$r > \beta \left(\int_{\underline{x}}^{\overline{x}} \alpha(x)^{\epsilon-1} \, \mathrm{d}\, x \right)^{\frac{1}{\epsilon-1}}.$$

This is because final good and task production are linear in capital while capital production is linear in final good. Such linearity in circular production may enable infinite output, analogously to unbounded growth of the AK-type in a dynamic model, if the marginal cost of capital is too low.

2.5.2 Automation

The threshold task $\widetilde{x}$ indicates the size of the set of tasks performed by capital, and hence measures the extent of automation in the model. To analyze how automation, as measured by the threshold task $\widetilde{x}$, responds to changes in the supply of skills, apply the concepts of exogenous and endogenous technology equilibrium introduced in Section 2.2. Thereby, $\widetilde{x}$ takes the role of the technology variable θ in the general analysis above. Thus, in an exogenous technology equilibrium $\widetilde{x}$ is fixed exogenously while capital and labor sort endogenously into the tasks below (in the case of capital) and above (in the case of labor) $\widetilde{x}$. This sorting is determined by conditions (E1) to (E5), while condition (E6), which determines $\widetilde{x}$, is dropped. An endogenous technology equilibrium in contrast corresponds exactly to the equilibrium definition above, characterized by the full set of conditions (E1) to (E6).

Given these refined equilibrium definitions, the following lemma verifies that the model fits into the class of models covered by the general results of the previous sections.

[19] Conditional on the threshold task $\widetilde{x}$, the assignment of labor is analogous to Costinot and Vogel (2010), who consider the same model but without capital. Hence the determination of relative wages, conditional on $\widetilde{x}$, is analogous to their analysis as well. In the proof of their Lemma 2, they prove differentiability of the wage function w_s.

Lemma 2. *For any $\widetilde{x} \in (\underline{x}, \overline{x})$ there exists a unique exogenous technology equilibrium. Let $F(L, \widetilde{x})$ denote aggregate net production, that is, $Y - rK$, and $w(L, \widetilde{x})$ denote wages in this equilibrium. Then:*

1. *$F(L, \widetilde{x})$ is linear homogeneous in L.*
2. *Wages correspond to marginal products in F, that is, $w(L, \widetilde{x}) = \nabla_L F(L, \widetilde{x})$.*
3. *Any increase in $\widetilde{x}$ raises all skill premia, that is, $\widetilde{x} \leq \widetilde{x}' \Leftrightarrow \widetilde{x} \preceq^b \widetilde{x}'$ according to Definition 2 of skill-biased technical change.*

Moreover, for any labor supply L there exists a unique endogenous technology equilibrium with automation threshold $\widetilde{x}^*(L)$ such that

4. *$\widetilde{x}^*(L) \in \operatorname{argmax}_{\widetilde{x} \in X} F(L, \widetilde{x})$.*

Proof. See Appendix A.1.1. □

Most of the points of the lemma are straightforward up to some technical details. The economically most relevant result is that automation, represented by an increase in $\widetilde{x}$, raises all skill premia and therefore constitutes a skill-biased technical change. This is intuitive: since capital performs the least complex tasks in the economy, any expansion in the set of automated tasks directly displaces low-skilled workers from their tasks. In search for new tasks, low-skilled workers turn towards more complex tasks, propagating the effects through the skill distribution. But since all workers eventually end up at more complex tasks (where more skilled workers have comparative advantage), skill premia must rise throughout the wage distribution.

Induced Technical Change Effect Consider now an increase in relative skill supply as in Definition 1, that is, an increase in skill supply ratios along the entire skill set. Theorem 1 implies that such an increase in relative skill supply induces skill-biased technical change, or, in the present context, automation.

Corollary 3. *Any increase in relative skill supply induces automation, which itself raises all skill premia, that is,*

$$L \preceq^s L' \Rightarrow \widetilde{x}^*(L) \leq \widetilde{x}^*(L') \Rightarrow w(L', \widetilde{x}^*(L)) \preceq^p w(L', \widetilde{x}^*(L')).$$

Proof. Lemma 2 establishes that all conditions of Theorem 1 are satisfied here, so Corollary 3 follows directly from Theorem 1, given that an increase in $\widetilde{x}$ corresponds to skill-biased technical change in the current model, $\widetilde{x} \leq \widetilde{x}' \Leftrightarrow \widetilde{x} \preceq^b \widetilde{x}'$ (point 3 in Lemma 2). □

Since machines and workers are perfect substitutes in the production of tasks, the interaction between labor supply and automation runs via task prices. In particular, an increase in relative skill supply raises the prices of tasks performed by low-skilled workers, which makes it more attractive for firms to automate these tasks. The more general force behind Corollary 3, as described in Section 2.3.1, is that the production sector responds to the decrease in the relative supply of less skilled workers by switching to technologies that are less reliant on low-skilled labor. Here, low-skilled labor is less important the more tasks are automated, as indicated by the positive effect of automation on the returns to skill. So, firms automate additional tasks in order to minimize adverse effects from decreased (relative) availability of low-skilled workers.

Total Effect The total effect of an increase in relative skill supply on relative wages combines the direct effect (at constant $\widetilde{x}$) and the effect of the induced automation. By Theorem 3, whether the direct of the induced technical change effect dominates depends crucially on the curvature of the isoquants of aggregate net production in the endogenous technology equilibrium, $\overline{F}(L) := F(L, \widetilde{x}^*(L))$.

Here, firms make their automation decisions individually and independently of each other, as in the baseline model of endogenous technology choices in Section 2.2. Therefore, aggregate production is quasiconcave in labor supply and by Theorem 3 the strong bias phenomenon, whereby all skill premia rise with relative skill supply, cannot occur. In addition, a direct consequence of the fact that any increase in relative skill supply induces automation (Corollary 3) while capital productivity remains unchanged is that low-skilled workers' wages can never fall in absolute terms in response to an increase in relative skill supply.

Lemma 3. *The endogenous technology net aggregate production function,* $\overline{F}(L) := F(L, \widetilde{x}^*(L))$ *(where F is as defined in Lemma 2), is quasiconcave.*

Proof. See Appendix A.1.1. □

Corollary 4. *There is no increase in relative skill supply that raises all skill premia after adjustment of the degree of automation $\widetilde{x}^*$. That is,*

$$L \preceq^s L' \Rightarrow w(L, \widetilde{x}^*(L)) \npreceq^b w(L', \widetilde{x}^*(L')).$$

Moreover, any increase in relative skill supply raises the least skilled worker's wage, that is,

$$L \preceq^s L' \Rightarrow w_{\underline{s}}(L, \widetilde{x}^*(L)) \leq w_{\underline{s}}(L', \widetilde{x}^*(L')).$$

Proof. The first part follows directly from Lemma 3 and Theorem 3. The second part follows from Corollary 3 and the fact that

$$w_{\underline{s}}(L, \widetilde{x}^*(L)) = \frac{\gamma(\underline{s}, \widetilde{x}^*(L))}{\alpha(\widetilde{x}^*(L))} r$$

by (E6), noting that $\gamma(s,x)/\alpha(x)$ increases in x for any s by comparative advantage (Assumption 2). □

The result is intuitive: first, the fact that automation is induced by an increase in the prices of tasks performed by low-skilled workers implies that the induced automation can never fully offset this increase in task prices; for if it did, automation would not occur in the first place. Moreover, since task production is linear, the increase in task prices is fully passed through to low-skilled workers' wages (the second part of Corollary 4). If now skill premia increased as well, all wages and hence all workers' marginal products would go up. But then individual firms could choose a greater relative skill input and a higher automation threshold already in the initial equilibrium, and thereby raise profits. Since this cannot be true by definition of equilibrium, the case where all skill premia rise with relative skill supply cannot occur (the second part of Corollary 4).

This reasoning points towards the general force behind Corollary 4, as described in Section 2.3.2: in settings where firms choose their technologies individually and independently of each other, if skill premia rose in relative skill supply, firms would demand more skilled workers in the initial equilibrium already and adjust their technology accordingly. An important reason why the present model of automation falls into this class of settings is that it describes a process of pure technology adoption: given the productivity of machines, firms decide for each task whether to use machines or not. The next section shows that once agents can invest into improving the productivity of machines, quasiconcavity of aggregate net production may fail and strong bias results can arise.

2.5.3 Automation and Machine Productivity

The decision whether to use machines or labor in a given set of tasks is clearly preceded by the decision (potentially by a different set of agents) to invest into developing machines with a certain set of abilities. A natural way to include such a decision in the model is to give agents the opportunity to invest into improving capital productivity $\alpha(x)$. When investing into $\alpha(x)$ is the only opportunity for agents to spend resources on research and development, the investment into $\alpha(x)$ is obviously equal to total R&D expenditure. A probably more realistic approach is to give agents the choice between different types of technologies in which to invest, thereby separating the factors that affect the direction of R&D spending from those that affect its overall amount. In the following, agents will therefore face the choice whether to invest resources into improving machine productivity α or final good productivity β.[20]

To endogenize α and β, the market structure must be adjusted, as both final good and task production exhibit increasing returns to scale in input factors and technology variables jointly. Following the monopolistic competition approach from endogenous growth theory, I therefore assume that α and β are aggregates of monopolistically supplied intermediate goods (see Appendix A.2.1 for a general version of monopolistic competition based models of directed technical change). The monopolistic suppliers then invest R&D resources to improve their products.

In particular, final good production now takes the form

$$Y = \int_0^1 \beta_i q_{\beta,i}^{\kappa} \, \mathrm{d}\, i \left(\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \, \mathrm{d}\, x \right)^{\frac{(1-\kappa)\epsilon}{\epsilon-1}}, \tag{2.9}$$

where the $q_{\beta,i}$ are technology-embodying intermediate goods to be described further below. Tasks are produced according to

$$Y_x = \int_0^1 \alpha_i q_{\alpha,i}^{\kappa} \, \mathrm{d}\, i \, \overline{\alpha}(x) K_x^{1-\kappa} + \int_{\underline{s}}^{\overline{s}} \gamma(s,x) L_{s,x} \, \mathrm{d}\, s,$$

where, again, the $q_{\alpha,i}$ are technology-embodying intermediate goods, which are required to produce tasks using machines. Assumption 2 about comparative advantage is maintained, now applying to $\overline{a}(x)$ and $\gamma(s,x)$. To reduce notation, normalize $\overline{\alpha}(x) \equiv 1$. As before, capital is produced at marginal cost r from final good. The markets for final good, capital, and tasks are still perfectly competitive.

The technology-embodying intermediates, in contrast, are supplied under monopolistic competition. In particular, there is a continuum of α-monopolists, indexed by $i \in [0,1]$, who produce $q_{\alpha,i}$ at marginal cost η_α from final good. Analogously, there is a continuum of β-monopolists who produce $q_{\beta,i}$ at marginal cost η_β from final good.[21] The inverse demand for $q_{\alpha,i}$, derived from task producer optimization, is given by

$$p_{\alpha,i} = \kappa \alpha_i q_{\alpha,i}^{\kappa-1} \int_0^{\widetilde{x}} p_x K_x^{1-\kappa} \, \mathrm{d}\, x, \tag{2.10}$$

[20]Increases in β may be thought of as a stylized description of the invention of new goods or higher quality versions of existing goods, which generate additional utility for consumers. Increases in α in contrast are process innovations that allow to produce a given set of goods with fewer inputs.

[21]In a slight abuse of notation, I use the same index to denote α- and β-monopolists. This shall not implicate that a given monopolist produces both $q_{\alpha,i}$ and $q_{\beta,i}$, although this would not change any argument.

which makes use of the result from Lemma 1 (which carries over to the present setting) that capital is used in a subset of tasks $[\underline{x}, \widetilde{x})$. Analogously, the inverse demand for $q_{\beta,i}$ is

$$p_{\beta,i} = \kappa \beta_i q_{\beta,i}^{\kappa-1} \left(\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \,\mathrm{d}\, x \right)^{\frac{(1-\kappa)\epsilon}{\epsilon-1}} . \tag{2.11}$$

Since both inverse demand functions are iso-elastic, monopolists will charge a constant markup over marginal cost. Intermediate good prices will thus be given by $p_{\alpha,i} = \eta_\alpha / \kappa$ and $p_{\beta,i} = \eta_\beta / \kappa$ in equilibrium.

In addition to supplying intermediate goods, α- and β-monopolists can also invest into the quality of their products. For a quality level α_i, an α-monopolist must employ R&D resources of $\alpha_i^{1/\rho}$, with $\rho \in (0, 1-\kappa)$. Analogously, a β-monopolist must employ $\beta_i^{1/\rho}$ units of R&D resources to obtain a quality level β_i. In order to isolate effects on the direction of technical change from effects on the aggregate amount of resources spent on R&D activities, fix the total amount of R&D resources at D.[22] This implies an R&D resource constraint of $\int_0^1 \left(\alpha_i^{1/\rho} + \beta_i^{1/\rho} \right) \mathrm{d}\, i = D$.

Denote the unit price of R&D resources by p_D. Each α-monopolist then chooses α_i to maximize profits

$$\pi_{\alpha,i}(\alpha_i) = \max_q \left\{ \kappa \alpha_i q^\kappa \int_{\underline{x}}^{\widetilde{x}} p_x K_x^{1-\kappa} \,\mathrm{d}\, x - \eta_\alpha q - p_D \alpha_i^{1/\rho} \right\} .$$

With $\rho \in (0, 1-\kappa)$, it can be verified that profits are pseudoconcave in α_i, so the first order condition for the choice of α_i is necessary and sufficient for an optimum:

$$\rho \kappa q_{\alpha,i}^\kappa \int_{\underline{x}}^{\widetilde{x}} p_x K_x^{1-\kappa} \,\mathrm{d}\, x = p_D \alpha_i^{\frac{1-\rho}{\rho}} . \tag{2.12}$$

Analogously, β-monopolists' profit maximization leads to the following first order condition for the choice of β_i:

$$\rho \kappa q_{\beta,i}^\kappa \left(\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \,\mathrm{d}\, x \right)^{\frac{(1-\kappa)\epsilon}{\epsilon-1}} = p_D \beta_i^{\frac{1-\rho}{\rho}} . \tag{2.13}$$

With this characterization of technology choices, the equilibrium definition from the previous section can be extended appropriately. Since all α-monopolists and all β-monopolists choose the same $q_{\alpha,i}$ and α_i, or, respectively, the same $q_{\beta,i}$ and β_i, it is convenient to define an equilibrium in terms of their symmetric choices q_α, α, q_β, and β.

An equilibrium consists of

[22] This is equivalent to the assumption of a fixed amount of "research labor" often made in dynamic models with endogenously directed technical change; see, for example, Acemoglu and Restrepo (2018a).

- an automation threshold $\widetilde{x}$, a matching function $m : S \to [\widetilde{x}, \overline{x}]$, an assignment of capital to tasks $\{K_x\}_{x \in X}$, task output $\{Y_x\}_{x \in X}$, technology intermediate quantities q_α and q_β, and productivity levels α and β;
- task prices $\{p_x\}_{x \in X}$, wages $\{w_s\}_{s \in S}$, a capital price p_c, technology intermediate prices p_α and p_β, and a price of R&D resources p_D;

such that

(E1)' $Y_x = \alpha q_\alpha^\kappa K_x^{1-\kappa}$ if $x < \widetilde{x}$ and $Y_x = \gamma(m^{-1}(x), x) L_{m^{-1}(x)} \frac{\mathrm{d}\, m^{-1}(x)}{\mathrm{d}\, x}$ if $x \geq \widetilde{x}$; (market clearing)

(E2)' $p_x = \frac{\partial Y}{\partial Y_x}$ for all x, where Y is given by (2.9);

(E3)' q_β satisfies equation (2.11);

} (final good firms)

(E4)' $m(s) \in \operatorname{argmax}_{x \in X} \gamma(s, x) p_x$ for all s;

(E5)' $w_s = \gamma(s, m(s)) p_{m(s)}$ for all s;

(E6)' $\left(\frac{p_\alpha}{\kappa \alpha}\right)^\kappa \left(\frac{p_c}{(1-\kappa)\alpha}\right)^{1-\kappa} = p_x$ for all $x < \widetilde{x}$ and $p_c = r$;

(E7)' $\frac{w_{\underline{s}}}{\gamma(\underline{s}, \widetilde{x})} = \left(\frac{p_\alpha}{\kappa \alpha}\right)^\kappa \left(\frac{r}{(1-\kappa)\alpha}\right)^{1-\kappa}$;

(E8)' q_α satisfies equation (2.10);

} (task producers)

(E9)' $p_\alpha = \frac{\eta_\alpha}{\kappa}$ and $p_\beta = \frac{\eta_\beta}{\kappa}$;

(E10)' α, β, and p_D satisfy equations (2.12), (2.13), and $\alpha^{\frac{1}{\rho}} + \beta^{\frac{1}{\rho}} = D$.

} (technology firms)

Compared to the previous section, conditions (E3)' and (E8)' to (E10)' are new. (E3)' and (E8)' determine the quantities of technology-embodying intermediate inputs as demanded by task producers or final good firms. (E9)' and (E10)' determine prices and productivity levels of technology intermediates as chosen by the corresponding monopolists. The remaining conditions are either unchanged or slightly adjusted to account for the fact that final good and task production now use the technology-embodying intermediate goods.

The full list of conditions (E1)' to (E10)' again defines an endogenous technology equilibrium, in the sense that the technology variables of interest α and β are determined endogenously. An exogenous technology equilibrium in contrast is characterized by conditions (E1)' to (E9)', given an exogenously fixed pair (α, β).

The following lemma verifies that the extended model is still covered by the general results obtained in Section 2.4.

Lemma 4. *For any* (α, β) *such that* $\alpha^{1/\rho} + \beta^{1/\rho} = D$*, there exists a unique exogenous technology equilibrium. Define* $F(L, \alpha, \beta)$ *as a "modified aggregate production function",*

$$F(L, \alpha, \beta) := Y - rK - \frac{\eta_\alpha}{\kappa} q_\alpha - \frac{\eta_\beta}{\kappa},$$

with Y*,* K*,* q_α*, and* q_β *being quantities in the exogenous technology equilibrium, and let* $w(L, \alpha, \beta)$ *denote wages in the exogenous technology equilibrium. Then:*

1. $F(L, \alpha, \beta)$ *is linear homogeneous in* L*.*

2. *Wages equal marginal products, that is,* $w(L, \alpha, \beta) = \nabla_L F(L, \alpha, \beta)$*.*

3. *For any* $(\alpha, \beta), (\alpha', \beta')$ *that satisfy the R&D resource constraint and* $\alpha \leq \alpha'$ *the following holds:*
$$F(L, \alpha, \beta) \leq F(L, \alpha', \beta') \Rightarrow w(L, \alpha, \beta) \preceq^p w(L, \alpha', \beta').$$

Moreover, for any L *and any*

$$(\alpha^*(L), \beta^*(L)) \in \underset{(\alpha,\beta)\in\mathcal{D}}{\operatorname{argmax}} F(L, \alpha, \beta),$$

where $\mathcal{D} = \left\{(\alpha, \beta) \in \mathbb{R}^2_+ \mid \alpha^{1/\rho} + \beta^{1/\rho} = D\right\}$ *is the innovation possibilities frontier, there exists an endogenous technology equilibrium with equilibrium productivity levels* $\alpha^*(L)$ *and* $\beta^*(L)$*.*

Proof. See Appendix A.1.1. The proof also shows that the endogenous technology equilibrium is unique whenever the innovation possibilities frontier is "sufficiently convex", as indicated by a sufficiently small ρ. Whenever the endogenous technology equilibrium is not unique, I select the equilibrium with the highest α^* in the following, in line with the selection rule imposed in Section 2.2. □

Note that $F(L, \alpha, \beta)$ here does not exactly correspond to net aggregate production in the model. Indeed, in the definition of F, the marginal costs of technology intermediates, η_α and η_β, are replaced by the intermediates' prices, η_α/κ and η_β/κ. The idea behind this change is that, when marginal costs are replaced in such a way, technology intermediates are supplied at the new marginal costs in equilibrium, and hence the exogenous technology equilibrium can be analyzed as if it were generated by perfect competition on all markets. This gives rise to the equality of wages and marginal products of labor in point 2 of the lemma (see also the analysis of the general version of monopolistic competition based models of endogenous technical change in Appendix A.2.1).

The second notable point of Lemma 4 is point 3: an increase in α, which here necessarily comes at the cost of a reduced β, raises skill premia whenever it raises F. Generally, an increase in α has two effects. On the one hand, machines become more productive, displace low-skilled workers, and therefore raise skill premia. On the other hand, the corresponding decrease in β reduces wages, while the price of capital remains constant. This stifles automation and hence reduces skill premia. Lemma 4 then shows that the former effect dominates when F increases. Hence, an increase in capital productivity α that raises "modified aggregate production" F is a skill-biased technical change.[23]

Induced Technical Change Effect Consider now an increase in relative skill supply. Lemma 4 establishes that all conditions for Theorem 2 are satisfied, so the theorem immediately implies the following result.

Corollary 5. *Any increase in relative skill supply induces an improvement in capital productivity, which itself raises all skill premia, that is,*

$$L \preceq^s L' \Rightarrow \alpha^*(L) \leq \alpha^*(L') \Rightarrow w(L', \alpha^*(L), \beta^*(L)) \preceq^p w(L', \alpha^*(L'), \beta^*(L')).$$

Proof. Lemma 4 establishes that all conditions of Theorem 2 are satisfied. Theorem 2 then immediately implies the corollary. □

The result is closely related to Corollary 3 of the previous section. Corollary 3 says that an increase in relative skill supply induces automation as it raises the prices of those tasks that are technologically most prone to automation. The increase in automation in turn raises the incentive to improve the productivity of machines, as they become more widely used.

Total Effect Indeed, automation and improvements in machine productivity reinforce each other: the more widely machines are used, the greater is the incentive to improve them; and the more productive machines are, the more widely they are used. This reinforcement mechanism tends to "convexify" aggregate production and may thus, following Theorem 3, generate strong bias results.

To see this concretely, consider the limit case where there is a subset of low-skilled workers $[\underline{s}, \widetilde{s}]$ who have no discernible comparative advantage over machines. Formally, for $s \in [\underline{s}, \widetilde{s}]$, $\gamma(s, x)$ is constant in x and hence proportional to $\overline{\alpha}(x)$ (recall the normalization $\overline{\alpha}(x) \equiv 1$). This case itself does not satisfy Assumption 2, but it is the limit of a sequence of cases all covered

[23] More precisely, any increase in α along the innovation possibilities frontier is a skill-biased technical change according to the alternative Definition 3.

by the assumption. Since the equilibrium is continuous in the relevant parameters, we can still analyze the limit case on the basis of Lemma 4 and Corollary 5. The only complication is that the absence of strict comparative advantage between capital and workers with skill below $\widetilde{s}$ means that the assignment of these factors to tasks is no longer uniquely determined. This indeterminacy, however, neither affects prices, nor task, nor final good quantities. So we can safely ignore it when analyzing the response of wages to changes in labor supply conditional on curvature properties of aggregate production.

It can now be verified that aggregate net production is not quasiconcave in the limit case when taking into account the endogenous adjustment of machine productivity.

Lemma 5. *The endogenous technology production function $\overline{F}(L) := F(L, \alpha^*(L), \beta^*(L))$, where F is defined as in Lemma 4, fails to be quasiconcave along some line in direction of $\preceq^s$.*

Proof. See Appendix A.1.1. □

According to Theorem 3, the failure of quasiconcavity established in Lemma 5 generates the potential for strong bias. That is, skill premia may rise in response to an increase in relative skill supply. As an example, consider a proportional increase in the supply of all skill levels above $\widetilde{s}$ by a factor of $\lambda > 1$. Holding machine productivity constant at its initial level, it is easy to see that all wages remain unchanged. In particular, let $\widetilde{x}'$ be the threshold task such that skills above (below) $\widetilde{s}$ sort into tasks above (below) $\widetilde{x}'$ before the labor supply change, and suppose that the task assignment for skills above $\widetilde{s}$ remains unchanged when labor supply changes. Then, capital adjusts in a way that raises all task quantities below $\widetilde{x}'$ by the factor λ. This holds all task ratios and hence all task prices unchanged, such that, given constant labor assignment, wages will be unchanged as well. Constancy of wages in turn confirms the initial assumption of an unchanged labor assignment. So, at constant machine productivity, skill premia do not change in response to the specific increase in relative skill supply described above. But by Corollary 5, machine productivity will increase in response to the increase in relative skill supply. This raises all skill premia above their initial level, because the increase in machine productivity constitutes a skill-biased technical change (by Lemma 4).

In addition, low-skilled workers' wages will fall in response to any increase in relative skill supply. This is because capital is a perfect substitute for low-skilled workers in all tasks, due to the absence of comparative advantage between these factors. Therefore, when the productivity of machines rises while their prices stay constant, low-skilled workers' wages must fall.

Corollary 6. *Consider the limit case where $\gamma(s,x)/\overline{\alpha}(x)$ is constant in x for all $s \leq \widetilde{s}$ for some $\widetilde{s} \in (\underline{s}, \overline{s})$. Then, skill premia may rise in relative skill supply. Consider for example an increase in* *relative skill supply from L to L' such that $L'_s = \lambda_1 L_s$ for all $s \leq \widetilde{s}$ and $L'_s = \lambda_2 L_s$ for all $s > \widetilde{s}$ with $\lambda_2 > \lambda_1$. This increase in relative skill supply raises all skill premia,*

$$w(L, \alpha^*(L), \beta^*(L)) \preceq^p w(L', \alpha^*(L'), \beta^*(L')).$$

Moreover, in the limit case low-skilled workers' wages fall in response to any increase in relative skill supply,

$$L \preceq^s L' \Rightarrow w_s(L, \alpha^*(L), \beta^*(L)) \geq w_s(L', \alpha * (L'), \beta^*(L'))$$

for all $s \leq \widetilde{s}$.

Proof. The first part follows from Lemma 5 and Theorem 3, the example is proven in the text for $\lambda_1 = 1$. It holds for arbitrary $\lambda_1 < \lambda_2$ by zero homogeneity of wages and technology in L. The second part follows from the fact that, by the equilibrium condition (E7)',

$$w_{\underline{s}} = \left(\frac{p_\alpha}{\kappa\alpha}\right)^\kappa \left(\frac{r}{(1-\kappa)\alpha}\right)^{1-\kappa} \gamma(\underline{s}, \widetilde{x}), \tag{2.14}$$

observing that $\gamma(\underline{s}, x)$ is constant in x in the limit case under consideration and that α increases in response to any increase in relative skill supply by Corollary 5. The result extends to all skills $s \leq \widetilde{s}$ by noting that the ratios $w_s/w_{\underline{s}}$ are fixed for all $s \leq \widetilde{s}$, due to the absence of strict comparative advantage between these skills in the limit case. □

The central mechanism behind the results of Corollary 6 is the reinforcement between automation and investment in machine productivity. At fixed machine productivity, the automation induced by an increase in relative skill supply never outweighs the direct effect of the increase in relative skill supply on the skill premium (see Section 2.5.2). But automation raises the incentives to improve the productivity of machines, which in turn reinforces automation in a way that may ultimately overturn the direct effect effect and lead the skill premium to increase in total.

The general force behind this result is the failure of quasiconcavity in aggregate production, which is enabled by the separation of technology and labor demand choices in the model (see Section 2.3.2). In the case of strong bias, individual task producers would like to increase skilled labor input, automation, and machine productivity jointly, as this would raise their profits (see the discussion after Corollary 4). But machine productivity is chosen by technology firms, and technology firms do not cater to an individual firm's demand but to the aggregate demand of all task producers. Aggregate technology demand of task producers, however, depends on aggregate labor input, and aggregate labor input is restricted by labor supply. Technology firms therefore choose machine productivity taking aggregate labor input as given, while task producers demand labor taking the available technology as given. Hence, even though all individual firms' objectives are concave, the aggregate production function may fail to be concave in labor and technology jointly.

While Corollary 6 is restricted to the limit case, continuity arguments imply that its results hold more broadly. In particular, strong bias and the drop in low-skilled workers' wages are generally likely whenever there are no tasks in which low-skilled workers maintain a strong comparative advantage over machines. But even if such tasks exist, it is not clear that they are of great help to the low-skilled. First, they may already be occupied by more skilled workers with a comparative advantage over low-skilled workers in these tasks. Second, their number may be small relative to the number of displaced workers, making their prices fall rapidly as low-skilled workers relocate. Finally, in reality, though not in the present model, limits to (for example, spatial) mobility may prevent low-skilled workers from accessing such tasks.

Discussion In summary, this section demonstrates that not only the use of automation technology but also its development responds to increases in the relative supply of skill in a way that is detrimental to low-skilled workers. An increase in relative skill supply induces automation, which in turn stimulates investment into improving the underlying technologies. Such improvements then further increase the incentives to automate tasks. In effect, low-skilled workers' wages may decline in total, both relative to more skilled workers' wages and in absolute terms, when the relative supply of skilled workers rises. This has potentially important implications for a rich set of policies that affect labor supply differentially at different points of the skill distribution. Minimum wages, for example, may reduce employment among low-skilled workers and thereby both provide incentives to replace such workers by machines and stimulate investment into improving these machines. When the technological feasibility of automation increases, these effects may create or exacerbate adverse employment effects of minimum wages. This is roughly in line with the results of Lordan and Neumark (2018) who find that minimum wage increases in the US over the last decades, while not having large effects on overall employment, have significantly reduced employment in occupations that are particularly vulnerable to automation in terms of their task mix. As another example, tax and benefit systems in many European countries impose particularly high marginal tax rates on low incomes (cf. OECD, 2011). This arguably restricts the labor supply of low-skilled workers and hence may intensify automation along the lines analyzed above. Such effects should clearly be taken into account when designing tax and transfer systems. A detailed analysis of these issues is left for future research, as well as the pursuit of empirical approaches to test the derived hypotheses.

2.5.4 Automation, Machine Productivity, and International Trade

The previous sections have analyzed how the use and development of automation technology depends on the supply of skills in the economy. The measure of skill supply in these analyses should clearly capture the entire pool of workers whose performance is accessible to firms via any type of (competitive) market. In a globalized world, however, firms do not only have access to domestic workers via the labor market, but also to the performance of foreign workers via international trade in tasks or, more broadly, intermediate goods. Therefore, the conditions under which countries trade with each other should have important effects on technologies used in general and on automation in particular. This section thus extends the model of the previous section to a two country setting and analyzes the interaction between trade and automation.

To this end, consider two countries called North and South. Under autarky, the Northern economy is described by the model of the previous section, where both the extent of automation and the productivity of machines are endogenous. The Southern economy differs from the North in exactly three aspects. First, it has no research sector but copies the technologies developed in the North with some loss in productivity. In particular, let α^N and β^N denote productivity levels in the North. Then, intermediate good firms in the South can produce goods of quality $\delta\alpha^N$ and $\delta\beta^N$ without incurring R&D costs, where $\delta \in (0, 1)$ measures the productivity loss relative to the North. In the absence of R&D costs, intermediate goods are supplied competitively and hence priced at marginal costs η_α and η_β in equilibrium. It follows that the Southern economy uses less advanced technologies than the North but does not feature R&D-related monopoly distortions. Let $\delta < \kappa$, such that the quality-adjusted price of intermediate goods in the South is greater than in the North, and the aggregate production process in the South is less efficient. The second difference between the two countries is that the South is skill-scarce relative to the North, that is, $L^S \preceq^s L^N$ with L^S and L^N denoting labor supply in the South and the North, respectively. Finally, I follow Costinot and Vogel (2010) and assume that labor productivity is lower in the South than in the North, $\gamma^S(s, x) = \Delta\gamma^N(s, x)$, with $\Delta \in (0, 1]$. While irrelevant for all results discussed below, this assumption allows for differences in the wage levels conditional on skill between North and South, even when tasks can be traded across countries.

An autarky equilibrium is defined as the union of: (i) an endogenous technology equilibrium as by conditions (E1)' to (E10)' for the North, and (ii) an exogenous technology equilibrium, characterized by conditions (E1)' to (E8)' plus the price condition $p^S_\alpha = \eta_\alpha$ and $p^S_\beta = \eta_\beta$ (replacing condition E9' due to the absence of monopoly distortions), for the South, with Southern technology (α^S, β^S) given by $(\delta_\alpha\alpha^N, \delta_\beta\beta^N)$.

Autarky is contrasted with a situation where all types of goods, that is, tasks, technology-embodying intermediates, and the final good, can be traded between the two countries.[24] In such a situation, Northern technology monopolists will serve the entire world market, since they produce output of higher quality at the same marginal cost as Southern technology firms. It follows that task and final good producers use the same technology in both countries, with the exception that Southern labor productivity is reduced by the factor Δ across all tasks. Under these conditions world production of the different types of goods and world prices will be the same as in a hypothetical scenario of full integration of both countries where Southern labor, scaled down by the productivity handicap Δ, moves to the North. This full integration scenario in turn is identical to an autarky equilibrium in the North with labor supply given by $L^N + \Delta L^S$ instead of L^N. We can hence equate the effects of trade integration on capital productivity α^N (the world technology frontier) and on Northern wages w^N with the effects of a change in labor supply from L^N to $L^N + \Delta L^S$. Appendix A.1.2 derives this equality formally, constructing equilibrium conditions for world quantities and prices under trade integration that are equivalent to conditions (E1)' to (E10)' from the closed economy setting. The only formal difference between the conditions for world quantities and prices and the conditions for an autarky equilibrium in the North is then that the former use world labor supply $L^N + \Delta L^S$ where the latter use Northern labor supply L^N only.

Induced Technical Change Effect Given that the effects of trade integration on capital productivity α^N are identical to the effects of increasing labor supply by ΔL^S, Corollary 5 implies the following result.

Corollary 7. *Trade integration with the South induces an improvement in the productivity of final good production β^N at the expense of reduced capital productivity α^N in the North; that is, $\alpha^{NT} \leq \alpha^N$, where α^{NT} denotes capital productivity under trade integration and α^N Northern capital productivity under autarky.*

Proof. It is easy to verify that $L^S \preceq^s L^N$ implies $L^N + \Delta L^S \preceq^s L^N$. Corollary 7 then follows as a consequence of Corollary 5, given that the effects of trade integration are equal to the effects of changing labor supply from L^N to $L^N + \Delta L^S$. This equality is derived formally in Appendix A.1.2. $\square$

To understand Corollary 7 on an intuitive level, note that the North imports tasks performed by low-skilled workers from the South, because the South is abundant in low-skilled labor. In exchange, the North exports technology-embodying intermediates, final good, and, potentially,

[24] The results are robust to alternative assumptions about which types of goods are tradable and which are not. See the discussion in footnote 26 below.

tasks performed by high-skilled workers.[25] The low-skill-intensive imports from the South reduce the prices of tasks performed by low-skilled workers in the North. This reduces the wages of Northern low-skilled workers, while the cost of capital remains constant. It follows that the incentive to automate tasks performed by low-skilled workers decreases. The thus induced reduction in the use of automation technology in turn also reduces investment into improving these technologies, hence α^N falls.

Total Effect The reduction in the use of automation technology and the decline in investment into its improvement reinforce each other. By the arguments provided in the preceding section, this reinforcement may lead to an overall increase in low-skilled workers' wages from trade integration in the North, both relative to high-skilled workers' wages and in absolute terms. Again, this is particularly likely if low-skilled workers and machines are highly substitutable, that is, if machines have no strong comparative advantage in the tasks they would perform in autarky (see Section 2.5.3). In particular, Corollary 6 implies the following results for the effects of trade in the North.

Corollary 8. *Consider the limit case where $\gamma^N(s,x)/\overline{\alpha}(x)$ is constant in x for all $s \leq \widetilde{s}$ for some $\widetilde{s} \in (\underline{s}, \overline{s})$. Then, skill premia may fall in the North in response to trade integration with the South. Consider for example a situation where $L_s^S = \lambda_1 L_s^N$ for all $s \leq \widetilde{s}$ and $L_s^S = \lambda_2 L_s^N$ for all $s > \widetilde{s}$ with $\lambda_1 > \lambda_2$. In this situation, trade integration reduces all skill premia in the North,*

$$w^{NT} \preceq^p w^N,$$

where w^{NT} denotes Northern wages under trade integration and w^N under autarky.

Moreover, in the limit case Northern low-skilled workers' wages rise in response to trade integration for any $L^S \preceq^s L^N$,

$$w_s^{NT} \geq w_s^N$$

for all $s \leq \widetilde{s}$.

Proof. Given that the effects of trade integration with the South are equivalent to the effects of a change in skill supply from L^N to $L^N + \Delta L^S$ (shown formally in Appendix A.1.2), Corollary 8 follows directly from its closed economy counterpart, Corollary 6. □

[25] Indeed, there is some degree of indeterminacy regarding the trade of final goods and technology intermediates in equilibrium. The South can either import technology goods from the North and produce the final good itself, or import the final good directly. To which extent the South makes use of either option is unclear. There is a continuum of possible outcomes, with two polar cases: first, the South imports all its final goods but no technology intermediates; second, it produces all its final good consumption itself, importing technology intermediates for that purpose. The indeterminacy, however, only affects the division of final good production between North and South. The overall production of goods in the world is unaffected, as are prices and wages.

Intuitively, trade with the South has two opposing effects on Northern low-skilled workers. First, they are exposed to import competition from the South as tasks produced by low-skilled workers are cheap in the South due to its abundance in low-skilled labor. This is the standard Heckscher-Ohlin effect, which puts downward pressure on low-skilled workers' wages in the North. Second, the reduction in automation, reinforced by the decline in machine productivity, expands employment opportunities for low-skilled workers and hence raises their wages. This effect is especially strong when the productivity profiles of low-skilled workers and machines are similar, such that low-skilled workers can benefit a lot from the retreat of machines. In this case, the automation effect dominates, such that the wages of Northern low-skilled workers rise, in relative and absolute terms, in response to trade integration, contrary to the standard Heckscher-Ohlin prediction.

The effect of trade integration on the Southern wage distribution is twofold as well. First, the standard Heckscher-Ohlin effect reduces skill premia, because the South is skill-scarce relative to the North. Second, the advanced Northern technology becomes available to the South, either directly via trade in technology-embodying intermediates or indirectly via trade in tasks. This exposes Southern low-skilled workers to competition from advanced Northern machines, and hence raises skill premia. The technology effect is likely to dominate when (i) the productivity difference between North and South is large under autarky, that is, δ is large; (ii) the reduction in Northern investment into automation technology induced by trade integration is small; and (iii) Heckscher-Ohlin effects are weak, for example, because the supply of skills is similar in the North and the South.

It is indeed straightforward to prove that trade integration can reduce low-skilled workers' wages in the South by constructing an extreme example. Suppose that skill supply is nearly identical between North and South, and that $\delta \ll \kappa$ such that the South uses much less advanced technology than the North in autarky. Then, Heckscher-Ohlin effects and the effect of trade integration on Northern capital productivity will be negligible, as there is hardly any difference in the relative supply of skills between the Northern and the world economy. The effect on Southern capital productivity, however, will be large, because trade makes the much more advanced Northern technology accessible to Southern firms. This effective increase in capital productivity will reduce low-skilled workers' wages in the South if the comparative advantage of low-skilled workers over capital is weak across tasks, as explained in more detail in Section 2.5.3.

Discussion In summary, trade with a skill-scarce country discourages both the use and the development of automation technology in a skill-abundant country. Moreover, if the skill-abundant country is technologically more advanced than the skill-scarce country, trade exposes low-skilled workers in the skill-scarce country to competition from the advanced machines of the skill-abundant country. These effects may overturn the standard Heckscher-Ohlin effects in both countries.[26]

From a theoretical perspective, it is insightful to compare the effects of trade on automation technology with the effects of trade on labor-augmenting technology from Acemoglu (2003). Acemoglu shows that trade with a skill-scarce country, assumed to have no independent R&D sector (as above), induces a skill-biased change in labor-augmenting technology. The reason is that with labor-replacing technology the interaction between labor and technology exclusively works via task prices. With labor-augmenting technology in contrast there is also a quantity-related effect, called the market size effect in Acemoglu (2002, 2003), because technology variables multiply with (instead of add to) labor supply.[27]

From an empirical perspective, the negative effect of trade on automation seems roughly in line with the low correlation between measures of exposure to industrial robots and exposure to Chinese imports across US commuting zones found in Acemoglu and Restrepo (2019).[28] While one might expect both industrial robots and Chinese imports to affect a similar set of industries (manufacturing industries intensive in low-skilled labor) and therefore a similar set of commuting zones, the correlation between the two exposure measures, conditional on a coarse set of covariates is even slightly negative (Acemoglu and Restrepo, 2019, p. 15). On the industry level, the automotive industry experienced by far the largest increase in the number of robots per worker between 1993 and 2007, but hardly any increase in the value of imports from China. The increase in the value of imports from China on the other hand was most pronounced in the textile industry, where the

[26] Note also that the results are robust to alternative assumptions about which types of goods can be traded and which not. Whether the final good is traded in equilibrium, is indeterminate anyway (see footnote 25). The results therefore do not change if the final good cannot be traded. In this case, both countries produce all their final good consumption themselves, and the South relies on technology imports from the North for that purpose. The difference to the baseline case (where all types of goods are traded) is only in the division of final good production between the two countries; world quantities and prices are unchanged. If, instead, technology intermediates cannot be traded, the South imports all its final good consumption from the North, but no technology goods. Again, only the division of final good production is affected, while world quantities and prices are unchanged relative to the baseline scenario. It is, however, crucial for the results that tasks can be traded. Without trade in tasks, wage structures and the extent of automation may differ strongly between the two countries. Technology firms will then cater to a weighted average of the two countries' demands, and it is unclear how this affects R&D investment and automation decisions relative to the autarky case.

[27] The first part of this paper shows that regarding the effects of changes in labor supply on the skill bias of production technology, there is essentially no difference between labor-augmenting and labor-replacing technology. Now I find that the effects of trade are opposite under these two regimes. The reconciliation is that for labor-replacing technology the effects of trade and changes in labor supply are the same, whereas this does not hold for labor-augmenting technology.

[28] These measures are constructed using changes in the number of robots (the value of Chinese imports) in a detailed set of industries between 1993 (1990) and 2007, and weighting these changes by the industries' employment shares at some prior date in each commuting zone.

number of robots per worker did virtually not increase. A loose interpretation of the developed theory would suggest the following explanation: since trade costs are higher for automobile parts than for textiles (due to the higher weight and volume), offshoring low-skill-intensive tasks to China is more attractive in the textile than in the automotive industry. Via the channels discussed above, this reduces the incentive to automate tasks in the textile relative to the automotive industry. Automation technologies such as industrial robots are therefore primarily used (and developed for use) in the automotive, not in the textile industry. A rigorous empirical analysis of these issues, building on a richer set of control variables and appropriate strategies to obtain exogenous identifying variation for the effect of trade on automation, is left for future research.

A further empirical observation that is broadly supportive of the predictions derived above comes from the debate around reshoring, which denotes the relocation of, primarily manufacturing, production from emerging or developing countries to developed economies. Backer, Menon, Desnoyers-James and Moussiegt (2016) report that such reshoring activities are related to increased capital investment but not to significant employment creation in the developed economy to which production relocates. This is in line with the model's prediction that tasks which are produced in the advanced economy instead of being offshored to a skill-scarce country are likely to be automated if they are intensive in low-skilled labor. Automation in response to reshoring would then explain the observation of increased capital investment without employment growth. Even more closely related to the predictions of the model, Krenz et al. (2018) find a positive correlation between reshoring and the use of industrial robots across several manufacturing industries in panel of mainly industrialized countries.

From a policy perspective, the negative relation between automation and trade is relevant for the design of policies regulating the trade between developed and emerging or developing countries. It casts some doubt on policies that aim to protect low-skilled workers in advanced economies by restricting trade with skill-scarce countries. In particular, the theoretical results suggest that such policies may seriously backfire: by stimulating use and development of automation technology, such policies may eventually leave low-skilled workers in the advanced economy no better or even worse off than before. A rigorous theoretical analysis of optimal trade policy when automation responds endogenously may be another promising task for future research (see Costinot, Donaldson, Vogel and Werning, 2015 for an optimal trade policy analysis in assignment models when labor is the only production factor).

The predicted negative effect of Northern automation technology on Southern low-skilled workers is related to recent estimates of the share of employment that is susceptible to automation from a technological point of view in different countries. The World Development Report 2016 (World Bank, 2016) estimates this share to be higher in developing than in developed countries.

The report also notes that barriers to and time lags in the adoption of new technologies are likely to mitigate the impact of automation on developing countries. To the extent that such barriers and time lags are related to trade restrictions, this is in line with the predictions of the theory.

A more concrete manifestation of the impact of Northern automation technology on Southern workers may be the persistent food trade deficit of many African countries that evolved in the mid 1970s (e.g. Rakotoarisoa, Iafrate and Paschali, 2012). While subsidization of food production in advanced economies is often cited as a reason for these deficits, the theory developed here suggests that they might even occur in the absence of policy interventions: since agriculture is highly automated in advanced economies, it may be at a comparative advantage relative to agricultural production in developing countries, which still largely relies on human labor.[29] Agricultural imports can then be expected to hurt the typically poor and uneducated rural population in developing countries. The impact will be particularly severe when opportunities to evade the competition from foreign machines are rare. In the case of agriculture workers, such opportunities may consist of manufacturing jobs, which require workers to migrate from rural to more urbanized areas. Impediments to this form of migration, such as a lack of infrastructure and (affordable) housing space in the urbanized areas, may then create the shortage of alternative employment possibilities emphasized by the theory.

2.6 Conclusion

The first part of the paper develops general results, based on simple concepts, about the effects of the supply of skills on the skill bias of technical change. The results are independent of the functional form of aggregate production, hold for a variety of different microfoundations of endogenous technology choices, for settings with more than two and potentially infinitely many different levels of skill, and apply to both discrete and infinitesimal changes in the supply of skills. They show that under a scale invariance restriction on the skill bias of technology any increase in the relative supply of skills induces skill-biased technical change. Moreover, the total effect of an increase in relative skill supply on skill premia, accounting both for the induced technical change effect and the direct effect, can be positive only if aggregate production fails to be quasiconcave. This generalizes upon existing results, which are limited to the special case of differentially labor-augmenting technology, two skill levels, and infinitesimal changes in the supply of skills.

[29] The expectation that free trade may not necessarily improve the food trade position of developing countries is also implicitly reflected in the series of WTO negotiations on agricultural trade. Both the WTO Agreement on Agriculture from the Uruguay Round of 1995 and the more recent Nairobi Package from 2015 provide comprehensive exemptions to developing countries from requirements to cut import tariffs and export subsidies.

The second part uses the developed theory to derive novel predictions on endogenous automation technology in assignment models of the type proposed by Teulings (1995). In the model investigated, a continuum of differentially skilled workers and capital, taking the form of machines that perfectly substitute for labor in the production of tasks, are assigned to a continuum of tasks, which in turn are combined to produce a single final good. Three results stand out. First, any increase in relative skill supply induces automation, as measured by the set of tasks performed by machines. Second, when machine productivity is endogenous, an increase in relative skill supply does not only stimulate automation but also investment into improving machine productivity. Such investments and automation reinforce each other, potentially leading to a situation where low-skilled workers' wages decrease, both relative to high-skilled workers' wages and in absolute terms, in response to an increase in relative skill supply. Third, in a two country setting the reinforcement mechanism between automation and investment into machine productivity may overturn the standard Heckscher-Ohlin effects from international trade. In particular, trade with a skill-scarce country reduces incentives for the use and development of automation technology in the skill-abundant country, potentially leading to (relative and absolute) increases in low-skilled workers' wages. In the skill-scarce country in contrast, low-skilled workers are exposed to competition from the advanced machines of the skill-abundant country, potentially causing their wages to decline in response to trade.

There are several starting points for future research. First, the results of the first part and the results on the effects of skill supply on automation may serve as a starting point for future explorations of the implications of endogenous technical change in general and endogenous automation in particular for the design of redistributive policies, such as redistributive labor income taxation. The results on the interaction of international trade and automation may as well be the starting point for an analysis of optimal trade policy along the lines of Costinot et al. (2015). Second, the predictions on determinants of the use and development of automation technology from the second part should be of interest for empirical work. Especially the predictions on the effects of trade on automation are testable once a suitable source of exogenous variation across observational units in the exposure to trade is found. Finally, moving beyond the analysis of low-skill automation by relaxing the assumption that machines always have comparative advantage versus workers in less complex tasks seems an important goal for future theory.

3 Redistributive Income Taxation with Directed Technical Change

Author: Jonas Löbbing

3.1 Introduction

Technical change is widely considered an important determinant of changes in the wage structure of an economy and hence of first-order importance for the design of redistributive tax schemes. Existing work analyzes how redistributive taxes respond optimally to exogenous changes in production technology that affect the wage distribution (e.g. Ales, Kurnaz and Sleet, 2015). But technologies are developed and adopted by firms pursuing economic objectives (e.g. Acemoglu, 1998, 2007), so they should respond to perturbations of the economy such as tax reforms. In particular, previous work on directed technical change has theoretically proposed and empirically substantiated that the supply of skills in an economy is an important determinant of the extent to which technology favors skilled workers and thereby raises wage inequality (e.g Acemoglu, 1998; Lewis, 2011; Carneiro et al., 2019). At the same time a large literature on redistributive taxation shows that (non-linear) labor income taxes distort the supply of labor at different levels of skill (e.g. Meghir and Phillips, 2010; Chetty, 2012). Changes in labor income taxes should thus be expected to induce changes in technology, which in turn affect pre-tax wage inequality. Taking into account these technology responses in the analysis of income tax policy seems an important task for taxation theory.[1]

This paper therefore analyzes the design of non-linear labor income taxes when technology is determined endogenously through the profit-maximizing decisions of firms. For that purpose, I develop a general but tractable model of the economy that features both endogenous labor supply of a continuum of differentially skilled workers and endogenous technology development and adoption choices of firms.

[1] An alternative approach is to consider direct taxes on specific technologies, as, for example, in Thuemmel (2018) and Costinot and Werning (2018). Yet, it is unlikely that in practice such direct taxes can be targeted perfectly to all technologies that raise wage inequality. Whenever the targeting of direct taxes is imperfect, there is scope for exploiting directed technical change effects via the income tax, which is what I focus on here.

In the model, technical change is driven by technology firms' decisions in which type of technology to invest. Some types are more complementary to high-skilled workers, some are more complementary to low-skilled workers. Technology firms' investment decisions depend on final good firms' demand for intermediate goods that embody the different types of technologies. This intermediate good demand in turn crucially depends on the structure of labor supply firms face on the labor market. If there is a relatively large supply of low-skilled workers, firms demand technologies that are relatively complementary to the low-skilled; if the supply of high-skilled workers is relatively large, firms demand more skill-biased technologies.

Income taxes interact with technology via the structure of labor supply. For example, raising marginal tax rates for high incomes and reducing them for low incomes discourages labor supply of high-skilled and encourages labor supply of low-skilled workers. This shifts firms' demand towards less skill-biased technologies, to which technology firms respond by shifting investment towards such technologies. Intuitively, progressive tax reforms should therefore induce technical change in favor of less skilled workers.

I examine this intuition formally and investigate its implications for the design of optimal taxes. To this end, I first show that the model's equilibrium has a parsimonious reduced form, which makes the tax analysis tractable. Importantly, the reduced form equations determining wages and technology are well studied by the theory of directed technical change (e.g. Acemoglu, 2007; Loebbing, 2018). Moreover, Acemoglu (2007) shows that they apply to a large set of directed technical change models studied in the literature. This makes my tax analysis, which is based exclusively on the reduced form, generic within the theory of directed technical change.

Turning to the analysis of income taxes, I first study the effects of tax reforms on the direction of technical change. In line with the intuition developed above, I find that, under certain conditions, progressive tax reforms induce technical change that compresses the pre-tax wage distribution. Hence, a progressive tax reform not only achieves a more equal distribution of post-tax incomes but potentially also lowers pre-tax wage inequality.

This is reflected in the shape of the optimal tax. Compared to a suitably defined benchmark with exogenous technology, the optimal tax is more progressive, featuring higher marginal tax rates at the top and lower marginal tax rates at the bottom of the income distribution. Intuitively, the optimal tax capitalizes on the reduction in pre-tax wage inequality brought about by the technical change induced by a more progressive tax.

The benchmark scenario with exogenous technology still features a non-linear, concave production structure, which gives rise to complementarities between different workers' labor inputs as in Stiglitz (1982) and Sachs, Tsyvinski and Werquin (2020). Comparing optimal taxes instead to an appropriately specified scenario with completely exogenous wages (as in Mirrlees, 1971, and most of the subsequent literature), the results depend on whether directed technical change dominates the effects from imperfect worker substitution within a given technology. If directed technical change dominates – which is theoretically possible and empirically plausible – the optimal tax is even more progressive than in the scenario with exogenous wages. Otherwise, the optimal tax is sandwiched between the case with exogenous wages and the case with exogenous technology. Hence, directed technical change not only mitigates but potentially even reverses the impact of within-technology complementarity between workers on the optimal tax.

To assess the quantitative relevance of these results, I calibrate the model based on the empirical literature on directed technical change. I first quantify the effects of tax reforms on the wage distribution. In this regard, my theoretical results imply that regressive tax reforms induce skill-biased technical change. I therefore ask whether the regressive reforms of the US tax system since the 1970s, via directed technical, have played a role in the concurrent rise of US wage inequality. For that, I simulate the effects of reversing the cumulative reforms of the US tax and transfer system between 1970 and 2005, by taking the tax system back from its 2005 to its 1970 state. With an optimistic calibration of directed technical change effects – in which directed technical change dominates the effects of within-technology complementarities – I find that the hypothetical reform reduces the 90-10-percentile ratio of the wage distribution by up to 2.6%. Inversely, by the same metric, the regressive reforms between 1970 and 2005 can account for up to 9% (2.6% out of a total of 30%) of the total concurrent rise in US wage inequality.[2]

Turning to optimal taxes, the impact of directed technical change is substantial. Relative to the exogenous technology benchmark, optimal marginal tax rates increase by 3 to 8 percentage points on high incomes and decrease by 5 to 17 percentage points on low incomes. With an optimistic, but still empirically plausible, calibration of directed technical change effects, the optimal tax is also significantly more progressive than in the benchmark with exogenous wages. With the same calibration, optimal marginal tax rates increase monotonically over the bulk of the income distribution, while they follow a pronounced U-shape when ignoring directed technical change (in the exogenous technology benchmark).

The structure of the paper is as follows. Section 3.3 presents the model, introduces special cases, and defines some key elasticity concepts. Section 3.4 states important results from the theory of directed technical change, which provide the basis for the analysis in the present paper. Section 3.5 contains the analysis of tax reforms, while Section 3.6 studies optimal taxes. Section 3.7 quantifies

[2]The worker side of the model is kept deliberately simple to enable a non-parametric (or, Mirrleesian) optimal tax analysis. Thereby, it ignores some adjustment margins to taxes, such as endogenous education or occupation choices, which might amplify the effects of tax reforms on the wage distribution. The true effect of the regressive tax reforms on US wage inequality may thus be well above 9% of the actually observed increase. A more comprehensive analysis of this question is left for future research.

the results from the preceding sections and Section 3.8 concludes.

3.2 Related Literature

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

In the literature on optimal taxation, it is closely related to Stiglitz (1982) and Sachs et al. (2020). These papers analyze the implications of complementarity between different types of workers for the design of non-linear labor income taxes. 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 their analysis to incorporate directed technical change effects. Directed technical change counteracts (within-technology) complementarity between workers. Hence, I find that accounting for directed technical change reduces optimal marginal tax rates at the bottom and increases them at the top. With a conservative calibration of directed technical change effects, directed technical change and within-technology complementarity effects offset each other approximately and the results are close to those obtained with exogenous wages (as, e.g., in Diamond, 1998; Saez, 2001). With a more optimistic calibration of directed technical change effects, directed technical change dominates and makes the optimal tax more progressive than with exogenous wages.

In a conceptually similar contribution, Rothschild and Scheuer (2013) extend a model à la Stiglitz (1982) to incorporate endogenous sorting into occupations. 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). The implications of directed technical change are qualitatively different, because with directed technical change, overcompensation of (within-technology) complementarity effects is a possibility and the optimal tax can be more progressive than with exogenous wages.

Ales et al. (2015) and Jacobs and Thuemmel (2018b) (see also Jacobs and Thuemmel 2018a) analyze the effects of skill-biased technical change on optimal taxes. They treat skill-biased technical change as an exogenous change in the production technology, whereas in my analysis, the degree to which technology is skill-biased is endogenous and responds to the tax system. My approach to the analysis of technical change and taxes is therefore conceptually different from theirs and produces different results.

Another set of related studies analyzes optimal direct taxes on specific technologies, such as industrial robots (Guerreiro, Rebelo and Teles, 2018; Thuemmel, 2018; Costinot and Werning, 2018).[3] In these studies, direct taxes can be perfectly targeted towards a particular technology, the labor market impact of which is known. While optimal in theory, this is challenging in practice: besides detailed information about the labor market impact of the technology and its various applications, it requires that the government be able to monitor which technology is used in which ways in any given firm. These requirements are unlikely to be satisfied for more than a few well-studied examples. My approach, in contrast, does not rely on this type of information or enforcement capabilities. Hence, it is complementary to work on direct technology taxes in that it applies to technologies for which a specific direct tax is not available.[4]

Finally, Jagadeesan (2019) and Jones (2019) analyze optimal taxation in settings where the speed of technical progress is endogenous. They find that optimal taxes on labor income are reduced relative to a setting where technical progress is exogenous. Yet, technical progress is always unbiased in their models, in the sense that it does not affect relative wages between workers. This precludes an analysis of the issues that are at the heart of the present paper. I therefore view my work as complementary to theirs.

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 (2018) 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 Section 3.7. In particular, I use estimates from Lewis (2011), Dustmann and Glitz (2015), Morrow and Trefler (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 3.7.

[3] See also Naito (1999), who provides the general theoretical argument for distorting production efficiency to reduce pre-tax wage inequality when the income tax cannot condition on worker types.

[4] Slavik and Yazici (2014) analyze the optimal differential taxation of structure and equipment capital, based on the assumption that equipment capital is skill-biased. In that case, the concern is not that the targeted technology is narrow but that the targeting is imprecise: there are 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.

3.3 Setup

I merge a standard Mirrlees (1971) model of optimal income taxation with a directed technical change model in the spirit of Acemoglu (2007).[5] While the model is specific in several respects, my tax analysis only uses a reduced subset of the model's equilibrium conditions. This subset is much more general than the model itself. In fact, it can be obtained from any of the models presented in Acemoglu (2007) and Loebbing (2018, Appendix B.1) once they are augmented to include endogenous labor supply.[6]

3.3.1 Model

The model features heterogeneous workers, perfectly competitive final good firms, monopolistically competitive technology firms, and a government that levies taxes.

Workers There is a continuum of workers with different types $\theta \in \Theta = [\underline{\theta}, \overline{\theta}] \subset \mathbb{R}$. 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.

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

$$c_\theta = R_T(w_\theta l_\theta) + S \ , \tag{3.1}$$

where S is a lump-sum transfer used to neutralize the government's budget.

Workers choose their labor supply to maximize utility, taking wages as given. The first-order

[5] The 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.

[6] The main reason for presenting a specific model, besides providing guidance to the reader, is that it allows to make explicit the restrictions on government policy by which the optimal tax in Section 3.6 becomes second best. See the description of the government below.

condition is (assuming that the tax function is differentiable):

$$v'(l_\theta) = R'_T(w_\theta l_\theta) w_\theta \ . \tag{3.2}$$

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 $G(L_i, Q_i)$. 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 $Q_i = \{Q_{i,j}\}_{j\in\{1,2,\ldots,J\}}$ collects the variables $Q_{i,j}$, each of which is an aggregate of a continuum of technology-embodying intermediate goods:[7]

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

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 with the same j-index. 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.

With this structure of final good production, we can write the output of firm i as $\widetilde{G}(L_i, \phi, q_i)$ where $\phi = \{\phi_{j,k}\}_{(j,k)\in\{1,2,\ldots,J\}\times[0,1]}$ and $q_i = \{q_{i,j,k}\}_{(j,k)\in\{1,2,\ldots,J\}\times[0,1]}$ collect qualities and quantities of all different intermediate inputs. I assume that the function $\widetilde{G}$ 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 $\widetilde{G}$ 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 labor:

$$w_\theta = D_{L_\theta} \widetilde{G}(L, \phi, q) \ . \tag{3.3}$$

The operator D_{L_θ} denotes Gateaux differentiation with respect to L in direction of the Dirac measure at θ, which I define rigorously in Section 3.3.2. Labor market clearing requires that the

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

aggregate labor demand L_θ equals the sum of individual workers' labor supply,

$$L_\theta = l_\theta h_\theta \quad \text{for all}\, \theta \ .$$

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

$$p_{j,k} = \alpha \phi_{j,k} q_{j,k}^{\alpha-1} \frac{\partial G(L,Q)}{\partial Q_j} \ , \tag{3.4}$$

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

Technology Firms The technology-embodying 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

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

subject to the demand from final good firms (equation (3.4)). 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:

$$p_{j,k} = \frac{\eta_j}{(1+\xi)\alpha} \ . \tag{3.5}$$

Technology firms can invest R&D resources to improve the quality of their output. In particular, 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 condition for the choice of quality is given by

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

where $q_{j,k}$ is assumed to take its profit-maximizing value implied by equation (3.5). 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 (3.3.1) 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 $\phi_{j,k}$, which ensures that the first-order condition identifies the unique profit maximum. Since profits are symmetric across all firms (j,k) with the same j-index, uniqueness of the profit maximum implies that the choices of all firms with index j are the same and we can drop the k-index henceforth.

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

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

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.[8]

Government The government levies an income tax, a profit tax, and a uniform tax/subsidy on intermediate goods, which cannot differentiate between intermediate good types.[9]

Since final good firms' production function $\widetilde{G}$ 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 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, levied on technology firms and the owners of R&D resources, is assumed to be confiscatory to avoid a role for the distribution of firm ownership without a meaningful theory of wealth formation in the model.[10]

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.[11]

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

[9] From 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.

[10] Note 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. Alternatively, I could assume that firm ownership and the ownership of R&D resources are uniformly distributed across workers without changing any of the results.

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

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

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

which is redistributed lump-sum across workers.

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 production at labor input l and a given set of quality levels ϕ can be written as (because intermediate good prices equal marginal cost):

$$F(l, \phi) \coloneqq \max_q \left\{ \widetilde{G}(\{h_\theta l_\theta\}_{\theta \in \Theta}, \phi, q) - \sum_{j=1}^J \eta_j q_j \right\} \ . \tag{3.6}$$

Note that I used labor market clearing (equation (3.3.1)) 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 (3.3) then implies that in equilibrium, wages are given by

$$w_\theta(l, \phi) = \frac{1}{h_\theta} D_{l_\theta} F(l, \phi) \ , \tag{3.7}$$

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 (equation (3.3.1)) 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 $\Phi = \left\{ \phi \in \mathbb{R}^J_+ \mid \sum_{j=1}^J C_j(\phi_j) \leq \overline{C} \right\}$. Thus,

$$\phi^*(l) \coloneqq \operatorname*{argmax}_{\phi \in \Phi} F(l, \phi) \tag{3.8}$$

is an equilibrium technology. In the following I focus on equilibria in which technology satisfies equation (3.8). Existence of other equilibria can be ruled out by imposing assumptions that guarantee strict quasiconcavity of F in ϕ under the constraint $\phi \in \Phi$.[12]

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:

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

It follows that the revenue from corporate taxes and the expenses on technology good subsidies offset each other exactly in equation (3.3.1), such that the expression for government revenue shrinks to

$$S(y) = \int_{\Theta} T(y_\theta) h_\theta \, d\theta \ . \tag{3.9}$$

The equilibrium values for wages w, technology ϕ, labor inputs l, consumption levels c, and government revenue S can now be characterized by the wage equation (3.7), the technology condition (3.8), workers' first-order conditions (3.2), their budget constraints (3.1), and the equation for government revenue (3.9). These equilibrium conditions provide the starting point for the tax analysis in the following sections.

The wage and technology equations (3.7) and (3.8) are identical to the conditions characterizing equilibrium in the directed technical change models presented in Acemoglu (2007) and Loebbing (2018, Appendix B.1). 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 present model but applies more generally within a large class of directed technical change models.

3.3.2 Derivative and Elasticity Concepts

The tax analysis uses functional derivatives and various elasticities. To simplify the exposition I define a specific notation for several frequently used expressions.

[12] In particular, if the constrained function

$$\widetilde{F}(l, \phi_{-J}) := F(l, \phi_{-J}, \widetilde{\phi}_J(\phi_{-J})), \text{ where } \phi_{-J} = \{\phi_j\}_{j \in \{1,2,\ldots,J-1\}} \text{ and } \widetilde{\phi}_J(\phi_{-J}) = C_J^{-1}\left(\overline{C} - \sum_{j=1}^{J-1} C_j(\phi_j)\right),$$

is strictly quasiconcave in ϕ_{-J}, the first-order conditions for a maximum of $\widetilde{F}$ in ϕ_{-J} are necessary and sufficient and there is a unique value $\phi^*_{-J}(l)$ that satisfies them. Equivalently, there is a unique value $\phi^*(l)$ satisfying the first-order conditions of the program (3.8), which are identical to the equilibrium condition (3.3.1), and this unique value indeed solves the program.

Functional Derivatives For derivatives in finite-dimensional spaces I use standard notation. For perturbations of the tax function T and labor input l I will frequently use the following functional derivatives.

Let $x : (T, z) \mapsto x(T, z)$ be a function of the tax T and, potentially, further variables z. Then,

$$D_\tau x(T,z) := \left.\frac{dx(T+\mu\tau, z)}{d\mu}\right|_{\mu=0}$$

denotes the directional derivative of x with respect to T in direction of the tax reform τ.

Similarly, let $x : (l, z) \mapsto x(l, z)$ be a function of labor input l and, potentially, further variables z. I formalize the derivative of x with respect to labor supply of a given type θ, l_θ, as[13]

$$D_{l_\theta} x(l,z) := \lim_{\Delta\to 0}\frac{1}{\Delta}\left.\frac{dx(l+\mu\widetilde{l}_{\Delta,\theta}, z)}{d\mu}\right|_{\mu=0},$$

where $\widetilde{l}_{\Delta,\theta} : \widetilde{\theta} \mapsto \widetilde{l}_{\Delta,\theta,\widetilde{\theta}}$ is a real-valued function on the type space. For interior types $\theta \in (\underline{\theta}, \overline{\theta})$ it is given by

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

for the highest type $\overline{\theta}$ by

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

and for the lowest type $\underline{\theta}$ by

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

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. Appendix B.1.1 demonstrates

[13]The derivative of a function with respect to aggregate labor supply L_θ is defined analogously.

that the thus defined derivative works in a natural way by showing in detail that

$$D_{L_\theta} \int_{\underline{\theta}}^{\overline{\theta}} w_{\widetilde{\theta}} L_{\widetilde{\theta}} \, d\widetilde{\theta} = w_\theta \quad \forall \theta \; .$$

This also proves the labor demand equation (3.3).

The tax analysis below often distinguishes between the direct effect of changes in T or l on an outcome x and the indirect effect mediated through the response of technology ϕ^*. In particular, suppose $x : (T, \phi) \mapsto x(T, \phi)$ depends (directly) on taxes T and technology ϕ. The direct effect of a tax reform in direction τ, holding technology fixed, is then given by $D_\tau x(T, \phi)$ as defined above. For the indirect effect of the tax reform via technology (the directed technical change effect, henceforth) I introduce the following notation:

$$D_{\phi,\tau} x(T, \phi^*(T)) := \left. \frac{dx(T, \phi^*(T + \mu\tau))}{d\mu} \right|_{\mu=0} \; .$$

Here, $\phi^*(T)$ denotes the equilibrium technology at tax function T. The total effect of the reform on x is then obtained as the sum of the direct and the directed technical change effect. Writing $x^*(T) := x(T, \phi^*(T))$, we get

$$D_\tau x^*(T) = D_\tau x(T, \phi^*(T)) + D_{\phi,\tau} x(T, \phi^*(T)) \; .$$

Analogously, if $x : (l, \phi) \mapsto x(l, \phi)$ is a function of labor input l and technology ϕ, the induced technical change effect of a labor input change in direction l_θ is

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

where $\phi^*(l)$ is given by equation (3.8).

Wage Elasticities The response of wages to labor input changes plays a central role in the tax analysis. Consider wages as given by (3.7), that is, for each type θ the wage w_θ is a function of labor inputs l and technology ϕ.

The first set of wage elasticities is concerned with the direct effect of labor inputs on wages, holding technology constant. I call these elasticities the within-technology substitution elasticities (sometimes just substitution elasticities, for brevity), as they describe the changes in marginal productivities induced by factor substitution within a given technology.[14] The own-wage substitution

[14] Within-technology substitution elasticities are equivalent to the wage elasticities introduced by Sachs et al. (2020).

elasticity, that is, the elasticity of w_θ with respect to l_θ, is defined as

$$\gamma_{\theta,\theta} \coloneqq \frac{l_\theta}{w_\theta} \lim_{\Delta \to 0} \left. \frac{dw_\theta(l + \mu \widetilde{l}_{\Delta,\theta}, \phi)}{d\mu} \right|_{\mu=0} .$$

Alternatively, we could write the wage w_θ as a function of ϕ, l, and type θ's labor input l_θ separately, as typically a type's labor input affects its own wage in a way distinct from the labor input function l (see for example the CES case in Section 3.3.4). Then, the own-wage substitution elasticity is simply

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

The cross-wage substitution elasticity, that is, the elasticity of w_θ with respect to a different type's labor input $l_{\widetilde{\theta}}$ (with $\widetilde{\theta} \neq \theta$), is given by

$$\gamma_{\theta,\widetilde{\theta}} \coloneqq \frac{l_{\widetilde{\theta}}}{w_\theta} D_{l_{\widetilde{\theta}}} w_\theta(l, \phi) ,$$

with the derivative $D_{l_{\widetilde{\theta}}}$ as defined above.

The second set of wage elasticities captures the directed technical change effects of changes in labor inputs on wages. These elasticities are called technical change elasticities in the following. The own-wage technical change elasticity is defined as

$$\rho_{\theta,\theta} \coloneqq \frac{l_\theta}{w_\theta} \lim_{\Delta \to 0} \left. \frac{dw_\theta(l, \phi^*(l + \mu \widetilde{l}_{\Delta,\theta}))}{d\mu} \right|_{\mu=0} .$$

Again, the CES case in Section 3.3.4 clarifies why this is a natural definition of the own-wage technical change elasticity and how it can be expressed in terms of conventional partial derivatives.

The cross-wage technical change elasticity measures how wage w_θ is affected by a change in another type's labor supply $l_{\widetilde{\theta}}$ via directed technical change. Formally, it is given by

$$\rho_{\theta,\widetilde{\theta}} \coloneqq \frac{l_{\widetilde{\theta}}}{w_\theta} D_{\phi, l_{\widetilde{\theta}}} w_\theta(l, \phi^*(l)) ,$$

where the derivative $D_{\phi, l_{\widetilde{\theta}}}$ has been defined above.

Rate of Progressivity The rate of progressivity of a tax schedule T is defined as minus the elasticity of the marginal retention rate R'_T with respect to income,

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

It measures the progression of marginal tax rates as income increases. If the income tax is linear such that marginal tax rates are constant, $P_T(y)$ is zero. If the income tax is progressive (regressive) in the sense that marginal tax rates increase (decrease) with income, the rate of progressivity is positive (negative).

Labor Supply Elasticities To express the effects of tax reforms compactly, I also define some standard concepts of labor supply elasticities. The first is the hypothetical elasticity of labor supply with respect to the marginal retention rate that would obtain if the retention function were linear:

$$e_\theta(l) := \frac{v'(l_\theta)}{v''(l_\theta)l_\theta} \ .$$

Consider now the labor supply of an arbitrary worker type θ, given by workers' first-order condition (3.2), as a function of T and w_θ. The true elasticity of labor supply with respect to the marginal retention rate must take into account potential non linearities of the retention function, which cause the worker's marginal retention rate to change as labor supply changes (e.g. Sachs et al., 2020). This elasticity is given by

$$\epsilon_\theta^R(T, l, w) := \frac{R'_T(w_\theta l_\theta)}{l_\theta} D_{\widetilde{\tau}} l_\theta(T, w_\theta) \ ,$$

where the auxiliary tax reform $\widetilde{\tau}$ 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:

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

Inserting this into workers' first-order condition and differentiating with respect to μ (at $\mu = 0$) then gives exactly the local response of individual labor supply to a one unit increase in the marginal retention rate. This leads to the following expression for the elasticity of labor supply with respect to the marginal retention rate (see Appendix B.1.1 for details):

$$\epsilon_\theta^R(T, l, w) = \frac{e_\theta(l)}{1 + e_\theta(l) P_T(w_\theta l_\theta)} \ . \tag{3.10}$$

For a locally linear tax function, that is, for $P_T(w_\theta l_\theta) = 0$, the elasticity coincides with the hypothetical elasticity e_θ defined above.

Finally, define the elasticity of labor supply with respect to the wage as

$$\epsilon^w_\theta(T, l, w) := \frac{w_\theta}{l_\theta} \frac{\partial l_\theta(T, w_\theta)}{\partial w_\theta} \; .$$

It is a standard result that this elasticity can be written as (see Appendix B.1.1 for details)

$$\epsilon^w_\theta(T, l, w) = \frac{(1 - P_T(w_\theta l_\theta))e_\theta(l)}{1 + e_\theta(l) P_T(w_\theta l_\theta)} \; . \tag{3.11}$$

3.3.3 Assumptions for the Tax Analysis

The following assumptions are maintained throughout the paper.

Assumption 3.

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

The first three parts of the assumption ensure that the wage elasticities $\gamma_{\theta,\widetilde{\theta}}$ and $\rho_{\theta,\widetilde{\theta}}$ are well defined. The last part guarantees that workers' second-order conditions are satisfied strictly under a given tax T, such that the labor supply elasticities ϵ^R_θ and ϵ^w_θ are well defined (see Appendix B.1.1).

3.3.4 Special Cases

Under further restrictions, we can obtain particularly tractable special cases of the model, which allow to derive additional qualitative insights in the tax analysis.

CES Production An important special case is obtained when the aggregate production function F features a constant elasticity of substitution (CES) between worker types while the research cost functions are isoelastic. In this case, aggregate production takes the form

$$F(l, \phi) = \left[\int_{\underline{\theta}}^{\overline{\theta}} (\kappa_\theta \phi_\theta l_\theta h_\theta)^{\frac{\sigma - 1}{\sigma}} \, d\theta \right]^{\frac{\sigma}{\sigma - 1}} \; , \tag{3.12}$$

where $\sigma > 0$ is the elasticity of substitution and κ is a continuously differentiable function that assigns an exogenous productivity level to each type of worker. The endogenous technology ϕ similarly takes the form of a function assigning (endogenous) productivity levels to workers. The set of feasible technologies is given by

$$\Phi = \left\{ \phi : \theta \mapsto \phi_\theta \in \mathbb{R}_+ \mid \int_{\underline{\theta}}^{\overline{\theta}} \phi_\theta^\delta \, d\theta \leq \overline{C} \right\} , \tag{3.13}$$

where δ governs the substitutability of productivity levels across worker types.

I derive the expressions for aggregate production and the set of feasible technologies from restrictions on the fundamentals of the model in Appendix B.1.1. There, I also show that wages in the CES case are given by

$$w_\theta(l, \phi) = (\kappa_\theta \phi_\theta)^{\frac{\sigma-1}{\sigma}} (l_\theta h_\theta)^{-\frac{1}{\sigma}} F(l, \phi)^{\frac{1}{\sigma}} . \tag{3.14}$$

Accordingly, the own-wage substitution elasticity becomes[15]

$$\gamma_{\theta,\theta} = -\frac{1}{\sigma} =: \gamma^{CES} , \tag{3.15}$$

while the cross-wage substitution elasticity is

$$\gamma_{\theta,\widetilde{\theta}} = \frac{1}{\sigma} \frac{l_{\widetilde{\theta}} w_{\widetilde{\theta}} h_{\widetilde{\theta}}}{F(l, \phi)} . \tag{3.16}$$

The own-wage technical change elasticity is given by

$$\rho_{\theta,\theta} = \frac{(\sigma - 1)^2}{(\delta - 1)\sigma^2 + \sigma} =: \rho^{CES} , \tag{3.17}$$

and the cross-wage technical change elasticity becomes

$$\rho_{\theta,\widetilde{\theta}} = -\frac{(\sigma - 1)^2}{(\delta - 1)\sigma^2 + \sigma} \frac{l_{\widetilde{\theta}} w_{\widetilde{\theta}} h_{\widetilde{\theta}}}{F(l, \phi)} . \tag{3.18}$$

Isoelastic Disutility of Labor When the disutility of labor is isoelastic, workers' utility functions take the form

$$u_\theta = c_\theta - \frac{e}{e+1} l_\theta^{\frac{e+1}{e}} .$$

[15] See again Appendix B.1.1 for the derivations of the wage elasticities in the CES case.

In this case, the hypothetical labor supply elasticity $e_\theta(l)$ is constant across θ and l:

$$e_\theta(l) = e \quad \text{for all } \theta, l.$$

Constant-Rate-of-Progressivity Taxes A constant-rate-of-progressivity (CRP) tax function takes the form (e.g. Feldstein, 1969; Heathcote, Storesletten and Violante, 2017)

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

For any CRP tax schedule T the rate of progressivity P_T is constant across income levels:

$$P_T(y) = P \quad \text{for all } y.$$

This special case, when combined with isoelastic disutility of labor, ensures that the labor supply elasticities ϵ_θ^R and ϵ_θ^w are constant in θ.

3.4 Directed Technical Change

Directed technical change theory makes predictions about the relationship between labor inputs, technology, and wages as governed by equations (3.7) and (3.8) (copied here for convenience):

$$w_\theta(l, \phi) = \frac{1}{h_\theta} D_{l_\theta} F(l, \phi)$$
$$\phi^*(l) := \operatorname*{argmax}_{\phi \in \Phi} F(l, \phi) \ .$$

The theory requires the following assumption.

Definition 5. A technology ϕ is more skill-biased than another technology $\widetilde{\phi}$ if, for any labor input l, all skill premia are greater under ϕ than under $\widetilde{\phi}$, that is,

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

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

We write $\phi \succeq^{sb} \widetilde{\phi}$.

Assumption 4. *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 $\widetilde{\phi}$, if ϕ weakly raises output* *relative to all technologies $\underline{\phi}$ that are less skill-biased than both ϕ and $\widetilde{\phi}$, then there must exist a technology $\overline{\phi}$ that weakly raises output relative to $\widetilde{\phi}$ and is more skill-biased than both ϕ and $\widetilde{\phi}$.*[16]

Quasisupermodularity requires that changes in technology that raise skill premia in different parts of the wage distribution must not be substitutes. To illustrate, suppose there are two new technologies, one that raises skill premia in the top half of the wage distribution and one that raises skill premia in the bottom half. Assumption 4 now requires that, if the new technology that raises skill premia at the top leads to an increase in output absent the other technology, then it must also increase output when the other technology is already implemented.

The CES production function naturally satisfies quasisupermodularity. Reshuffling productivity levels in the upper half of the type space has no bearing on whether a certain rearrangement of productivity levels in the lower half enhances output or not.

In Loebbing (2018, Section 5), I present further important examples that satisfy quasisupermodularity. In particular, an assignment model à la Costinot and Vogel (2010) satisfies quasisupermodularity in various forms. First, when treating the matching function from skills to tasks as the endogenous technology variable ϕ, the baseline model by Costinot and Vogel (2010) satisfies equations (3.7) and (3.8) plus Assumption 4. Hence, my analysis encompasses the case where directed technical change takes the form of endogenous changes in the allocation of workers to tasks.[17]

Second, when augmenting the baseline assignment model to include capital as an additional production factor, the extent of automation, as measured by the set of tasks performed by capital, can take the place of the endogenous technology variable ϕ while preserving all of the above conditions. Finally, we can endogenize the productivity of capital in such a setting and study its endogenous adjustments as directed technical change (Loebbing, 2018).[18]

3.4.1 Weak Relative Bias

Under Assumption 4, any increase in the relative supply of skill in the economy induces skill-biased technical change.

[16]Note 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 $\widetilde{\phi}$ 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 $\widetilde{\phi}$. In particular, for any l and any $\phi, \widetilde{\phi}$, if $F(l, \underline{\phi}) \leq F(l, \phi)$, then $F(l, \overline{\phi}) \geq F(l, \widetilde{\phi})$, where $\underline{\phi}$ and $\overline{\phi}$ denote infimum and supremum of ϕ and $\widetilde{\phi}$. My definition is slightly less restrictive (and sufficiently restrictive for the present purpose).

[17]Yet, in this case there won't be strong relative bias, as there is no source of (quasi-)convexity in the aggregate production function, see below.

[18]Note also that the predictions derived under quasisupermodularity (see Lemma 6 below) receive support in the empirical literature, as discussed in Section 3.7.

Lemma 6. *Take any labor input l and let dl be a change in the labor input such that dl_θ/l_θ increases in θ. Then, the technical change induced by dl raises more skilled workers' wages relative to less skilled workers' wages, that is,*

$$\frac{1}{w_\theta}\left.\frac{dw_\theta(l,\phi^*(l+\mu dl))}{d\mu}\right|_{\mu=0} - \frac{1}{w_{\widetilde{\theta}}}\left.\frac{dw_{\widetilde{\theta}}(l,\phi^*(l+\mu dl))}{d\mu}\right|_{\mu=0} \geq 0 \tag{3.19}$$

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

Proof. See Appendix B.1.2. □

The intuition behind this result is that an increase in relative skill supply raises the profitability of technologies that are relatively complementary to high-skilled workers and these technologies in turn raise the relative productivity of the high-skilled.

The within-technology substitution effect of an increase in relative skill supply typically has the opposite direction and reduces the relative wages of high-skilled workers. An important question is then whether the directed technical change effect (raising skill premia) or the within-technology substitution effect (reducing skill premia) dominates.

3.4.2 Strong Relative Bias

When the directed technical change effect dominates, we say that there is strong relative bias of technology (Acemoglu, 2002). Formally, strong relative bias of technology at a given labor input l means that

$$\frac{1}{w_\theta}\left.\frac{dw^*_\theta(l+\mu dl)}{d\mu}\right|_{\mu=0} - \frac{1}{w_{\widetilde{\theta}}}\left.\frac{dw^*_{\widetilde{\theta}}(l+\mu dl)}{d\mu}\right|_{\mu=0} \geq 0 \tag{3.20}$$

for all $\theta \geq \widetilde{\theta}$ and for any labor supply change dl such that dl_θ/l_θ increases in θ. In Loebbing (2018), I show that strong relative bias occurs if and only if aggregate production becomes quasiconvex in labor when taking into account the endogenous adjustment of technology.

In the CES case, there is a simple parametric condition for strong relative bias. In particular, using the notation for wage elasticities introduced above, the total effect of a labor supply change dl on the relative wage between types $\theta \geq \widetilde{\theta}$ is

$$(\rho_{\theta,\theta}+\gamma_{\theta,\theta})\frac{dl_\theta}{l_\theta} + \int_{\underline{\theta}}^{\overline{\theta}}(\rho_{\theta,\theta'}+\gamma_{\theta,\theta'})\frac{dl_{\theta'}}{l_{\theta'}}\,d\theta' - (\rho_{\widetilde{\theta},\widetilde{\theta}}+\gamma_{\widetilde{\theta},\widetilde{\theta}})\frac{dl_{\widetilde{\theta}}}{l_{\widetilde{\theta}}} - \int_{\underline{\theta}}^{\overline{\theta}}(\rho_{\widetilde{\theta},\theta'}+\gamma_{\widetilde{\theta},\theta'})\frac{dl_{\theta'}}{l_{\theta'}}\,d\theta'$$
$$= (\rho^{CES}+\gamma^{CES})\left(\frac{dl_\theta}{l_\theta}-\frac{dl_{\widetilde{\theta}}}{l_{\widetilde{\theta}}}\right).$$

For an increase in relative skill supply ($dl_\theta/l_\theta \geq dl_{\widetilde{\theta}}/l_{\widetilde{\theta}}$), the effect on the relative wage is positive if and only if

$$\rho^{CES} + \gamma^{CES} \geq 0 \ . \tag{3.21}$$

Hence, if condition (3.21) is satisfied, directed technical change effects dominate within-technology substitution effects and skill premia increase with relative skill supply.

3.5 Tax Reforms

Starting from a given tax T, a tax reform is represented by the change from T to $T + \mu\tau$, where $\mu \in \mathbb{R}_+$ 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.

3.5.1 Progressive and Regressive Reforms

The curvature of τ, relative to the curvature of T, governs the progressivity of the reform. More precisely, I call a reform progressive if the post-reform tax schedule has a higher rate of progressivity than the pre-reform schedule everywhere.

Definition 6. Starting from tax T the reform (τ, μ) is progressive if and only if

$$P_{\widetilde{T}}(y) \geq P_T(y) \quad \forall\, y \ ,$$

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

This definition is equivalent to the following characterizations of progressivity.

Lemma 7. *Take any tax function T. The following statements are equivalent.*

1. *The reform (τ, μ) is progressive according to Definition 6.*

2. *The post-reform tax $\widetilde{T} = T + \mu\tau$ can be obtained by taxing post-tax income under the initial tax in a progressive way, that is, by means of a tax function with increasing marginal tax rates. Formally,*

$$R_{\widetilde{T}} = r \circ R_T$$

for some concave function r.

3. *The reform* (τ, μ) *satisfies*

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

Proof. See Appendix B.1.3. □

The first equivalence provides an intuitive interpretation of progressivity: a reform is progressive if and only if it can be obtained by augmenting the initial tax by an additional tax on post-tax income that features increasing marginal tax rates. The second equivalence shows that a progressive reform raises the marginal tax rate relative to the initial marginal retention rate by more for higher incomes. This equivalence will turn out useful in the analysis below.

A regressive reform is defined as the inverse of a progressive reform: it reduces the rate of progressivity of the tax schedule everywhere.

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 $\mu \to 0$. Note that this does not lead to confusion with the definition of progressivity, because, as indicated by the second equivalence in Lemma 7, 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 wages under the initial tax schedule, that is, $w_\theta \leq w_{\widetilde{\theta}}$ if $\theta \leq \widetilde{\theta}$ under the initial tax.

To describe the effects of tax reforms on economic outcomes formally, 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]

3.5.2 Effects on Labor Inputs

A key step in the analysis is to characterize the responses of labor inputs to a given tax reform. Let

$$\widehat{l}_{\theta,\tau}(T) := \frac{1}{l_\theta} D_\tau l_\theta(T)$$

[19] Note that in some cases this involves an abuse of notation. I write for example $w_\theta(l, \phi)$ in equation (3.7) to denote wages as a function of labor inputs and technology; now I use $w_\theta(T, \phi^*(T))$ 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.

denote the relative change in the labor input of type θ in response to reform τ. Relative labor input changes must satisfy the following fixed point equation (see Appendix B.1.3 for details):[20]

$$\widehat{l}_{\theta,\tau}(T) = -\epsilon_\theta^R \frac{\tau'(y_\theta(T))}{1-T'(y_\theta(T))} + \epsilon_\theta^w(\gamma_{\theta,\theta}+\rho_{\theta,\theta})\widehat{l}_{\theta,\tau}(T) + \epsilon_\theta^w \int_{\underline{\theta}}^{\overline{\theta}} (\gamma_{\theta,\widetilde{\theta}}+\rho_{\theta,\widetilde{\theta}})\widehat{l}_{\widetilde{\theta},\tau}\, d\widetilde{\theta}\,. \tag{3.22}$$

In equilibrium, labor inputs respond directly to the tax reform (the first term in equation (3.22)) but they also cause wages to adjust. These wage adjustments in turn feed back to labor inputs, which is captured by the second and third terms in equation (3.22). Accounting for these feedback effects gives rise to the fixed point character of equation (3.22).

I characterize the fixed point of equation (3.22) by an iteration procedure. Within the iteration steps I disentangle the feedback effects purely transmitted via directed technical change from those transmitted via within-technology factor 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 structure of directed technical change effects predicted by the theory of directed technical change.[21]

Lemma 8. *Fix an initial tax T and suppose that*

$$\sup_{\theta\in\Theta}\left[(\epsilon_\theta^w\rho_{\theta,\theta})^2\right] + \int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}} (\epsilon_\theta^w\rho_{\theta,\widetilde{\theta}})^2\, d\widetilde{\theta}\, d\theta + 2\sqrt{\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}} \left(\epsilon_\theta^w\rho_{\theta,\theta}\epsilon_\theta^w\rho_{\theta,\widetilde{\theta}}\right)^2 d\widetilde{\theta}\, d\theta} < 1 \tag{3.23}$$

$$\sup_{\theta\in\Theta}\left[(\epsilon_\theta^w\zeta_{\theta,\theta})^2\right] + \int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}} (\epsilon_\theta^w\zeta_{\theta,\widetilde{\theta}})^2\, d\widetilde{\theta}\, d\theta + 2\sqrt{\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}} \left(\epsilon_\theta^w\zeta_{\theta,\theta}\epsilon_\theta^w\zeta_{\theta,\widetilde{\theta}}\right)^2 d\widetilde{\theta}\, d\theta} < 1\,, \tag{3.24}$$

where $\zeta_{\theta,\widetilde{\theta}} := \gamma_{\theta,\widetilde{\theta}} + \rho_{\theta,\widetilde{\theta}}$.[22]

[20] All elasticities in this section are evaluated at the equilibrium under the initial tax T. I do not write this dependence explicitly to save on notation.

[21] The representation of labor input responses in Lemma 8 is different from that provided by Sachs et al. (2020) even when ignoring directed technical change (i.e., when setting $\rho_{\theta,\widetilde{\theta}} = 0$ for all $\theta, \widetilde{\theta}$). I discuss the relationship between Lemma 8 and the results of Sachs et al. (2020) in Appendix B.3.1. In short, my approach has the advantage that, after decomposing the total effect, it allows me to derive analytical insights into the structure of the directed technical change component.

[22] Conditions (3.23) and (3.24) ensure that the series in equations (3.25) and (3.26) 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.

Then, the effect of tax reform τ on the labor input of type θ can be written as

$$\widehat{l}_{\theta,\tau}(T) = \sum_{n=0}^{\infty} \widehat{l}^{(n)}_{\theta,\tau}(T) \tag{3.25}$$

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

$$\begin{aligned} \widehat{l}^{(0)}_{\theta,\tau}(T) &= -\epsilon^R_\theta \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} \\ \widehat{l}^{(n)}_{\theta,\tau}(T) &= \epsilon^w_\theta \zeta_{\theta,\theta} \widehat{l}^{(n-1)}_{\theta,\tau}(T) + \epsilon^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \zeta_{\theta,\widetilde{\theta}} \widehat{l}^{(n-1)}_{\widetilde{\theta},\tau}(T)\, d\widetilde{\theta} \quad \forall n > 0 \,. \end{aligned}$$

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

$$\widehat{l}_{\theta,\tau}(T) = -\epsilon^R_\theta \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} + \underbrace{\sum_{n=1}^{\infty} \widetilde{TE}^{(n)}_{\theta,\tau}(T)}_{=:\widetilde{TE}_{\theta,\tau}(T)} + \underbrace{\sum_{n=1}^{\infty} \widetilde{SE}^{(n)}_{\theta,\tau}(T)}_{=:\widetilde{SE}_{\theta,\tau}(T)} \,, \tag{3.26}$$

where (omitting the argument T)

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

If ϵ^w_θ and ϵ^R_θ are constant in θ (e.g. because the disutility of labor is isoelastic and T is CRP), then $\widetilde{TE}_{\theta,\tau}$ is decreasing in θ for any progressive reform τ.

Proof. See Appendix B.1.3. □

Equation (3.25) 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 $\widehat{l}^{(0)}_{\theta,\tau}(T)$ is the direct effect of the reform on labor supply, holding wages constant. The direct adjustment of labor supply in turn changes wages, which then feeds back into labor supply. This first-round feedback effect is captured by $\widehat{l}^{(1)}_{\theta,\tau}(T)$. The labor supply change $\widehat{l}^{(1)}_{\theta,\tau}(T)$ then induces another adjustment of wages, which again affects labor supply, and so on.[23]

Equation (3.26) 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.

With constant labor supply elasticities across workers, the directed technical change component $\widetilde{TE}_{\theta,\tau}(T)$ is decreasing in θ for any progressive tax reform; that is, the directed technical change component reduces the labor supply of more relative to less skilled workers. This follows from the weak bias result of directed technical change theory. Intuitively, with constant labor supply elasticities, the direct effect of a progressive tax reform on labor supply reduces relative skill supply. By weak bias, this induces technical change reducing 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 $\widetilde{TE}_{\theta,\tau}(T)$, which must therefore reduce relative skill supply (i.e., decrease in θ) as well.

3.5.3 Directed Technical Change Effects

Consider now the relative wage changes that are caused by the technical change induced by a reform τ. Using the derivative $D_{\phi,\tau}$ introduced in Section 3.3.2, these relative wage changes are given by

$$\frac{1}{w_\theta} D_{\phi,\tau} w_\theta(T, \phi^*(T)) \; .$$

They can be expressed in terms of directed technical change elasticities and labor input responses as follows:

$$\frac{1}{w_\theta} D_{\phi,\tau} w_\theta(T, \phi^*(T)) = \rho_{\theta,\theta} \widehat{l}_{\theta,\tau}(T) + \int_{\underline{\theta}}^{\overline{\theta}} \rho_{\theta,\widetilde{\theta}} \widehat{l}_{\widetilde{\theta},\tau}(T) \, d\widetilde{\theta} \; . \tag{3.27}$$

[23] Mathematically, the series representation in equation (3.25) is the von Neumann series expansion of the solution to the fixed point equation (3.22). 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 von Neumann series expansion, this is equivalent to $\widehat{l}_\tau = \sum_{n=0}^{\infty} X^n Z$.

Inserting expression (3.26) from Lemma 8 into equation (3.27) yields an expression for the directed technical change effects of reform τ, consisting of three terms with intuitive interpretations. The slope of two of these terms can be signed using the structure of directed technical change effects imposed by weak bias.

Proposition 3. *Fix an initial tax T and let conditions* (3.23) *and* (3.24) *be satisfied.*

Then, the relative effect of the technical change induced by tax reform τ on wages can be written as

$$\begin{aligned}\frac{1}{w_\theta} D_{\phi,\tau} w_\theta(T, \phi^*(T)) = & \underbrace{\rho_{\theta,\theta}(-\epsilon_\theta^R) \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} + \int_{\underline{\theta}}^{\overline{\theta}} \rho_{\theta,\widetilde{\theta}}(-\epsilon_{\widetilde{\theta}}^R) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1 - T'(y_{\widetilde{\theta}}(T))} \, d\widetilde{\theta}}_{:=DE_{\theta,\tau}(T)} \\ & + \underbrace{\rho_{\theta,\theta} \widetilde{TE}_{\theta,\tau}(T) + \int_{\underline{\theta}}^{\overline{\theta}} \rho_{\theta,\widetilde{\theta}} \widetilde{TE}_{\widetilde{\theta},\tau}(T) \, d\widetilde{\theta}}_{=:TE_{\theta,\tau}(T)} \\ & + \underbrace{\rho_{\theta,\theta} \widetilde{SE}_{\theta,\tau}(T) + \int_{\underline{\theta}}^{\overline{\theta}} \rho_{\theta,\widetilde{\theta}} \widetilde{SE}_{\widetilde{\theta},\tau}(T) \, d\widetilde{\theta}}_{=:SE_{\theta,\tau}(T)} ,\end{aligned} \tag{3.28}$$

for all $\theta \in \Theta$, where $\widetilde{TE}_{\theta,\tau}(T)$ and $\widetilde{SE}_{\theta,\tau}(T)$ are defined in Lemma 8.

If ϵ_θ^w and ϵ_θ^R are constant in θ (e.g. because the disutility of labor is isoelastic and T is CRP), then $DE_{\theta,\tau}(T)$ and $TE_{\theta,\tau}(T)$ are decreasing in θ for any progressive reform τ.

Proof. See Appendix B.1.3. □

The terms in equation (3.28) follow directly from Lemma 8. The first line of equation (3.28) is the technical change effect on wages induced by the direct component of the labor supply response to the tax reform τ. It decreases in θ for any progressive reform (under constant labor supply elasticities) because, by Lemma 8, the direct effect of a progressive reform reduces relative skill supply; and by weak bias, a reduction in relative skill supply induces equalizing technical change.

The term $TE_{\theta,\tau}(T)$ captures the technical change effect induced by the component $\widetilde{TE}_{\theta,\tau}(T)$ of the labor supply response to τ. Recall from Lemma 8 that this component decreases in θ for any progressive reform (with constant labor supply elasticities). Hence, by weak bias, it induces equalizing technical change. The term $TE_{\theta,\tau}(T)$ must therefore decrease in θ. Intuitively, it captures the successive rounds of general equilibrium feedback from directed technical change to labor supply and back to technical change. The direct response of labor supply to a progressive reform τ induces equalizing technical change (see above). 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).

Finally, the slope of the term $SE_{\theta,\tau}(T)$ cannot be signed without further restrictions. The reason is that this term includes within-technology substitution effects. To sign within-technology substitution effects, however, we have imposed too little structure on the aggregate production function F so far.

3.5.4 CES Case

The CES case provides the additional structure needed to sign not only the slope of the term $SE_{\theta,\tau}(T)$ but also, more importantly, the total directed technical change effect of a progressive tax reform. In particular, when aggregate production takes the CES form and labor supply elasticities are constant across workers, any progressive tax reform induces equalizing technical change.

Corollary 9. *Fix an initial tax T and assume that F and Φ are CES as introduced in Section 3.3.4. 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 the technical change induced by tax reform τ satisfies*

$$\frac{1}{w_\theta} D_{\phi,\tau} w_\theta(T, \phi^*(T)) = \rho^{CES}(-\bar{\epsilon}^R) \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} - \rho^{CES} \int_{\underline{\theta}}^{\overline{\theta}} \frac{w_{\widetilde{\theta}} l_{\widetilde{\theta}} h_{\widetilde{\theta}}}{F(l(T), \phi^*(T))} (-\bar{\epsilon}^R) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1 - T'(y_{\widetilde{\theta}}(T))} \, d\widetilde{\theta} \tag{3.29}$$

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

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

Hence, any progressive tax reform induces technical change that reduces all skill premia.

Proof. See Appendix B.1.3. □

Inversely, regressive reforms induce skill-biased technical change under the conditions of Corollary 9.

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 3.7.

Second, directed technical change effects have implications for the welfare assessment of progressive and regressive tax reforms. I investigate these implications in Appendix B.3.2.

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

3.6 Optimal Taxes

To characterize optimal taxes, it is convenient to use the following notation. For a function $x : (\theta, z) \mapsto x_\theta(z)$ (e.g., wages or labor inputs) that depends on θ and potentially 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

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

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

3.6.1 Welfare

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 the derivative D_{u_θ} is defined analogously to the definition of D_{l_θ} in Section 3.3.2. Let the average marginal welfare weight of all workers above a given type θ be denoted by

$$\widetilde{g}_\theta := \frac{1}{1 - H_\theta} \int_{\underline{\theta}}^{\overline{\theta}} g_{\widetilde{\theta}} h_{\widetilde{\theta}} \, d\widetilde{\theta} .$$

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

Assumption 5 ensures that redistributing consumption from workers with high utility to workers with low utility improves welfare.

3.6.2 Optimal Tax Formula

To derive optimal tax rates, I follow the mechanism design approach to optimal taxation.[24] For that, write welfare as a function of consumption and labor allocations instead of utility levels:

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

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

$$\int_{\underline{\theta}}^{\overline{\theta}} c_\theta h_\theta \, d\theta = F(l, \phi^*(l)) \ . \tag{3.30}$$

Incentive compatibility requires

$$u_\theta = \max_{\widetilde{\theta}\in\Theta} \left\{ c_{\widetilde{\theta}} - v\left(\frac{w_{\widetilde{\theta}} l_{\widetilde{\theta}}}{w_\theta}\right)\right\} \quad \forall\, \theta \ .$$

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 of labor inputs extends to wages as follows.

Assumption 6. *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.[25]

[24] The alternative approach would be to use the formulas for the welfare effects of tax reforms from Appendix B.3.2 and impose that these effects are zero for all reforms at the optimum. The two approaches yield the same results.
[25] The worker density being C^1 ensures that the first part of Assumption 6 is satisfied in the CES case.

Under this restriction and with the wage function w_θ increasing in θ at the optimum, the incentive compatibility constraint is equivalent to the following conditions:

$$c'_\theta = v'(l_\theta)(w'_\theta l_\theta + w_\theta l'_\theta)\frac{1}{w_\theta} \quad \text{for almost every } \theta, \tag{3.31}$$

$$y'_\theta \geq 0 \quad \text{for almost every } \theta\ . \tag{3.32}$$

As is usual in the literature, I drop the monotonicity requirement (3.32) and study the relaxed problem of maximizing welfare subject to (3.30) and (3.31).[26]

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 the following expression for optimal marginal tax rates.

Proposition 4. *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.*

$$\frac{T'(y_\theta)}{1-T'(y_\theta)} = PE^*_\theta + TE^*_\theta + SE^*_\theta\ ,$$

where

$$PE^*_\theta = \left(1+\frac{1}{e_\theta}\right)\frac{1-N_{w_\theta}}{n_{w_\theta} w_\theta}(1-\widetilde{g}_\theta)$$

$$TE^*_\theta = \lim_{\Delta\to 0}\frac{1}{\Delta}\int_{\underline{\theta}}^{\overline{\theta}} \frac{1-T'(y_{\widetilde{\theta}})}{1-T'(y_\theta)}\frac{1-H_{\widetilde{\theta}}}{h_\theta w_\theta}(1-\widetilde{g}_{\widetilde{\theta}})y_{\widetilde{\theta}}\left.\frac{d\widehat{w}_{\widetilde{\theta}}(l,\phi^*(l+\mu\widetilde{l}_{\Delta,\theta}))}{d\mu}\right|_{\mu=0} d\widetilde{\theta}$$

$$SE^*_\theta = \lim_{\Delta\to 0}\frac{1}{\Delta}\int_{\underline{\theta}}^{\overline{\theta}} \frac{1-T'(y_{\widetilde{\theta}})}{1-T'(y_\theta)}\frac{1-H_{\widetilde{\theta}}}{h_\theta w_\theta}(1-\widetilde{g}_{\widetilde{\theta}})y_{\widetilde{\theta}}\left.\frac{d\widehat{w}_{\widetilde{\theta}}(l+\mu\widetilde{l}_{\Delta,\theta},\phi^*(l))}{d\mu}\right|_{\mu=0} d\widetilde{\theta}\ ,$$

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.

Moreover, if $\limsup_{\theta\to\overline{\theta}} l'_\theta < \infty$ and $\liminf_{\theta\to\underline{\theta}} l'_\theta > -\infty$ under the optimal tax,[27] then the following holds:

1. $TE^*_{\underline{\theta}} \leq 0$ *and, if there is strong bias (see Section 3.4.2),* $TE^*_{\underline{\theta}} + SE^*_{\underline{\theta}} \leq 0$.

[26] In all simulations of optimal taxes, I verify that the monotonicity condition (3.32) holds at the optimum.

[27] This assumption guarantees 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 in turn allows to evaluate the effects of the labor input perturbations $\widetilde{l}_{\Delta,\overline{\theta}}$ and $\widetilde{l}_{\Delta,\underline{\theta}}$ on relative labor inputs and invoke the directed technical change results of Section 3.4.

2. $TE^*_{\underline{\theta}} \geq 0$ *and, if there is strong bias,* $TE^*_{\overline{\theta}} + SE^*_{\overline{\theta}} \geq 0$.

Proof. See Appendix B.1.4. □

Proposition 4 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 second 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. Intuitively, 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 third term, SE^*_θ, stems from within-technology substitution effects. To sign this term, further structure on aggregate production is required.

Yet, whenever directed technical change dominates within-technology substitution effects (i.e., there is strong bias), the sum of the terms SE^*_θ and TE^*_θ has the same sign as TE^*_θ itself. Hence, with strong bias, the optimal marginal tax is positive at the top and negative at the bottom. This reverses the result from Stiglitz (1982) by which the optimal marginal tax at the top is negative with a general production structure that features complementarity between worker types.

The results for the marginal top tax rate extend as follows to the asymptotic marginal tax rate that is obtained when the upper tail of the income distribution has a Pareto shape.

Corollary 10. *Let the conditions of Proposition 4 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*

$$\lim_{\theta \to \overline{\theta}} \frac{1 - N_{w_\theta}}{n_{w_\theta} w_\theta} = \frac{1}{a} \ ,$$

the terms TE^*_θ and SE^*_θ satisfy

$$\lim_{\theta \to \overline{\theta}} TE^*_\theta = \overline{TE} \quad \textit{and} \quad \lim_{\theta \to \overline{\theta}} SE^*_\theta = \overline{SE} \ ,$$

and welfare weights satisfy

$$\lim_{\theta \to \overline{\theta}} g_\theta = g^{top}$$

at the optimal tax.[28]

Then, the optimal tax T *satisfies*

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

where $\overline{TE} \geq 0$ *and, if there is strong bias,* $\overline{TE} + \overline{SE} \geq 0$.

Proof. See Appendix B.1.4. □

Hence, directed technical change effects provide a force for higher marginal tax rates in the upper Pareto tail of the income distribution ($\overline{TE} \geq 0$). Moreover, under strong bias, they lead to an upwards adjustment of marginal tax rates even relative to the formula obtained in a setting with exogenous wages (e.g. Diamond, 1998).

3.6.3 CES Case

In the CES case, the optimal tax formula takes the following, particularly transparent form.[29]

Proposition 5. *Suppose the conditions of Proposition 4 are satisfied and* F *and* Φ *take the CES form introduced in Section 3.3.4. Then, at every type* θ*, optimal marginal tax rates satisfy the following conditions:*

$$\frac{T'(y_\theta)}{1-T'(y_\theta)} = \left(1+\frac{1}{e_\theta}\right)\frac{1-B_{\beta_\theta}}{b_{\beta_\theta}\beta_\theta}(1-\widetilde{g}_\theta) + \gamma^{CES}(1-g_\theta) + \rho^{CES}(1-g_\theta) \ , \tag{3.33}$$

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 B.1.4. □

Here, the optimal marginal tax at the top is

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

[28] To guarantee existence of these limits under the optimal tax, additional structure on aggregate production is needed. The CES structure imposed by Corollary 11 below is sufficient, but clearly not necessary.

[29] See Appendix B.1.4 for the CES versions of the terms TE^*_θ and SE^*_θ.

which is positive whenever there is strong bias ($\gamma^{CES} + \rho^{CES} \geq 0$). Without strong bias, the optimal marginal top tax is negative, following the logic of Stiglitz (1982).

The optimal marginal tax at the bottom is given by

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

Since $g_{\underline{\theta}} \geq 1$, this is negative if there is strong bias and positive otherwise.

The optimal marginal tax in the Pareto tail of incomes takes the following form in the CES case.

Corollary 11. *Let the conditions of Proposition 4 be satisfied and F and Φ take the CES form introduced in Section 3.3.4. Suppose at a tax $\overline{T}$, with $\overline{T}'(y) = \tau^{top}$ for all $y \geq \widetilde{y}$ and some threshold $\widetilde{y}$, the income distribution satisfies*

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

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

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

at the optimal tax.

Then, the optimal tax T satisfies

$$\lim_{\theta \to \bar{\theta}} \frac{T'(y_\theta)}{1 - T'(y_\theta)} = \frac{1 - g^{top}}{ae} + \frac{a-1}{a}\gamma^{CES}(1 - g^{top}) + \frac{a-1}{a}\rho^{CES}(1 - g^{top}) . \tag{3.34}$$

Proof. See Appendix B.1.4. □

As predicted by Corollary 10, directed technical change effects, captured here by the technical change elasticity ρ^{CES}, increase the optimal asymptotic tax in the Pareto tail. Under strong bias ($\gamma^{CES} + \rho^{CES} \geq 0$), this even exceeds the optimal asymptotic tax from a corresponding setting with exogenous wages.

Remarkably, Proposition 5 and Corollary 11 provide expressions for optimal marginal tax rates that have closed form up to welfare weights. This enables a precise analysis of the impact of directed technical change on the optimal tax.

3.6.4 Comparison to Exogenous Technology Planner

To cleanly identify the role of directed technical change, I compare the optimal tax to the one perceived as optimal by an exogenous technology planner. The exogenous technology planner observes the economy under some initial tax $\overline{T}$, correctly infers all parameters of the economy, but mistakenly believes that technology remains fixed at its current state $\phi^*(\overline{T})$, irrespectively of the tax schedule. Formally, the exogenous technology planner bases his computation of optimal taxes on the equilibrium conditions (3.2), (3.1), (3.7), and (3.9), but replaces the equilibrium technology condition (3.8) by the "wrong" equation

$$\phi^*(l) = \phi^*(l(\overline{T})) = \operatorname*{argmax}_{\phi \in \Phi} F(l(\overline{T}), \phi) \quad \forall l \,.$$[30]

Proposition 12 in Appendix B.1.4 shows that the tax $T^{ex}_{\overline{T}}$ perceived as optimal by the exogenous technology planner satisfies the following condition:[31]

$$\frac{T^{ex\prime}_{\overline{T}}(y_\theta)}{1 - T^{ex\prime}_{\overline{T}}(y_\theta)} = \left(1 + \frac{1}{e_\theta}\right) \frac{1 - \overline{B}_{\overline{\beta}_\theta}}{\overline{b}_{\overline{\beta}_\theta} \overline{\beta}_\theta} (1 - \widetilde{g}_\theta) + \gamma^{CES}(1 - g_\theta) \,, \tag{3.35}$$

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

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

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

Comparing the exogenous technology planner's tax rates (3.35) with the optimal tax in equation (3.33), 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.

The second difference is that the optimal tax formula features the hazard ratio of β whereas the exogenous technology planner uses that of $\overline{\beta}$. 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. The function $\overline{\beta}$ instead is the exogenous technology planner's wrong

[30] Note that the set of optimal taxes computed by the exogenous technology planner for arbitrary initial taxes $\overline{T}$ strictly includes the self-confirming policy equilibrium of Rothschild and Scheuer (2013). Specifically, when setting $\overline{T}$ 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.

[31] I restrict attention to the CES case here, as this allows for sharp analytical conclusions.

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 $\overline{\beta}$ instead of β.

It can be shown that the exogenous technology planner's measure of exogenous inequality is larger than the true one (see Appendix B.1.4):

$$\frac{1-\overline{B}_{\overline{\beta}\theta}}{\overline{b}_{\overline{\beta}_\theta}\overline{\beta}_\theta} > \frac{1-B_{\beta_\theta}}{b_{\beta_\theta}\beta_\theta} \quad \forall\, \theta\ . \tag{3.36}$$

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 he mistakenly believes that the skill bias of the equilibrium technology under the initial tax $\overline{T}$ is exogenous. Since more exogenous inequality calls for higher marginal tax rates, the exogenous technology planner chooses elevated marginal tax rates everywhere.[32]

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 in equations (3.33) and (3.35) vanishes, such that the only remaining difference is the (positive) term $\rho^{CES}(1-g_{\overline{\theta}})$. 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, the exogenous technology planner computes optimal marginal tax rates according to (see Corollary 13 in Appendix B.1.4)

$$\lim_{\theta\to\overline{\theta}} \frac{T^{ex\prime}_{\overline{T}}(y_\theta)}{1-T^{ex\prime}_{\overline{T}}(y_\theta)} = \frac{1-g^{top}}{ae} + \frac{a-1}{a}\gamma^{CES}(1-g^{top})\ ,$$

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 11. Hence, conditional on the limit g^{top}, directed technical change unambiguously raises the optimal marginal tax in the Pareto tail.

[32] An 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 at a lower efficiency loss. See also Proposition 15 in Appendix B.3.3 and the subsequent discussion.

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 Pareto tail of the income distribution.

3.7 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.

3.7.1 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 $\overline{T}$ (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 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 the 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 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.2.1.

Table 3.1 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 3.1, 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). Hence, accounting for directed technical change reduces the analysis approximately to the case with exogenous wages studied extensively in the literature on optimal taxation. 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 bias (see Section 3.4). I therefore call this the strong bias case.

The conservative case is supported by all four studies in Table 3.1. Moreover, there are at least two further papers that, for different reasons, do not provide estimates that could be used to infer

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 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.2.1)	38 countries	-0.53
Morrow and Trefler (2017)	Some tertiary versus no tertiary education	Long (see description in Appendix B.2.1)	38 countries	-0.11

Table 3.1. 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.2.1.

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 (2018) 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 et al. (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 3.1 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 3.1, 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.[33] The model developed in Section 3.3.1 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 3.1 likely miss part of ρ^{CES} and therefore underestimate it.[34]

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 of technology.[35]

Labor Supply Elasticity I assume an isoelastic disutility of labor as introduced in Section 3.3.4. This necessitates calibration of the hypothetical labor supply elasticity along the linearized budget set represented by the parameter e. I choose e such that the elasticity of taxable income with respect to changes in the marginal retention rate implied by the model matches empirical estimates of this elasticity. Starting from e, the elasticity of taxable income has to account for

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

[34] A 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 3.1, 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.

[35] Note 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 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.

potential non-linearities in the tax scheme and for the equilibrium response of wages to changes in the aggregate labor supply of workers of a given type. The first adjustment is accommodated by the elasticity ϵ^R, as explained in Section 3.3.2. The second adjustment leads to the following expression for the model-implied elasticity of taxable income (see Appendix B.3.1 for a more detailed explanation of this elasticity):

$$\frac{\epsilon_\theta^R}{1-(\gamma^{CES}+\rho^{CES})\epsilon_\theta^w} = \frac{e}{1+eP_{\overline{T}}(y_\theta)-(\gamma^{CES}+\rho^{CES})(1-P_{\overline{T}}(y_\theta))e}\ .$$

The wage response to a change in a type's labor supply is again likely to differ between the short and the long run. The above expression incorporates the long-run response, as evidenced by its use of $\gamma^{CES}+\rho^{CES}$. Most reliable estimates of the elasticity of taxable income, however, measure income responses over rather short periods of time (Saez, Slemrod and Giertz, 2012). Hence, a more appropriate theoretical counterpart of these estimates is given by the expression

$$\tilde{\epsilon}_\theta^R := \frac{e}{1+eP_{\overline{T}}(y_\theta)-\gamma^{CES}(1-P_{\overline{T}}(y_\theta))e}\ ,$$

which only includes the wage effects of within-technology factor substitution. I use this expression to calibrate e given estimates of $\tilde{\epsilon}_\theta^R$. In doing so, I set $P_{\overline{T}}$ to 0, as empirical tax systems are piece-wise linear.

Saez et al. (2012) propose a value of 0.25 for $\tilde{\epsilon}_\theta^R$, while Gruber and Saez (2002) find a value of 0.57. These estimates map into values for e of 0.27 and 0.64. This range includes many other estimates from the extensive literature on labor supply elasticities (e.g. Meghir and Phillips, 2010; Chetty, 2012). I choose an intermediate value of 0.5 for my baseline calibration. The main insights are robust to other values in this range.

Initial Tax System I set the initial tax system, denoted by $\overline{T}$, as an approximation to the US income tax in 2005. I follow Heathcote et al. (2017), who show that a constant-rate-of-progressivity schedule as introduced in Section 3.3.4 provides a good approximation. 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 $\overline{T}$ calibrated, the exogenous technology parameter κ is identified by the earnings distribution under the initial tax system $\overline{T}$. 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 $[\underline{\theta}, \overline{\theta}] = [0, 1]$.

In the CES case, this assumption is insubstantial, because the cross-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 (3.2) and the wage equation (3.14). The procedure is described in more detail in Appendix B.2.2.

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

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

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 $r = 1$, such that the (income-weighted) average of optimal marginal tax rates in the conservative case (i.e., $\rho^{CES} = 0.4$) is the same as in the US 2005 tax and transfer system.[36]

3.7.2 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 3.5 and 3.6.

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.[37]

I compute the wage effects of the described reform using the expressions provided by Corollaries 9 and 12. For the exogenous technology planner, who ignores directed technical change, the effects are given by the expression in Corollary 12 alone.[38]

[36] I report results for $r = 50$ (close to Rawlsian) in Appendix B.3.5.

[37] I 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

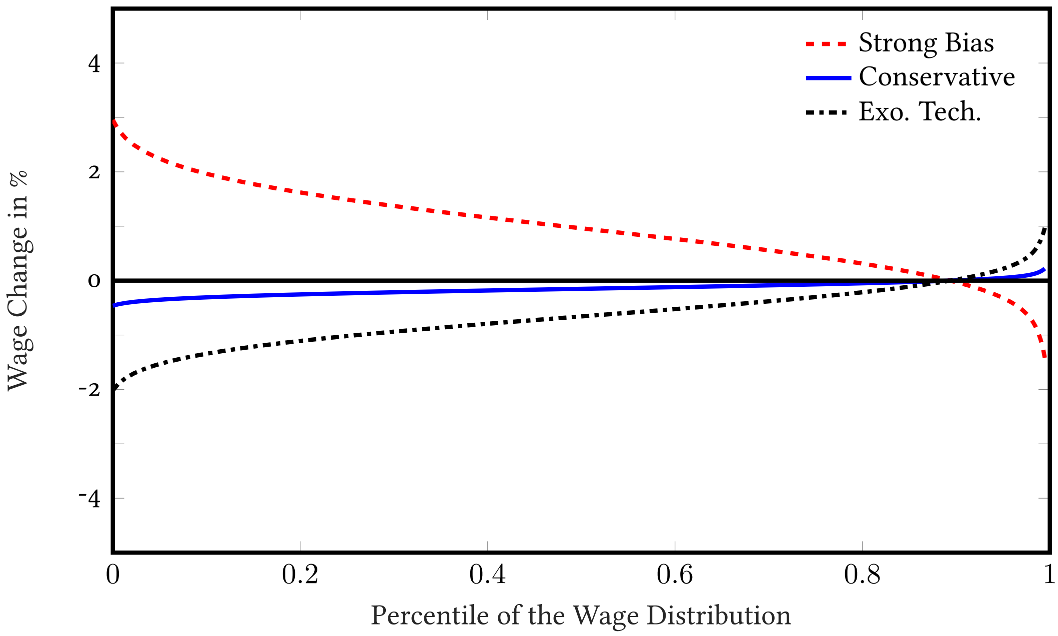

Figure 3.1. The figure displays the total wage changes in log points induced by the progressive tax reform described in the text. Wage changes are shown for workers at each percentile of the wage distribution. The red dashed line and the blue line are for the strong bias and, respectively, the conservative case of the baseline calibration described in the text. The black line indicates wage changes as predicted by the exogenous technology planner.

The results are displayed in Figure 3.1. 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 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 $e = 0.64$ (at the high end of the range discussed in Section 3.7.1), the reduction in the 90-10-percentile ratio becomes 2.6%.

unchanged.

[38]Corollaries 9 and 12 provide 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.

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 2.6% of this total of 30%, that is, 9% of the total increase in relative terms.

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 imagine 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 next step but outside the scope of the present paper, which focuses on the implications of directed technical change for the (non-parametric) design of optimal taxes.

Optimal Taxes I compute optimal marginal tax rates according to Proposition 5 and, for comparison, the marginal tax rates preferred by the exogenous technology planner as given by Proposition 12. 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 3.2 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.[39]

[39] At the very bottom of the income distribution (below the 10th percentile), optimal marginal tax rates are high in all cases. This is because, for comparability with other quantitative optimal tax studies (e.g. Mankiw, Weinzierl and Yagan, 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 5, which require strictly positive incomes

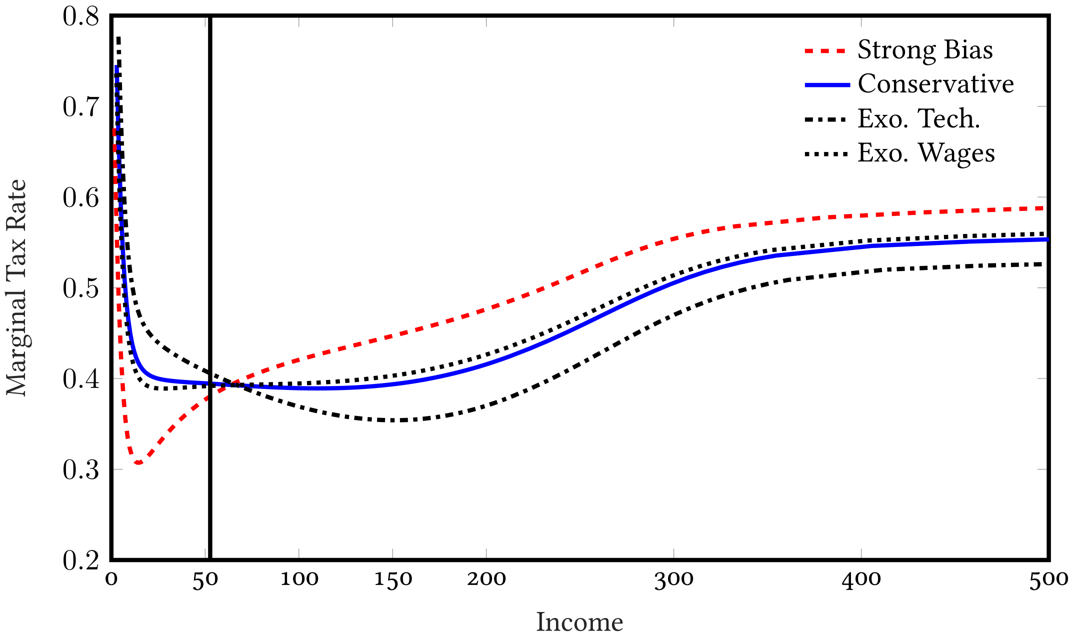

Figure 3.2. The figure displays optimal marginal tax rates by income level. The red dashed line and the blue line are for the strong bias and, respectively, the conservative case of the baseline calibration described in the text. The black lines are for the exogenous technology planner and for the case with a fully exogenous wage distribution. The vertical line indicates the US median earnings level in 2005 of about $52k.

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 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 B.3.5).

for the least skilled workers, do not apply.

Interestingly, 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.

3.8 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 rather 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 when measured by the 90-10-percentile ratio of the wage distribution. The impact of directed technical change on optimal marginal tax rates, however, is substantial. 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.

Future work may elaborate on the empirical content 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, potentially, education decisions to taxes.

4 Inefficiency and Regulation of Credence Goods Markets with Altruistic Experts

Authors: Razi Farukh, Anna Kerkhof, and Jonas Löbbing

4.1 Introduction

Market regulation is a pervasive feature of the economy in virtually all countries. In general, it appears to be more prevalent in developing countries and has consequently been associated with poor economic performance (e.g. Djankov, La Porta, de Silanes and Shleifer, 2002).

Yet, even in highly developed countries, a certain set of service sector industries exhibits a particularly high degree of regulation. In these industries, often highly qualified experts provide specialized services to consumers, who are unable to reliably assess the quality of the service provided. In its purest form, the resulting information asymmetry requires that the consumer trusts the expert to provide an appropriate service. Hence, such services have been termed credence goods (e.g. Darby and Karni, 1973; Dulleck and Kerschbamer, 2006). Existing regulation of credence goods markets often entails a combination of price controls and entry restrictions.[1] Given their potentially detrimental effect on efficiency, it is important to understand whether such regulations can be justified by the specific features of credence goods markets.[2]

Addressing this issue, we provide a novel rationale for price and entry regulation on markets for credence goods, based on considerations of economic efficiency.

In particular, we consider a setting where consumers demand a good of variable quality and cannot write contracts contingent on quality or on a signal thereof. Producers (experts, henceforth)

[1]The European Economic and Social Committee (2014) provides a comprehensive description of the various types of regulations imposed on credence goods markets in the European Union. For a detailed overview of the regulation of health care markets (arguably one of the most important credence goods markets) in OECD countries, see Paris, Devaux and Wei (2010).

[2]The professions related to credence goods markets, such as physicians or lawyers, consistently rank among the top-earning occupations in most advanced economies. Arguably, their high incomes partly reflect the regulations imposed on their markets. See, for example, Kleiner and Krueger (2013) for evidence supporting that occupational-level entry restrictions substantially increase earnings of incumbent workers. The question for justification of these regulations is therefore also relevant from a distributional perspective.

are altruistic in the sense that they value both their own monetary income and their consumers' well-being.

We impose two key assumptions. First, experts' preferences are convex in a way that makes their marginal rate of substitution between income and consumer utility decline in income. Put differently, experts' valuation of additional money relative to their consumers' utility decreases in the amount of income already earned. Second, there is a common agency structure, whereby many consumers (the principals) are served by a single expert (the agent).

In combination, these two assumptions give rise to an externality across consumers: the payment of a given consumer raises the expert's income, which in turn increases the relative importance of the other-regarding part of the expert's preferences. This improves the service quality received by all consumers served by the expert.

We study the implications of this externality in the setting that allows to expose our main results in the most transparent way. In particular, we assume that consumers are matched randomly to experts (in a many-to-one fashion) and make a take-it-or-leave-it price offer to the matched expert. Experts then decide whether to accept the offers and, in case of acceptance, covertly choose the quality of the good supplied to the respective consumer.[3]

Our first set of results shows that consumers' equilibrium price offers are inefficiently low. When making offers, consumers do not internalize the positive effect of their payment on the quality received by other consumers. Consequently, raising prices above the (unregulated) equilibrium level can make all consumers better off. Since experts are trivially better off when prices increase, introducing a fixed price or a price floor above the equilibrium price can achieve a Pareto improvement. We also show that there is no need to consider policies other than the regulation of prices in our baseline setting. Price regulation can implement all allocations that are constrained efficient in an appropriate sense.

Next, we endogenize the entry decisions of experts. We introduce a fixed cost of entry and decreasing returns in experts' technology, such that entry costs are financed out of inframarginal rents. The unregulated equilibrium is still (constrained) inefficient. With endogenous entry, however, price regulation alone does not suffice to overcome this inefficiency. Indeed, price regulation alone can lead to a Pareto deterioration: Elevated prices draw additional experts into the market until profits (net of the cost of entry) are close to zero again. Thus, the desirable effect of a price floor on profits, and thereby on experts' social behavior, vanishes. This leaves the increase in price and a congruent increase in total entry costs as the only essential allocation changes. Yet, when price regulation is combined with entry restrictions, its efficiency-enhancing

[3] In Appendix C.2, we show that our main results are unchanged in a setting where experts post prices and consumers subsequently choose between experts.

effect is re-established. A cap on the number of active experts prevents the dilution of profits through entry after prices have been raised, such that profits and the extent of experts' prosociality increase as desired.

Key to our results is the assumption that experts' preferences give rise to income effects on social behavior. We discuss evidence for this assumption in Section 4.8 at length. In a nutshell, we describe three types of evidence from existing work that support our assumption. First, results from numerous dictator games show that the level of giving strongly increases in the overall amount of money to be distributed (e.g. Engel, 2011). Second, Bartling, Valero and Weber (2019) present results from a more focused experiment, showing that increases in (experimental) income raise participants' willingness to forgo additional income to the benefit of others. Finally, various forms of correlational evidence on real-world giving behavior support the notion that giving increases with income (e.g. List, 2011).

We contribute to the existing literature by providing a novel rationale for price and entry regulation in credence goods markets. This complements Pesendorfer and Wolinsky (2003) who provide an alternative argument for price (but not entry) regulation in markets for credence goods. Other theoretical analyses of quality-related entry or price regulation, such as Atkeson, Hellwig and Ordonez (2015), deviate more strongly from the pure credence goods case and thus have different applications. Existing studies of credence goods markets with socially motivated experts (e.g. Kerschbamer, Sutter and Dulleck, 2017) and, more generally, in behavioral contract theory have not discovered the cross-consumer externality central to our results, because they either lack the common agency structure or the non-linear structure of (social) preferences.

The relation of our work to the existing literature is discussed in more detail in the next section. Section 4.3 introduces our model. In Section 4.4, we discuss a benchmark without common agency to clearly lay out the key mechanism in the model. Section 4.5 analyzes a market setting with common agency and Section 4.6 analyzes regulatory intervention. In Section 4.7, we extend the analysis to include endogenous market entry of experts and, correspondingly, study the effects of entry regulation. In Section 4.8, we describe evidence from existing work that supports our assumption that social behavior depends on income. Finally, Section 4.9 concludes.

4.2 Related Literature

In studying the regulation of credence goods markets, our work is closely related to Pesendorfer and Wolinsky (2003). They also provide a rationale for the introduction of price floors on credence goods markets. Their argument is based on a setting where consumers can consult multiple experts sequentially to learn about the service most appropriate to their needs. In this setting, an externality arises from experts' efforts to identify the need of a consumer: if other experts identify the consumer's need with high probability, the consumer can verify any given expert's recommendation with high precision by consulting a second expert. Price competition then leads any given expert to reduce price and effort, which erodes effort incentives for all other experts. A price floor stops this process and sustains high diagnostic effort by all. Our rationale for regulation is different, building on experts' social preferences. It is complementary to Pesendorfer and Wolinsky (2003) in the sense that, incorporating non-linear social preferences into their setup would give rise to the same considerations as in our analysis. In particular, this would arguably strengthen the case for a price floor and introduce benefits from entry restrictions.[4]

Other theoretical work on market regulation with the goal to promote quality deviates more strongly from the pure credence goods case analyzed here. Atkeson et al. (2015), for example, assume that consumers receive an imperfect signal of quality after their purchase, which allows for reputation building by suppliers. They also find a rationale for joint entry and price regulation, as this incentivizes sellers to undertake ex-ante investments into their quality. But again, if experts had social preferences as in our analysis, the cross-consumer externality from our setting would also arise in theirs and our implications for regulation would complement their results.

More generally, whenever the monitoring of quality is imperfect and experts have non-linear social preferences, our reasoning applies and creates a rationale for regulation. Yet, it is arguably most relevant in the pure credence goods case, where social behavior of suppliers becomes crucial because other mechanisms, such as reputation building or explicit monetary incentives, are not available.[5]

The theoretical literature on credence goods mainly focuses on relaxing the informational restrictions of the pure credence goods case in various ways and studies how this affects the ability of private contracts to overcome the remaining informational problems. Dulleck and Kerschbamer (2006) provide a useful taxonomy of informational assumptions and the associated results, giving a comprehensive overview of the corresponding studies.[6] With the exception of Pesendorfer

[4]Note that the reason for price regulation identified by Pesendorfer and Wolinsky (2003) critically depends on consumers being able to consult multiple experts. This excludes a variety of settings, in which our analysis remains applicable. These are (i) settings with a need for immediate service delivery, such as medical emergencies; (ii) situations where recommendation and execution of the service cannot be well separated; and (iii) situations where separation is feasible but the execution cannot be monitored.

[5]It is, however, important for our results that consumers have a restricted set of contracts at their disposal. Prescott and Townsend (1984) show that unrestricted private contracts achieve a constrained efficient outcome in a wide range of moral hazard settings. Their results do not apply in our case because we do not allow consumers to propose contracts contingent on experts' interaction with other consumers. For example, consumers might overcome the inefficiency in our setting by offering prices conditional on experts not accepting lower prices by other consumers. We consider this less realistic than the analyzed regulatory interventions. See Arnott, Greenwald and Stiglitz (1994) for a similar view.

[6]For examples, see Pitchik and Schotter (1987), Wolinsky (1993), and Emons (1997). An important more recent

and Wolinsky (2003) (see above), these studies do not analyze the scope for public regulation. In contrast, Mimra, Rasch and Waibel (2016) study the effects of price regulation on quality in an experiment on credence goods provision. They find that fixed prices lead to higher quality than price competition, but do not offer a theoretical explanation for their results.

Kerschbamer et al. (2017) propose social preferences as an explanation for deviations from theoretical predictions identified in experimental work by Dulleck, Kerschbamer and Sutter (2011). Yet, neither these authors nor subsequent work studies (non-linear) social preferences in a market setting with common agency. Hence, they do not discover the externality that is at the core of our results.

The same holds, more generally, for the entire literature on behavioral contract theory (see Kőszegi (2014) for a survey). Englmaier and Wambach (2010), for example, study moral hazard with inequity-averse agents, but they do not embed their analysis in a common agency framework. Therefore, they do not obtain externalities across principals.

Studies of common agency, in contrast, have identified externalities across principals in various settings (e.g. Dixit, Grossman and Helpman, 1997). Yet, these papers do not consider non-linear social preferences. Hence, their externalities are different from the one in our analysis.

4.3 Setup

We set up a model with many consumers who need a service and many experts who can provide this service. Experts covertly choose the quality of the service, which creates moral hazard. Moreover, consumer utility is not contractible, which makes the service a credence good (e.g., Dulleck and Kerschbamer, 2006).

4.3.1 Consumers

There is a continuum of consumers (or, buyers) indexed by $b \in B$. The mass of consumers $|B|$ is denoted M. Consumer b's utility is

$$u_b = v(a_b) - p_b \tag{4.1}$$

if the consumer receives a service of quality a_b and pays p_b in return. If the consumer receives no service, he gets outside utility $\underline{v}$.[7]

contribution to this line of research is Bester and Dahm (2018).

[7] We use 'he' when we speak of a consumer and 'she' when we speak of an expert.

We assume that v is C^2, with $v' > 0$ and $v'' < 0$ everywhere. For interior solutions, let $v'(a) \to 0$ as $a \to \infty$.

4.3.2 Experts

There is a finite set of experts indexed by $e \in E := \{1, 2, ..., N\}$. To reduce notation, let the number of experts equal the mass of consumers, $N = M$. Expert e earns an income of

$$y_e = \int_{B_e} [p_b - c(a_b)] \; db \; ,$$

where $B_e \subset B$ is the set of consumers served by expert e and $c(a_b)$ denotes the cost of providing a service of quality a_b. The cost function is C^2 with $c > 0$, $c' > 0$, and $c'' > 0$ everywhere. We restrict the quality variable to take positive values, such that 0 is the minimum quality an expert can provide.[8]

Note that we do not explicitly model the expert's opportunity cost of service provision. Hence, the cost function c is best thought of as including this opportunity cost. Income is then measured net of opportunity costs. If $y_e = 0$, the expert does therefore not literally earn nothing, but she earns the same amount she could earn from alternative uses of her time.

Expert e's utility is given by

$$u_e = W(y_e) + \int_{B_e} [v(a_b) - p_b] \; db \; . \tag{4.2}$$

Hence, experts care about their material payoff y_e but also about the utility of their clients. The function W is C^2 with $W' > 1$. This ensures that the expert always values her own income more than her clients' incomes at the margin. Crucially, we also assume that the marginal utility from income is decreasing, that is, $W'' < 0$ everywhere. This makes the expert's degree of selfishness contingent on her income level. If the expert earns little, she will focus on increasing her income with little regard to consumers' utility. If in contrast the expert is financially well situated, she will pay more attention to her clients' needs.

We impose two further sensible assumptions on preferences to simplify the analysis. Our main results do not depend on these assumptions. First, we transform consumers' utility function such that $v(0) - c(0) = 0$. This implies that experts do not derive moral satisfaction (i.e., utility through the non-selfish part of their preferences) by serving consumers the minimum quality 0 at the price

[8] We interpret 0 as a quality threshold such that consumers can observe whether the quality they receive exceeds 0 or not. Consumers can then condition payments on this, making experts always provide at least 0 quality. Alternatively, take 0 as a minimum service that is costless to the expert, such that she is always willing to provide this minimum.

of its cost. Second, let consumers' outside utility be small, $\underline{v} \leq 0$. This excludes uninteresting cases where consumers refuse to participate in the market.

4.3.3 Information

We assume throughout the paper that only experts themselves observe the quality of their services. Thus, consumers cannot enforce contracts that make payments contingent on quality. Moreover, we assume that consumer utility is not contractible either.[9] This precludes standard approaches to moral hazard problems.

With purely selfish preferences, these assumptions would make the case for consumers hopeless. Experts would never have an incentive to provide more than the minimum level of quality. Non-selfish experts, however, may provide higher quality services because they care for their clients. This makes our setup well-suited to study the impact of non-selfish preferences on credence goods provision in isolation from other considerations.

Note at this point that, in contrast to standard moral hazard and credence goods problems, our setting does not include a stochastic, potentially unobservable state. We can easily incorporate such a state in the analysis, but this does not add any relevant insights.

4.4 Bilateral Trade

To prepare the analysis of trading mechanisms for many consumers and many experts, consider first a bilateral setting with a single expert e and a single consumer b. The consumer is as described above. The expert, however, does not perceive the consumer as atomistic, because he is her only client. Hence the expert's utility is

$$\tilde{u}_e = W(p_b - c(a_b)) + v(a_b) - p_b$$

if she provides her service to the consumer, and $W(0)$ otherwise. In relation to the common agency setting studied in the remainder of the paper, this may best be thought of as a situation where all consumers perfectly cooperate and are replaced by a representative consumer who follows their jointly optimal strategy.

Suppose now the consumer offers a payment p_b to the expert, who can then accept or reject the offer. If the expert accepts the offer, she chooses the quality a_b and provides the service.

If the expert accepts an offer p_b, she will choose the quality a_b of her service to maximize

[9] In the jargon of the credence goods literature, we consider a setting without verifiability (of treatments) and liability (e.g., Dulleck and Kerschbamer, 2006).

utility. Expert utility is strictly concave in a_b and a_b must be non-negative by assumption, so the following Kuhn-Tucker conditions uniquely determine the optimal quality $\tilde{a}^{IC}(p_b)$:

$$\begin{aligned} \left[W'(p_b - c(\tilde{a}_b^{IC}))c'(\tilde{a}_b^{IC}) - v'(\tilde{a}_b^{IC})\right] \tilde{a}_b^{IC} &= 0 \\ W'(p_b - c(\tilde{a}_b^{IC}))c'(\tilde{a}_b^{IC}) - v'(\tilde{a}_b^{IC}) &\geq 0 \\ \tilde{a}_b^{IC} &\geq 0 \ . \end{aligned} \tag{4.3}$$

For concreteness, assume now that

$$W'(0)c'(0) \geq v'(0) \ . \tag{4.4}$$

This implies that the expert chooses the minimum quality of 0 if her income is zero. In particular, she will not incur monetary losses (relative to her outside option) to provide a quality higher than necessary.

Consider now the expert's acceptance decision. Suppose the offer is $p_b = c(0)$. If accepting this offer, the expert will choose a quality of 0 and obtain utility $W(0)$, equal to her outside option. For simplicity we assume throughout the paper that, when indifferent between two actions one of which leads to the outside option, all individuals decide against the outside option. Hence, the expert accepts the payment $c(0)$. Moreover, her utility strictly increases in p_b (recall that $W' > 1$), so she accepts all offers above $c(0)$ and rejects all offers below.

Anticipating these decisions of the expert, the consumer chooses his payment offer. In particular, he takes into account the effect of his payment on service quality. By condition (4.3), this effect is positive: a higher payment raises the expert's income, which reduces the marginal utility of income and makes the expert pay more attention to consumer utility. Thus, the consumer's offer choice is non-trivial; he may well choose a payment above $c(0)$ to receive a service of higher quality.

Let p^* denote the optimal offer for the consumer, that is,

$$p^* \in \underset{p_b \geq c(0)}{\operatorname{argmax}} \left\{ v\left(\tilde{a}_b^{IC}(p_b)\right) - p_b \right\} \ . \tag{4.5}$$

To focus on the most interesting case, we assume henceforth that v, W, and c indeed leave some scope for mutually beneficial exchange above the minimum quality 0. Formally, the minimum offer $c(0)$ (and the resulting minimum quality service) shall not maximize consumer utility:

$$c(0) \notin \underset{p_b \geq c(0)}{\operatorname{argmax}} \left\{ v\left(\tilde{a}_b^{IC}(p_b)\right) - p_b \right\} \ . \tag{4.6}$$

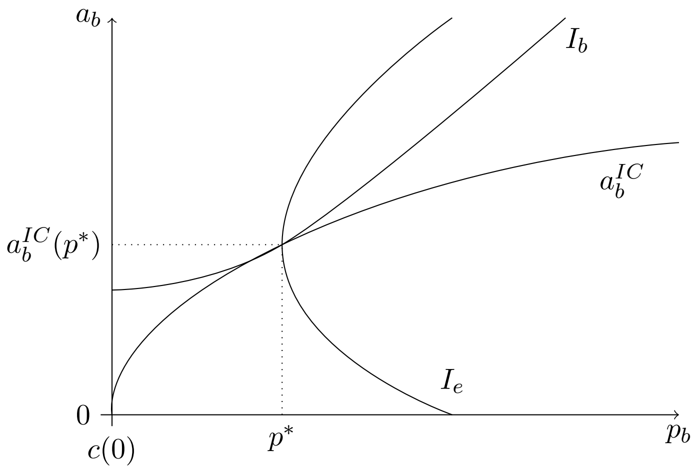

Figure 4.1. The figure displays indifference curves of the expert, I_e, and of the consumer, I_b, together with the graph of expert's quality choices $\tilde{a}_b^{IC}$. The point $(p^*, \tilde{a}_b^{IC}(p^*))$ maximizes consumer utility on the curve $\tilde{a}_b^{IC}$.

In Appendix C.1.1, we provide an exact condition showing that assumption (4.6) holds if the expert's marginal cost does not increase too quickly in quality at $a_b = 0$.

Figure 4.1 illustrates the results of the bilateral setting. The curve $\tilde{a}_b^{IC}$ marks the set of feasible allocations from the consumer's perspective. The consumer chooses the point $\left(p^*, \tilde{a}_b^{IC}(p^*)\right)$ on the curve, where his indifference curve I_b is tangent to the graph of $\tilde{a}_b^{IC}$. The expert's indifference curves I_e are such that expert utility is maximized at $\tilde{a}_b^{IC}(p_b)$ for any p_b. Hence they have slope infinity at any point $\left(p_b, \tilde{a}_b^{IC}\right)$.

4.5 Market Trade

Consider now again the setup with a finite number of experts and a continuum of consumers. As in the bilateral setting we study a trading mechanism in which consumers offer payments in exchange for the expert service and experts accept or reject.

In Appendix C.2 we analyze a mechanism where experts offer prices and consumers decide which offer to accept. This mechanism yields essentially the same outcome as the consumer-proposing mechanism studied here. The only difference is that the expert-proposing mechanism gives rise to additional equilibria (with different outcomes), which heavily rely on coordination across consumers. We argue in the appendix that these equilibria are not very plausible and provide two selection criteria, restricting consumers' ability to coordinate. Both criteria leave only the equilibrium that replicates the outcome of the consumer-proposing mechanism. To avoid these complications here, we focus directly on the consumer-proposing mechanism.

In particular, consider the following mechanism.

Stage 1 Each consumer b is matched randomly to an expert e and offers a payment p_b to the expert.[10]

Stage 2 Experts accept or reject the payments offered to them. If a consumer b's offer is rejected, he obtains outside utility $\underline{v}$. If b's offer is accepted, the accepting expert chooses a quality level a_b, and consumer b receives utility (4.1). Each expert e receives utility (4.2), where B_e is the set of consumers whose offers the expert accepted.[11]

Stages 1 and 2 describe a sequential game with complete information. We study its subgame perfect equilibria by backward induction. For that, suppose payments $\{p_b\}_{b\in B}$ and acceptance sets B_e are given. Then, experts choose quality levels a_b to maximize utility subject to the non-negativity constraint $a_b \geq 0$ for all b. Let a_b^{IC} denote the optimal quality choice of expert e for consumer $b \in B_e$. As in the bilateral setting, this quality is uniquely determined by the following Kuhn-Tucker conditions:[12]

$$\begin{aligned} \left[W'(y_e)c'(a_b^{IC}) - v'(a_b^{IC})\right] a_b^{IC} &= 0 \\ W'(y_e)c'(a_b^{IC}) - v'(a_b^{IC}) &\geq 0 \\ a_b^{IC} &\geq 0 \; . \end{aligned} \tag{4.7}$$

Before choosing quality, experts decide which offers to accept. Formally, each expert e assesses for each of her offers the marginal utility of adding the offer to her acceptance set B_e. The set B_e must therefore satisfy the following conditions:

$$W'(y_e)\left(p_b - c(a_b^{IC})\right) + v\left(a_b^{IC}\right) - p_b \begin{cases} \geq 0 & \forall\, b \in B_e \\ < 0 & \text{for all } b \text{ whose offer } e \text{ rejects.} \end{cases} \tag{4.8}$$

Using experts' quality choices, these conditions lead to a simple characterization of acceptance decisions contingent on an expert's income.

[10] We assume that for each consumer the matching probability is uniform across experts. Thus, each expert will be matched to a mass M/N of consumers.

[11] Note that consumers cannot condition their payments on the service quality they receive. This follows from our assumption that quality is hidden to consumers and final outcomes are not contractible.

[12] Expert utility is strictly concave in $\{a_b\}_{b\in B_e}$, such that the Kuhn-Tucker conditions identify a unique maximizer.

Lemma 9. *Given payment offers $\{p_b\}_{b \in B}$, any expert e's acceptance set B_e and income y_e must satisfy, for any b matched to e on stage 1,*

$$b \in B_e \quad \Leftrightarrow \quad p_b \geq \begin{cases} c(0) & \text{if } y_e \leq 0 \\ \tilde{p}(y_e) & \text{if } y_e > 0 \end{cases}$$

with $\tilde{p} : y_e \mapsto \tilde{p}(y_e)$ decreasing in y_e and $\tilde{p}(y_e) \leq c(0)$ for all $y_e > 0$.

Proof. See Appendix C.1.2. □

Lemma 9 provides an acceptance threshold for consumers' offers. Anticipating this threshold and experts' subsequent quality choices, consumers decide about their offers.

Importantly, here the quality provided by expert e does not depend on any individual payment p_b. In particular, by condition (4.7) the quality an expert provides is fully determined by her income. But since consumers are atomistic, they perceive their contribution to the expert's income as negligible. Hence, in contrast to the bilateral setting, consumers have no incentive to raise their payment above the acceptance threshold. The following proposition shows that the relevant piece of the threshold then becomes $c(0)$.

Proposition 6. *Consider the game described by stages 1 and 2. In any subgame perfect equilibrium all consumers offer $c(0)$ and receive the minimum quality, that is, $p_b = c(0)$ and $a_b = 0$ for all $b \in B$.*[13,14]

Proof. See Appendix C.1.3. □

Proposition 6 stands in stark contrast to the result from the bilateral setting. Intuitively, this discrepancy stems from an externality across buyers. If other buyers raised their payments, experts' incomes would increase and so would the service quality that any given buyer receives.

Note that the key assumption for this result is that experts' preferences over income and consumer utility are convex in a way that makes the marginal rate of substitution between the two goods decreases in income. This induces experts to care more for their consumers and provide higher quality services when their income is high.

[13] Our propositions focus on equilibrium outcomes instead of on the equilibria themselves, because there may be multiplicity in the latter. This multiplicity, however, purely arises from off-equilibrium actions.

[14] A formal complication arises from the assumption of a consumer continuum: If experts change their actions towards a measure zero of consumers, this does not affect experts' utilities. We ignore this uninteresting issue throughout the paper. Specifically, we dismiss any equilibrium in which some expert chooses a special action for a measure zero subset of consumers.

4.6 Regulation and Efficiency

The cross-buyer externality suggests to study regulation policy. We study price regulation that fixes consumers' payments at a prescribed level.[15]

In particular, consider the game described by stages 1 and 2 but with buyers' offers p_b fixed at the level $\overline{p}$. Since buyers then have no decisions left, the game collapses to experts' acceptance and quality decisions. These must again satisfy conditions (4.7) and (4.8).

From Lemma 9 we already know that experts accept all offers if the regulation $\overline{p}$ is greater or equal to $c(0)$. Otherwise, they reject all offers. We can therefore implement an allocation $\{p_b\}_{b\in B}$, $\{B_e\}_{e\in E}$, $\{a_b\}_{b\in\cup_{e\in E}B_e}$ via price regulation if and only if it satisfies the following conditions.[16]

(i) Payments are uniform across buyers, $p_b = p_{b'}$ for all $b, b' \in B$, and $p_b \geq c(0)$ for all $b \in B$.

(ii) The sets B_e have equal size, $|B_e| = 1$ for all $e \in E$, and they are disjoint, $B_e \cap B_{e'} = \emptyset$ for all $e \neq e'$.

(iii) Service quality is uniform across buyers, $a_b = a_{b'}$ for all $b, b' \in B$, and satisfies the Kuhn-Tucker conditions (4.7).

We call such allocations implementable. In an implementable allocation, consumer utility is given by

$$v\left(\overline{a}^{IC}(\overline{p})\right) - \overline{p}\ ,$$

where the quality level $\overline{a}^{IC}(\overline{p})$ follows from the Kuhn-Tucker conditions (4.7). Using the symmetry of implementable allocations implied by (i) and (ii), the Kuhn-Tucker conditions simplify to

$$\begin{aligned}\left[W'\left(\overline{p} - c(\overline{a}^{IC})\right)c'(\overline{a}^{IC}) - v'(\overline{a}^{IC})\right]\overline{a}^{IC} &= 0\\ W'\left(\overline{p} - c(\overline{a}^{IC})\right)c'(\overline{a}^{IC}) - v'(\overline{a}^{IC}) &\geq 0\\ \overline{a}^{IC} &\geq 0\ .\end{aligned}$$

The thus defined quality $\overline{a}^{IC}$ is identical to the quality $\tilde{a}^{IC}$ from the bilateral setting. Hence, consumer utility as a function of the regulated price $\overline{p}$ is identical to consumer utility as a function of the consumer's payment offer in the bilateral setting. This identity implies that the price p^* (as defined by equation (4.5)) maximizes consumer utility among all implementable allocations.

[15]If payments were restricted by a lower bound instead of fixed, consumers would set their offers at the lower bound as long as the lower bound does not fall short of the competitive level $c(0)$. Hence, a price floor yields essentially the same results as a fixed price.

[16]Via $\overline{p} < c(0)$ we can also implement the trivial allocation where $B_e = \emptyset$ for all $e \in E$. We ignore this allocation here.

Turning to experts' utility under regulation $\overline{p}$, we obtain

$$\max_{a\geq 0}\left\{W(\overline{p}-c(a))+v(a)-\overline{p}\right\} \ .$$

This is strictly increasing in $\overline{p}$. Since $p^* > c(0)$ by assumption (4.6), experts prefer the regulation p^* to the competitive equilibrium outcome (described in Proposition 6).[17] We have therefore established that price regulation at p^* Pareto-improves upon the competitive outcome.[18]

Proposition 7. *The allocation implemented by price regulation p^* (defined in equation (4.5)) Pareto-dominates the competitive equilibrium outcome described in Proposition 6.*

Intuitively, price regulation forces consumers to raise their payments as if internalizing the externality they impose on other consumers. This counteracts the inefficiency that arises in the competitive equilibrium.

Note at this point that a subsidy could not achieve such efficiency gains. A subsidy would lower experts' acceptance thresholds. Anticipating this, consumers would reduce their offers, leaving producer prices at $c(0)$. The incidence of the subsidy therefore falls completely on consumers. It thereby fails to raise experts' profits such that service quality remains unchanged.

To understand the potential of price regulation more completely, consider the set of constrained efficient allocations. This is the set of implementable allocations that are not Pareto-dominated by any other implementable allocation.

Since the regulation p^* maximizes consumer utility, the allocation induced by p^* is constrained efficient. When raising the price above p^*, experts gain and consumers lose. Hence, regulation levels $\overline{p} > p^*$ are constrained efficient as well. Any allocation implemented by $\overline{p} < p^*$ in contrast is not constrained efficient, as both consumers and experts prefer the allocation under p^*. The set of contrained efficient allocations is therefore the set of allocations implementable by a fixed price $\overline{p} \geq p^*$.[19]

Compare now the set of constrained efficient allocations to the set of fully efficient allocations. An allocation is fully efficient if and only if it is not Pareto-dominated by any other allocation. In the proof of Proposition 8 below, we show that an allocation is fully efficient if and only if

[17] We use the term competitive (equilibrium) outcome for the allocation described in Proposition 6, because it is identical to the outcome obtained under (perfect) price competition between experts in Appendix C.2.

[18] We say that an allocation Pareto-dominates another allocation, if no agent is worse off and a non-zero measure of agents is strictly better off in the first allocation.

[19] By the way we set up the analysis of price regulation, we ignore participation constraints of consumers. If we were to include such constraints, they would imply an upper bound on the regulation $\overline{p}$, beyond which consumers no longer participate. Otherwise, the results would remain unchanged.

$a_b = a^{**}$ for all consumers b, where the (fully) efficient quality a^{**} is given by

$$v'(a^{**}) = c'(a^{**}) \ .$$

Intuitively, fully efficient allocations maximize surplus, defined as $\int_B (v(a_b) - c(a_b)) \ db$. Starting from an allocation that does not maximize surplus, we can move to a surplus-maximizing allocation and redistribute the gains over experts and consumers to make everyone better off.

Inspecting the Kuhn-Tucker conditions for experts' quality choices, we find that expert e chooses the fully efficient quality a^{**} if and only if $W'(y_e) = 1$. In words, to provide fully efficient quality, experts must be indifferent regarding marginal redistribution of money between them and their consumers. Since we excluded this by assumption ($W' > 1$), we can never achieve fully efficient service quality without interfering with experts' quality choices directly. So, the sets of constrained efficient and fully efficient allocations are disjoint; price regulation never achieves full efficiency.

We summarize our findings on the structure of efficient allocations as follows.

Proposition 8. *The set of constrained efficient allocations equals the set of allocations implementable by price regulation $\overline{p} \geq p^*$, where p^* is given by equation* (4.5).

The regulation p^* maximizes consumer utility. Expert utility increases strictly in the regulation $\overline{p}$. Moreover, the sets of constrained efficient and fully efficient allocations are disjoint.

Proof. See Appendix C.1. □

Proposition 8 is illustrated by Figure 4.2. The figure focuses on symmetric allocations, represented by a common payment p and a common quality level a across consumers.

The curve $\overline{a}^{IC}$ marks all allocations implementable via price regulation. Of these, all allocations on the red (dashed) part of the curve are constrained efficient, as they have $p \geq p^*$. There is no intersection with the set of fully efficient symmetric allocations marked by the blue (dotted) line. The competitive outcome CE at $(0, c(0))$ is neither constrained nor fully efficient.

In short, raising prices up to p^* is Pareto-improving. Raising prices further benefits experts and hurts consumers.

4.7 Endogenous Entry

When price regulation raises experts' profits it may incentivize new experts to enter the market. This may dilute profits and thereby undermine the desired consequences of regulation. To address this concern we extend the analysis to a setting with endogenous entry.

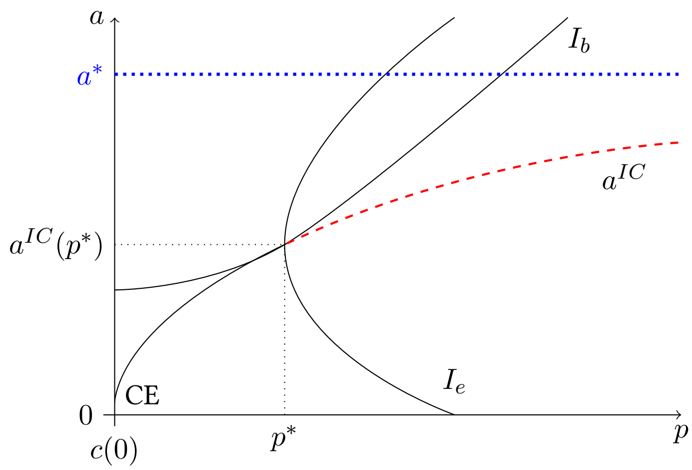

Figure 4.2. The figure displays indifference curves of experts, I_e, and of consumers, I_b, among symmetric allocations represented by a common payment p and a common service quality a. The function $\overline{a}^{IC}$ returns experts' optimal quality choice given a common payment offer p. The point CE marks the competitive equilibrium outcome from Proposition 6, the red dashed segment of $\overline{a}^{IC}$ is the set of symmetric constrained efficient allocations, and the blue dotted line is the set of symmetric fully efficient allocations.

In particular, suppose now that there is a (countably) infinite set of experts who initially decide whether to enter the market at a fixed cost $F > 0$ or not. To finance the entry cost even in a situation where prices equal marginal cost, suppose that experts operate decreasing returns to scale technologies. Formally, let the income of an expert e who entered the market be

$$\hat{y}_e = \int_{B_e} [p_b - c(a_b)] \; db - k(|B_e|) - F \; , \tag{4.9}$$

where all recurrent variables have the same meaning as before. The new cost function k is C^2 and satisfies $k(0) = 0$, $k' > 0$, and $k'' > 0$. Without loss of generality we can now impose the normalization $c(0) = 0$. The function k then measures a fixed cost per consumer served that is independent of service quality. It is convex in the mass of consumers served to capture decreasing returns to scale.[20]

Expert e's utility becomes

$$\hat{u}_e = W(\hat{y}_e) + \int_{B_e} \big(v(a_b) - p_b - v(0) + k'(|B_e|)\big) \; db \; . \tag{4.10}$$

[20]Decreasing returns to scale may for example stem from increasing difficulties to coordinate appointments with consumers, frictional interaction with a growing number of employees, or disproportional wear and tear of equipment.

Compared to the previous sections we adjust the other-regarding part of experts' utility by $|B_e|\,(-v(0)+k'(|B_e|))$. This adjustment ensures that experts do not derive immaterial benefits or losses from serving a consumer the minimal quality at marginal cost. It mirrors our assumption of $v(0)-c(0)=0$ from the previous sections. As in the previous sections, the assumption serves to simplify the analysis without substantively changing the results.

Consumers are modeled exactly as before (see section 4.3), except for that we replace the assumption $\underline{v} \leq 0$ by

$$\underline{v} \leq v(0) - k'(M)\ .$$

This again ensures that consumers' outside utility is small enough to exclude uninteresting cases where consumers refuse to participate in the market.

4.7.1 Market Trade with Endogenous Entry

We consider now the following timing of events.

Stage 1' Experts decide whether to enter the market or not. If they do not enter, they receive utility $W(0)$.

Stage 2' Denote by $E = \{1, 2, ..., N\}$ the set of experts who enter the market. Each consumer $b \in B$ is matched randomly to an expert $e \in E$ and offers a payment p_b to the expert.[21]

Stage 3' Experts accept or reject offers. If an offer p_b is rejected, consumer b receives the outside option $\underline{v}$. If p_b is accepted, the corresponding expert chooses a_b and the consumer receives utility (4.1). Finally, each expert $e \in E$ receives utility according to (4.10), where B_e is the set of consumers whose offers e accepts.

This defines a sequential game with complete information and we again study its subgame perfect equilibria by backward induction.

Given a set of active experts E, payment offers $\{p_b\}_{b\in B}$ and a matching $\{B_e\}_{e\in E}$, experts' quality choices $\hat{a}_b^{IC}$ are determined by the Kuhn-Tucker conditions (4.7) as in Section 4.5. The only difference is that income y_e is replaced by $\hat{y}_e$ as given by equation (4.9).

Moving backwards, the acceptance decisions of each expert $e \in E$ must satisfy

$$W'(\hat{y}_e)\left(p_b - c(\hat{a}_b^{IC}) - k'(|B_e|)\right) + v(\hat{a}_b^{IC}) - p_b - v(0) + k'(|B_e|) \begin{cases} \geq 0 & \forall\, b \in B_e \\ < 0 & \text{for all } b \text{ whose offer } e \text{ rejects.} \end{cases}$$

[21]Let the matching probability again be uniform, such that each expert is matched to mass M/N of consumers.

The condition computes the marginal benefit from expanding the set B_e by consumer b. If this marginal benefit is positive, the expert accepts b's offer, otherwise not. The condition leads to the following intermediate result.

Lemma 10. *Given payment offers $\{p_b\}_{b\in B}$, each active expert e's acceptance decisions B_e and income $\hat{y}_e$ must satisfy, for any consumer b matched to e on stage 2',*

$$b \in B_e \quad \Leftrightarrow \quad p_b \geq \begin{cases} k'(|B_e|) & \text{if } \hat{y}_e \leq 0 \\ \hat{p}(y_e, B_e) & \text{if } \hat{y}_e > 0 \end{cases}$$

with $\hat{p} : (\hat{y}_e, B_e) \mapsto \hat{p}(\hat{y}_e, B_e)$ decreasing in $\hat{y}_e$ and $\hat{p}(\hat{y}_e, B_e) \leq k'(|B_e|)$ for all $\hat{y}_e > 0$ and all B_e.

Proof. See Appendix C.1. □

Lemma 10 provides an acceptance threshold, which consumers anticipate when making their offers on stage 2'. Determining equilibrium offers is now complicated by inframarginal rents, which may induce positive profits. We therefore proceed with a case distinction.

Lemma 11. *Take a non-empty set of active experts E and consider the subgame after E described by stages 2' and 3'. Distinguish the following cases.*

1. *If*

$$\frac{M}{N} k'\left(\frac{M}{N}\right) - k\left(\frac{M}{N}\right) - F > 0 \, ,$$

payment offers and expert utilities must satisfy

$$p_b \leq k'\left(\frac{M}{N}\right) \quad \text{and} \quad \hat{u}_e > W(0)$$

for all $b \in B$ and $e \in E$.

2. *If*

$$\frac{M}{N} k'\left(\frac{M}{N}\right) - k\left(\frac{M}{N}\right) - F = 0 \, ,$$

payment offers and expert utilities must satisfy

$$p_b = k'\left(\frac{M}{N}\right) \quad \text{and} \quad \hat{u}_e = W(0)$$

for all $b \in B$ and $e \in E$.

3. *If*

$$\frac{M}{N}k'\left(\frac{M}{N}\right) - k\left(\frac{M}{N}\right) - F < 0\,,$$

payment offers and expert utilities must satisfy

$$p_b = k'\left(\frac{M}{N}\right) \quad \textit{and} \quad \hat{u}_e < W(0)$$

for all $b \in B$ *and* $e \in E$.

Proof. See Appendix C.1. □

Case 3 is not compatible with entry decisions on stage 1', as experts' utility falls short of their outside option. Hence, the equilibrium number of experts $\hat{N}$ must satisfy the conditions of cases 1 or 2. At $\hat{N}+1$, however, we need case 3, such that expert $\hat{N}+1$ finds it unprofitable to enter:

$$\frac{M}{\hat{N}}k'\left(\frac{M}{\hat{N}}\right) - k\left(\frac{M}{\hat{N}}\right) - F \geq 0 \tag{4.11}$$

$$\frac{M}{\hat{N}+1}k'\left(\frac{M}{\hat{N}+1}\right) - k\left(\frac{M}{\hat{N}+1}\right) - F < 0\,. \tag{4.12}$$

To resolve the cumbersome case distinction, suppose now that the mass of consumers is large, $M \to \infty$. Then, conditions (4.11) and (4.12) imply $M/\hat{N} \to m$, where m satisfies

$$mk'(m) - k(m) - F = 0\,. \tag{4.13}$$

Hence,

$$\frac{M}{\hat{N}}k'\left(\frac{M}{\hat{N}}\right) - k\left(\frac{M}{\hat{N}}\right) - F \to 0$$

as $M \to 0$. In words, when we get rid of the integer problem with finite N, we approach case 2 of Lemma 11, where experts make zero profits and payments equal marginal cost.

Proposition 9. *Consider the game described by stages 1' to 3'. Suppose $M \to \infty$. Then, in any subgame perfect equilibrium consumers' offers approach marginal cost and quality levels approach zero, that is, $p_b \to k'(m)$ and $a_b \to 0$ for all $b \in B$, where m is defined by equation* (4.13).

Proof. See Appendix C.1. □

Proposition 9 shows that for large M the equilibrium allocation with market entry approaches the competitive outcome of minimal quality and marginal cost pricing familiar from Section 4.5. The only difference is that here marginal cost is given by $k'(m)$ instead of $c(0)$.

4.7.2 Regulation with Endogenous Entry

We consider now a joint regulation of prices and entry, represented by the tuple $(\overline{p}, \overline{N})$. Such a regulation induces a game described by stages 1' to 3' with two modifications. First, only a number of $\overline{N}$ experts decides whether to enter the market on stage 1'. This caps the number of active experts at $\overline{N}$. Second, as in Section 4.6 price regulation fixes buyers' offers at $\overline{p}$.

Hence under regulation $(\overline{p}, \overline{N})$, experts decide whether to enter the market, whether to accept the fixed payment offers, and which quality to provide. Consumers have no choices. In the following we construct a regulation that Pareto-improves upon the competitive outcome of Proposition 9.

Note first that for a given number of active experts $\tilde{N}$, experts accept all offers if $\overline{p} \geq k'(M/\tilde{N})$. In such a situation, condition (4.7) for experts' quality choices simplifies to

$$\left[W'\left(\frac{M}{\tilde{N}}\overline{p} - \frac{M}{\tilde{N}}c(\hat{\overline{a}}^{IC}) - k\left(\frac{M}{\tilde{N}}\right) - F\right)c'(\hat{\overline{a}}^{IC}) - v'(\hat{\overline{a}}^{IC})\right]\hat{\overline{a}}^{IC} = 0$$
$$W'\left(\frac{M}{\tilde{N}}\overline{p} - \frac{M}{\tilde{N}}c(\hat{\overline{a}}^{IC}) - k\left(\frac{M}{\tilde{N}}\right) - F\right)c'(\hat{\overline{a}}^{IC}) - v'(\hat{\overline{a}}^{IC}) \geq 0$$
$$\hat{\overline{a}}^{IC} \geq 0\,.$$

This defines the quality $\hat{\overline{a}}^{IC}(M/\tilde{N}, \overline{p})$ as a function of the consumer to expert ratio $M/\tilde{N}$ and the price level $\overline{p}$. Consumer utility then also becomes a function of $M/\tilde{N}$ and $\overline{p}$. We denote the price that maximizes consumer utility at a given consumer to expert ratio by $\hat{p}^*(M/\tilde{N})$:

$$\hat{p}^*\left(\frac{M}{\tilde{N}}\right) \in \max_{\overline{p} \geq k'\left(\frac{M}{\tilde{N}}\right)} \left\{v\left(\hat{\overline{a}}^{IC}\left(\frac{M}{\tilde{N}}, \overline{p}\right)\right) - \overline{p}\right\}\,. \tag{4.14}$$

Assume now that for large M and at the unregulated expert number $\hat{N}$ (as given by conditions (4.11) and (4.12)), there is scope for trade above the minimum quality level of zero. Formally, if the expert to consumer ratio approaches its limit value m from the unregulated case (as given by equation (4.13)), marginal cost pricing is not collectively optimal for consumers:

$$k'(m) \notin \max_{\overline{p} \geq k'(m)} \left\{v\left(\hat{\overline{a}}^{IC}(m, \overline{p})\right) - \overline{p}\right\}\,. \tag{4.15}$$

This assumption is analogous to assumption (4.6) in the setting without entry.

As a consequence of assumption (4.15), if we can regulate entry such that the number of active experts remains the same as in the unregulated equilibrium, we can Pareto-improve upon the unregulated outcome by raising prices to $\hat{p}^*(m)$ when M is large. Proposition 10 shows that capping entry at the number of experts from the unregulated outcome, $\overline{N} = \hat{N}$, yields the desired result.[22] In addition, Proposition 10 shows that the entry-related component of the regulation is important.

Proposition 10. *Consider the regulation $(\hat{p}^*(m), \hat{N})$, where $\hat{p}^*$ is the consumer-optimal price given by equation (4.14) and $\hat{N}$ is the number of active experts in the unregulated equilibrium given by conditions (4.11) and (4.12). There exists a value $\overline{M}$ such that for all $M > \overline{M}$, the allocation implemented by the described regulation Pareto-dominates the unregulated equilibrium outcome described in Proposition 9.*

Consider in contrast the pure price regulation $(\hat{p}^*(m), \infty)$. There exists a value $\overline{M}'$ such that for all $M > \overline{M}'$, the allocation implemented by the pure price regulation is Pareto-dominated by the allocation implemented by the joint price and entry regulation described above.

Proof. See Appendix C.1. □

Proposition 10 shows that price regulation should be accompanied by entry regulation when entry is endogenous. Adding the entry regulation $\hat{N}$ to the pure price regulation $(\hat{p}^*(m), \infty)$ yields a Pareto-improvement.

To understand this result, note that the purpose of price regulation is to make experts behave less selfishly by raising their profits. But with endogenous entry, any attempt to raise profits via price regulation attracts new entrants, which counteracts the increase in profits. The desired effect on service quality is therefore mitigated. Entry regulation solves this problem by capping the number of active experts. Those who are still allowed to enter benefit from the increased prices and decide, non-selfishly, to provide higher quality services. Thus, entry regulation restores the effectiveness of price regulation.

Whether the price regulation alone already achieves a Pareto-improvement over the competitive outcome is unclear. For large M, experts' utility is approximately unaffected by pure price regulation, because entry drives down experts' utility to their outside option. For consumers the effect is ambiguous. On the one hand, increased prices reduce utility. On the other hand, although mitigated by entry, the pure price regulation can still have a positive effect on service quality. This is because the regulation raises prices above marginal cost, which has a negative effect on experts' utility through the non-selfish part of their preferences: experts feel bad because consumers pay "too much" for what they receive. This immaterial utility loss must be compensated by material gains to make experts enter the market. Hence, entry stops before the income level drops to zero. Since income is positive, service quality can be positive as well.

[22] Intuitively, raising prices above the marginal cost $k'(m)$ makes entry more attractive, such that the cap at $\hat{N}$ is binding and therefore equal to the actual number of active experts.

4.8 Does Social Behavior Depend on Income?

Our theory builds on the assumption that there are positive income effects on social behavior. In the experimental and empirical literature, there are three types of evidence supporting this assumption.

First, experimental evidence from dictator games consistently shows that individuals give more to others when their endowment increases.[23] Hence, as individuals' income in the experiment goes up, so does their willingness to forgo additional income to the benefit of others. This exactly replicates the crucial behavioral property implied by our assumption on experts' preferences. The finding that the absolute level of giving in dictator games increases in the endowment is uncontroversial in the experimental literature and therefore typically receives little attention. We view this as an indication that, at least qualitatively, our preference assumption is quite modest.

Bartling et al. (2019) question the informativeness of dictator games for whether social behavior is income-dependent or not, based on the assertion that there are strong social norms regarding the share of income to be kept in the dictator game.[24] They propose an alternative experiment, mimicking a market situation where participants decide between buying a good that inflicts externalities on others and one that does not. They find that the premium individuals are willing to pay for the externality-free good increases in their experimental income, in line with our preference assumption.

Finally, there is correlational evidence from the field. Many studies find that charitable giving significantly increases in household income (e.g. Smith, Kehoe and Cremer, 1995; List, 2011). Wiepking and Bekkers (2012) review over 50 studies showing that income and wealth have a positive effect on the level of philantrophic donations.[25] Moreover, Andreoni, Nikiforakis and Stoop (2017) demonstrate that rich households are more likely to return misdelivered envelopes with money than poor households.

Particularly insightful in our context is a study by Rasch and Waibel (2018). Using data on car repairs – i.e., expert services – in Germany, they find that a critical financial situation of a car garage is associated with a higher amount of overcharging incidences.

[23] See, for example, Carpenter, Verhoogen and Burks (2005), Chowdhury and Jeon (2014), Korenok, Millner and Razzolini (2012), and the comprehensive meta study on dictator games by Engel (2011).

[24] They argue that many individuals adhere to the norm that the money should be divided equally between dictator and recipient. Indeed, many individuals seem to follow this norm.

[25] Conducting dictator games with millionaires, Smeets, Bauer and Gneezy (2015) find that the level of giving by millionaires is "much higher than in other experiments we are aware of" (p. 10641).

4.9 Conclusion

We propose that income-dependence of social behavior creates an externality across principals in a common agency framework. This externality is most relevant in environments where the scope for monetary incentives is limited and social behavior plays a critical role. The propotypical case of such an environment is a market for credence goods.

We show that the externality creates a rationale for regulatory intervention in credence goods markets. Regulation that raises producer prices above their competitive level can achieve Pareto improvements. Examples are price floors and fixed prices. When market entry of experts is endogenous, price regulation must be accompanied by entry restrictions to seize Pareto gains.

Regarding their practical implications, our results provide a novel perspective on discussions about the dismissal of existing regulations in markets for expert services. While we believe that decisions about such deregulation must be made on a case-by-case basis, accounting for the idiosyncrasies of each market, our results should be considered as an input into these decisions.

5 Conclusion

This thesis provides three main conclusions. First, weak relative bias of technology – the result according to which a relative increase in the supply of a certain group of workers induces technical change that raises this group's relative wage – is a robust feature of neoclassical economic theory. The result is driven by the simple and powerful logic of complementarity: an increase in the relative supply of a certain input factor raises the returns to relatively complementary factors; since complementarity is, under mild conditions, a symmetric relation, the ensuing increase in the use of complementary factors in turn raises the relative return to the input that initially became more abundant.

Second, directed technical change has important implications for the design of redistributive labor income taxes. In particular, if the objective is to redistribute income from high- to low-income earners, the presence of directed technical change calls for a more progressive tax system. This insight is driven by a pre-distributive effect of directed technical change: a more progressive tax induces technical change that reduces pre-tax wage inequality. This reduction in pre-tax wage inequality makes more progressive taxes attractive for any social objective with a desire for equity.

Finally, regulation of prices and entry on credence goods markets can serve the public interest if suppliers' social behavior is income-dependent. At the heart of this result is an externality across consumers: one consumer's payment raises the supplier's income, which makes the supplier behave less selfishly and thereby improves service quality for all consumers. On a credence goods market, where informational restrictions preclude the use of explicit incentive schemes for quality provision, this externality is an important lever for regulation policy. By fixing prices above their competitive level and restricting entry, regulation can ensure high incomes for suppliers and thereby raise service quality.

Future work on technical change and inequality should focus on the following issues. First it should elaborate on the possibilities for policy to target technology more directly than via the income tax. While direct taxes on technology use have been investigated and found to be rather ineffective, it would be interesting to explore whether corporate taxes can be re-designed for this purpose. What, for example, would be the effects of a profit tax based on labor's share of firm revenue or on the spread of incomes paid by the firm?

Second, the theory of capital taxation must be revisited. With the share of income accruing to capital growing over the past 50 years, this is clearly an important topic. At the same time, there is a discrepancy between public opinion and economic theory: while public opinion often considers high capital taxes as fair, at least implicitly arguing that capital income is more indicative of luck than of strain, the arguments of economic theory for positive taxes on capital (e.g., systematic heterogeneity in time preference, or arguments based on the inverse Euler equation) are remarkably detached from such plain ideas. This discrepancy may well be indicative of a blind spot in current theory.

A Appendix to Chapter 2

A.1 Omitted Proofs and Derivations

A.1.1 Proofs

Proof of Corollary 1. The proof replicates the proof of Proposition 1 with the tools of differential calculus.

Since $\theta^*(L)$ is homogeneous of degree zero, we can restrict attention to a local increase in relative skill supply in direction of the isoquant of $F(L, \theta^*(L))$, that is, to $\mathrm{d}\, L$ such that $w_1(L, \theta^*(L))\,\mathrm{d}\, L_1 + w_2(L, \theta^*(L))\,\mathrm{d}\, L_2 = 0$ and $\mathrm{d}\, L_1 < 0$. Let $\mathrm{d}\,\theta^* := \nabla_L \theta^*(L)\,\mathrm{d}\, L$ be the direction of the response of $\theta^*(L)$ to the change $\mathrm{d}\, L$. The marginal output effect of a technical change in direction $\mathrm{d}\,\theta^*$ at $\theta^*(L)$ must increase with the local labor supply change $\mathrm{d}\, L$:

$$\nabla_L \left[\nabla_\theta F(L, \theta^*(L))\,\mathrm{d}\,\theta^*\right] \mathrm{d}\, L \geq 0. \tag{A.1}$$

But now suppose that $\mathrm{d}\,\theta^*$ has a negative effect on the skill premium:

$$\nabla_\theta \frac{w_2(L, \theta^*(L))}{w_1(L, \theta^*(L))}\,\mathrm{d}\,\theta^* < 0.$$

This implies:

$$\begin{aligned}
0 &> \nabla_\theta w_2(L, \theta^*(L))\,\mathrm{d}\,\theta^* - \nabla_\theta w_1(L, \theta^*(L))\,\mathrm{d}\,\theta^* \frac{w_2(L, \theta^*(L))}{w_1(L, \theta^*(L))} \\
&= -\frac{w_2(L, \theta^*(L))}{w_1(L, \theta^*(L))} \nabla_{L_1} \left[\nabla_\theta F(L, \theta^*(L))\right] \mathrm{d}\,\theta^* + \nabla_{L_2} \left[\nabla_\theta F(L, \theta^*(L))\right] \mathrm{d}\,\theta^* \\
&= \nabla_L \left[\nabla_\theta F(L, \theta^*(L))\right] \mathrm{d}\, L,
\end{aligned}$$

where the second line changes the order of differentiation and the last line uses that the vector $(-w_2(L, \theta^*(L))/w_1(L, \theta^*(L)), 1)$ is proportional to $\mathrm{d}\, L$. Strict negativity of the last line gives a contradiction to equation (A.1) above. □

Proof of Proposition 2. Proposition 2 treats a special case of Theorem 3. I nevertheless present a separate proof here, as it is simpler and may, as an intermediate step, facilitate reading the more general proof below.

Part 1. ($\Rightarrow$) I first show that, if there are $L \preceq^s L'$ such that $w(L, \theta^*(L)) \prec^p w(L', \theta^*(L'))$, then $\overline{F}$ cannot be quasiconcave. Let $H(L) = \{l \mid \nabla_L F(L, \theta^*(L))(l - L) = 0\}$ be the line tangent to the isoquant of F at L, holding θ fixed at $\theta^*(L)$. Since the endogenous technology wages $w(L, \theta^*(L))$ are homogeneous of degree zero in L, we can restrict attention to cases where $L' \in H(L)$. Let $l(\tau)$ parameterize the line $H(L)$ such that $l(0) = L$ and $l(1) = L'$.

Now suppose that $\overline{F}$ is quasiconcave. Then, $H(L)$ must be tangent to the (convex) upper contour set of $\overline{F}$ at L. Hence, the restriction of $\overline{F}$ to $H(L)$ must attain its maximum at L. Quasiconcavity then requires that $\overline{F}(l(\tau))$ decreases in τ. But by hypothesis, we have

$$\frac{w_2(l(1), \theta^*(l(1)))}{w_1(l(1), \theta^*(l(1)))} + \frac{\mathrm{d}\, l_1(1)/\,\mathrm{d}\,\tau}{\mathrm{d}\, l_2(1)/\,\mathrm{d}\,\tau} > \frac{w_2(l(0), \theta^*(l(0)))}{w_1(l(0), \theta^*(l(0)))} + \frac{\mathrm{d}\, l_1(1)/\,\mathrm{d}\,\tau}{\mathrm{d}\, l_2(1)/\,\mathrm{d}\,\tau} = 0,$$

where the equality follows from the construction of $l(\tau)$. Rearranging yields

$$\nabla_L F(l(1), \theta^*(l(1))) \frac{\mathrm{d}\, l(1)}{\mathrm{d}\,\tau} > 0,$$

so there exists $\tau' > 1$ such that $F(l(\tau'), \theta^*(l(1))) > F(l(1), \theta^*(l(1)))$. Finally, because $F(l(\tau), \theta^*(l(1)))$ is a lower bound of $\overline{F}(l(\tau))$ and both are equal at $\tau = 1$, we must also have that $\overline{F}(l(\tau')) > \overline{F}(l(1))$, contradicting quasiconcavity.

($\Leftarrow$) If $\overline{F}$ is not quasiconcave, it has an upper contour set that is not convex. Hence, there exists a line parameterized by $l(\tau)$ such that

$$\overline{F}(l(0)) = \overline{F}(l(1)) > \overline{F}(l(\overline{\tau}))$$

for some $\overline{\tau} \in (0, 1)$. Suppose now without loss of generality that relative skill supply increases in direction of τ, and apply the envelope theorem in Corollary 4 of Milgrom and Segal (2002) to obtain:

$$\overline{F}(l(\overline{\tau})) - \overline{F}(l(0)) = \int_0^{\overline{\tau}} \nabla_L F(l(\tau), \theta^*(l(\tau)) \frac{\mathrm{d}\, l(\tau)}{\mathrm{d}\,\tau}\, \mathrm{d}\,\tau < 0$$
$$\overline{F}(l(1)) - \overline{F}(l(\overline{\tau})) = \int_{\overline{\tau}}^1 \nabla_L F(l(\tau), \theta^*(l(\tau)) \frac{\mathrm{d}\, l(\tau)}{\mathrm{d}\,\tau}\, \mathrm{d}\,\tau > 0.$$

It follows that there exist $\tau_1 < \tau_2$ such that

$$\nabla_L F(l(\tau_1), \theta^*(l(\tau_1))) \frac{\mathrm{d}\, l(\tau_1)}{\mathrm{d}\,\tau} < 0 < \nabla_L F(l(\tau_2), \theta^*(l(\tau_2))) \frac{\mathrm{d}\, l(\tau_2)}{\mathrm{d}\,\tau}.$$

But since $\mathrm{d}\, l(\tau_1)/\,\mathrm{d}\,\tau$ is proportional to $\mathrm{d}\, l(\tau_2)/\,\mathrm{d}\,\tau$ (because $l(\tau)$ is a line), and because $\mathrm{d}\, l_2(\tau)/\,\mathrm{d}\,\tau > 0$ (since relative skill supply increases in direction of τ), the inequalities can be rearranged to yield

$$\frac{w_2(l(\tau_1),\theta^*(l(\tau_1)))}{w_1(l(\tau_1),\theta^*(l(\tau_1)))} < \frac{w_2(l(\tau_2),\theta^*(l(\tau_2)))}{w_1(l(\tau_2),\theta^*(l(\tau_2)))},$$

which establishes the first part of the theorem.

Part 2. ($\Rightarrow$) I first show that, if any increase in relative skill supply raises the skill premium, $\overline{F}$ must be quasiconvex. The proof is by contradiction and proceeds symmetrically to the proof of ($\Leftarrow$) above.

Suppose that $\overline{F}$ is not quasiconvex. Then it has a lower contour set that is not convex. Hence, there exists a line parameterized by $l(\tau)$ such that

$$\overline{F}(l(0)) = \overline{F}(l(1)) < \overline{F}(l(\overline{\tau}))$$

for some $\overline{\tau} \in (0,1)$. Suppose without loss of generality that relative skill supply increases in direction of τ, and apply the envelope theorem in Corollary 4 of Milgrom and Segal (2002) to obtain:

$$\begin{aligned}
\overline{F}(l(\overline{\tau})) - \overline{F}(l(0)) &= \int_0^{\overline{\tau}} \nabla_L F(l(\tau),\theta^*(l(\tau)))\frac{\mathrm{d}\, l(\tau)}{\mathrm{d}\,\tau}\,\mathrm{d}\,\tau > 0\\
\overline{F}(l(1)) - \overline{F}(l(\overline{\tau})) &= \int_{\overline{\tau}}^1 \nabla_L F(l(\tau),\theta^*(l(\tau)))\frac{\mathrm{d}\, l(\tau)}{\mathrm{d}\,\tau}\,\mathrm{d}\,\tau < 0.
\end{aligned}$$

It follows that there exist $\tau_1 < \tau_2$ such that

$$\nabla_L F(l(\tau_2),\theta^*(l(\tau_2)))\frac{\mathrm{d}\, l(\tau_2)}{\mathrm{d}\,\tau} < 0 < \nabla_L F(l(\tau_1),\theta^*(l(\tau_1)))\frac{\mathrm{d}\, l(\tau_1)}{\mathrm{d}\,\tau}.$$

But since $\mathrm{d}\, l(\tau_1)/\,\mathrm{d}\,\tau$ is proportional to $\mathrm{d}\, l(\tau_2)/\,\mathrm{d}\,\tau$ (because $l(\tau)$ is a line), and because $\mathrm{d}\, l_2(\tau)/\,\mathrm{d}\,\tau > 0$ (since relative skill supply increases in direction of τ), the inequalities can be rearranged to yield

$$\frac{w_2(l(\tau_2),\theta^*(l(\tau_2)))}{w_1(l(\tau_2),\theta^*(l(\tau_2)))} < \frac{w_2(l(\tau_1),\theta^*(l(\tau_1)))}{w_1(l(\tau_1),\theta^*(l(\tau_1)))},$$

which contradicts the hypothesis that any increase in relative skill supply raises the skill premium.

($\Leftarrow$) The proof is again by contradiction and proceeds symmetrically to the proof of ($\Rightarrow$) in part 1 above.

Suppose that there are $L \preceq^s L'$ such that $w(L',\theta^*(L')) \prec^p w(L,\theta^*(L))$. Let $H(L) = \{l \mid \nabla_L F(L,\theta^*(L))(l - L) = 0\}$ be the line tangent to the isoquant of F at L, holding θ fixed at

$\theta^*(L)$. Since the endogenous technology wages $w(L, \theta^*(L))$ are homogeneous of degree zero in L, we can restrict attention to cases where $L' \in H(L)$. Let $l(\tau)$ parameterize the line $H(L)$ such that $l(0) = L$ and $l(1) = L'$.

Now, since $\overline{F}$ is quasiconvex, $H(L)$ must be tangent to the (convex) lower contour set of $\overline{F}$ at L. Hence, the restriction of $\overline{F}$ to $H(L)$ must attain its minimum at L. Quasiconvexity then requires that $\overline{F}(l(\tau))$ increases in τ. But by hypothesis, we have

$$\frac{w_2(l(1), \theta^*(l(1)))}{w_1(l(1), \theta^*(l(1)))} + \frac{\mathrm{d}\, l_1(1)/\,\mathrm{d}\,\tau}{\mathrm{d}\, l_2(1)/\,\mathrm{d}\,\tau} < \frac{w_2(l(0), \theta^*(l(0)))}{w_1(l(0), \theta^*(l(0)))} + \frac{\mathrm{d}\, l_1(1)/\,\mathrm{d}\,\tau}{\mathrm{d}\, l_2(1)/\,\mathrm{d}\,\tau} = 0,$$

where the equality follows from the construction of $l(\tau)$. Rearranging yields

$$\nabla_L F(l(1), \theta^*(l(1))) \frac{\mathrm{d}\, l(1)}{\mathrm{d}\,\tau} < 0,$$

so there exists $\tau' < 1$ such that $F(l(\tau'), \theta^*(l(1))) > F(l(1), \theta^*(l(1)))$. Finally, because $F(l(\tau), \theta^*(l(1)))$ is a lower bound of $\overline{F}(l(\tau))$ and both are equal at $\tau = 1$, we must also have that $\overline{F}(l(\tau')) > \overline{F}(l(1))$, contradicting quasiconvexity. □

Proof of Theorem 1. The structure of the proof is the same as in the two skills case. First, since $\theta^*(L)$ is homogeneous of degree zero in L, we can restrict attention to labor supply changes along the exogenous technology isoquant of $F(L, \theta^*(L))$, that is, to changes from L to L' such that $F(L, \theta^*(L)) = F(L', \theta^*(L))$. This allows to construct a monotonic and differentiable path $l(\tau)$ from L to L' such that $l(0) = L$, $l(1) = L'$, and $F(l(\tau), \theta^*(L)) = F(L, \theta^*(L))$ for all $\tau \in [0, 1]$. (Monotonicity here means that each component $l_s(\tau)$ is monotonic in τ.)

Moreover, we can restrict attention to cases with $\theta^*(L) \not\approx^b \theta^*(L')$, because otherwise the statement of the theorem is trivially satisfied. In these cases, we have

$$F(l(0), \theta^*(L')) \leq F(l(0), \theta^*(L)) = F(l(1), \theta^*(L)) \leq F(l(1), \theta^*(L')), \tag{A.2}$$

with at least one of the inequalities being strict because we select the supremum of the maximizer set in equation (2.2). (If both inequalities were equalities, we would either select $\theta^*(L)$ at both $l(0)$ and $l(1)$, or $\theta^*(L')$.)

Using the mean value theorem, equation (A.2) implies that there is a $\tau' \in (0, 1)$ such that

$$\nabla_L F(l(\tau'), \theta^*(L')) \frac{\mathrm{d}\, l(\tau')}{\mathrm{d}\,\tau} > 0. \tag{A.3}$$

Let $\widetilde{s}$ denote a skill level such that $\mathrm{d}\, l_s(\tau')/\,\mathrm{d}\,\tau \geq 0$ for all $s > \widetilde{s}$ and $\mathrm{d}\, l_s(\tau')/\,\mathrm{d}\,\tau \leq 0$ for all

$s \leq \widetilde{s}$. Such a skill level exists because $l_s(\tau)$ is monotonic for each s and L' has greater relative skill supply than L. Recalling that

$$\nabla_L F(l(\tau'), \theta^*(L)) \frac{\mathrm{d}\, l(\tau')}{\mathrm{d}\, \tau} = 0,$$

we can extend equation (A.3) to

$$\left[\nabla_L F(l(\tau'), \theta^*(L')) - \frac{w_{\widetilde{s}}(l(\tau'), \theta^*(L'))}{w_{\widetilde{s}}(l(\tau'), \theta^*(L))} \nabla_L F(l(\tau'), \theta^*(L))\right] \frac{\mathrm{d}\, l(\tau')}{\mathrm{d}\, \tau} > 0. \tag{A.4}$$

The left-hand-side of this inequality is the product of two vectors with entries indexed by s. The second vector, $\mathrm{d}\, l(\tau')/\, \mathrm{d}\, \tau$, has weakly negative entries below and weakly positive entries above $\widetilde{s}$. If $\theta^*(L')$ were less skill-biased than $\theta^*(L)$, the opposite would hold for the first vector, that is, its entries are weakly positive below and weakly negative above $\widetilde{s}$, since

$$\frac{w_s(l(\tau'), \theta^*(L')}{w_s(l(\tau'), \theta^*(L))} \gtreqless \frac{w_{\widetilde{s}}(l(\tau'), \theta^*(L')}{w_{\widetilde{s}}(l(\tau'), \theta^*(L))} \quad \text{if} \quad s \lesseqgtr \widetilde{s}. \tag{A.5}$$

But this implies that the product of the two vectors is weakly negative, in contradiction to inequality (A.4). Finally, since by hypothesis we can order $\theta^*(L)$ and $\theta^*(L')$ according to their skill bias, we must have $\theta^*(L) \preceq^b \theta^*(L')$. □

Proof of Theorem 2. The proof is in large parts analogous to the proof of Theorem 1. We can again focus on L' such that $F(L, \theta^*(L)) = F(L', \theta^*(L))$, and we can again construct a path $l(\tau)$ from L to L', as in the proof of Theorem 1.

Now suppose, to derive a contradiction, that $\theta^*(L') \prec^{b'} \theta^*(L)$. This implies that there must exist a $\widetilde{\tau}$ such that

$$F(l(\widetilde{\tau}), \theta^*(L')) = F(l(\widetilde{\tau}), \theta^*(L)) = F(l(\tau), \theta^*(L)) < F(l(\tau), \theta^*(L'))$$

for all $\tau \in (\widetilde{\tau}, 1]$.

The mean value theorem then implies existence of a $\tau' \in (\widetilde{\tau}, 1)$ such that

$$\nabla_L F(l(\tau'), \theta^*(L')) \frac{\mathrm{d}\, l(\tau')}{\mathrm{d}\, \tau} > 0, \tag{A.6}$$

analogous to inequality (A.3) above. From here on, the proof follows exactly the proof of Theorem 1, starting at inequality (A.3). Note that inequality (A.5) holds here, because the choice of $\widetilde{\tau}$ guarantees that $F(l(\tau'), \theta^*(L)) \leq F(l(\tau'), \theta^*(L'))$. This, in combination with the initial supposition that $\theta^*(L') \prec^{b'} \theta^*(L)$, then implies inequality (A.5) and thus leads to a contradiction to the initial supposition. □

Proof of Part 1 of Theorem 3. The structure of the proof is the same as in the two skills case and hence follows closely Part 1 in the proof of Proposition 2.

Part 1. I first show that, if there are $L \preceq^s L'$ such that $w(L, \theta^*(L)) \prec^p w(L', \theta^*(L'))$, then $\overline{F}$ cannot be quasiconcave.

Let $H(L) = \{l \mid \nabla_L F(L, \theta^*(L))(l - L) = 0\}$ be the hyperplane tangent to the isoquant of F at L, holding θ fixed at $\theta^*(L)$. Since the endogenous technology wages $w(L, \theta^*(L))$ are homogeneous of degree zero in L, we can restrict attention to cases where $L' \in H(L)$. Let $l(\tau)$ parameterize the line through L and L', such that $l(0) = L$ and $l(1) = L'$.

Now suppose that $\overline{F}$ is quasiconcave. Then, $H(L)$ must be tangent to the (convex) upper contour set of $\overline{F}$ at L. Hence, the restriction of $\overline{F}$ to $H(L)$ must attain its maximum at L. Quasiconcavity then requires that $\overline{F}(l(\tau))$ decreases in τ. Let $\widetilde{s} \in (0,1)$ be the skill such that $\mathrm{d}\, l_s(1)/\,\mathrm{d}\,\tau > 0$ for all $s > \widetilde{s}$ and $\mathrm{d}\, l(1)/\,\mathrm{d}\,\tau \leq 0$ for $s \leq \widetilde{s}$. Such an $\widetilde{s}$ exists because L' has greater relative skill supply than L, both are on $l(\tau)$, which is tangent to the isoquant at L, and they must differ at a subset of skills of measure greater than zero because otherwise $w(L, \theta^*(L))$ and $w(L', \theta^*(L'))$ would be equal. Note that there must also exist an $\widetilde{s}' \in (0,1)$ such that $\mathrm{d}\, l_s(1)/\,\mathrm{d}\,\tau < 0$ if $s < \widetilde{s}'$. Moreover, by hypothesis, we have

$$\frac{w_s(l(1), \theta^*(l(1)))}{w_s(l(0), \theta^*(l(0)))} \gtreqless \frac{w_{\widetilde{s}}(l(1), \theta^*(l(1)))}{w_{\widetilde{s}}(l(0), \theta^*(l(0)))} \quad \text{if} \quad s \gtreqless \widetilde{s},$$

with strict inequality for a strictly positive measure of skills.[1]

Combining the information about $w(l(1), \theta^*(l(1)))$, $w(l(0), \theta^*(l(0)))$ and $\mathrm{d}\, l(1)/\,\mathrm{d}\,\tau$, we obtain

$$\left[w(l(1), \theta^*(l(1))) - \frac{w_{\widetilde{s}}(l(1), \theta^*(l(1)))}{w_{\widetilde{s}}(l(0), \theta^*(l(0)))} w(l(0), \theta^*(l(0))) \right] \frac{\mathrm{d}\, l(1)}{\mathrm{d}\,\tau} > 0,$$

because the left-hand side is the inner product of two vectors with positive (negative) entries for s above (below) $\widetilde{s}$, and these vectors are simultaneously different from zero at a subset of skills of strictly positive measure. It follows that

$$w(l(1), \theta^*(l(1))) \frac{\mathrm{d}\, l(1)}{\mathrm{d}\,\tau} > \frac{w_{\widetilde{s}}(l(1), \theta^*(l(1)))}{w_{\widetilde{s}}(l(0), \theta^*(l(0)))} w(l(0), \theta^*(l(0))) \frac{\mathrm{d}\, l(1)}{\mathrm{d}\,\tau} = 0, \tag{A.7}$$

where the equality follows from the construction of $l(\tau)$. Using that wages are identical to the

[1] Note that with a continuum of skills, the wage function $w(L, \theta) : S \to \mathbb{R}$ is determined uniquely up to a set of skills of measure zero by the Gateaux derivative of $F(L, \theta)$. Hence it is reasonable to treat wage functions that differ on a skill set of measure zero as equivalent. The notation $w(L, \theta^*(L)) \prec^p w(L', \theta^*(L'))$ is thus reserved for wage functions that differ on a strictly positive measure of skills.

L-derivative of F, inequality (A.7) yields

$$\nabla_L F(l(1), \theta^*(l(1))) \frac{\mathrm{d}\, l(1)}{\mathrm{d}\, \tau} > 0.$$

So, there exists $\tau' > 1$ such that $F(l(\tau'), \theta^*(l(1))) > F(l(1), \theta^*(l(1)))$.

Finally, because $F(l(\tau), \theta^*(l(1)))$ is a lower bound of $\overline{F}(l(\tau))$ and both are equal at $\tau = 1$, we must also have that $\overline{F}(l(\tau')) > \overline{F}(l(1))$, contradicting quasiconcavity.

Part 2. Parameterize the line along which $\overline{F}$ fails to be quasiconcave by $l(\tau)$, such that

$$\overline{F}(l(0)) = \overline{F}(l(1)) > \overline{F}(l(\overline{\tau}))$$

for some $\overline{\tau} \in (0, 1)$. Suppose now without loss of generality that relative skill supply increases in direction of τ, and apply the envelope theorem in Corollary 4 of Milgrom and Segal (2002) to obtain:

$$\begin{aligned}
\overline{F}(l(\overline{\tau})) - \overline{F}(l(0)) &= \int_0^{\overline{\tau}} \nabla_L F(l(\tau), \theta^*(l(\tau))) \frac{\mathrm{d}\, l(\tau)}{\mathrm{d}\, \tau}\ \mathrm{d}\, \tau < 0 \\
\overline{F}(l(1)) - \overline{F}(l(\overline{\tau})) &= \int_{\overline{\tau}}^{1} \nabla_L F(l(\tau), \theta^*(l(\tau))) \frac{\mathrm{d}\, l(\tau)}{\mathrm{d}\, \tau}\ \mathrm{d}\, \tau > 0.
\end{aligned}$$

It follows that there exist $\tau_1 < \tau_2$ such that

$$\nabla_L F(l(\tau_1), \theta^*(l(\tau_1))) \frac{\mathrm{d}\, l(\tau_1)}{\mathrm{d}\, \tau} < 0 < c \nabla_L F(l(\tau_2), \theta^*(l(\tau_2))) \frac{\mathrm{d}\, l(\tau_2)}{\mathrm{d}\, \tau},$$

with $c > 0$ some real number. Since $l(\tau)$ is a line, $\mathrm{d}\, l(\tau_1)/\, \mathrm{d}\, \tau$ and $\mathrm{d}\, l(\tau_2)/\, \mathrm{d}\, \tau$ are proportional. Thus, the two inequalities imply

$$[w(l(\tau_1), \theta^*(l(\tau_1))) - c w(l(\tau_2), \theta^*(l(\tau_2)))] \frac{\mathrm{d}\, l(\tau_1)}{\mathrm{d}\, \tau} < 0.$$

As in part 1 above, let $\widetilde{s}$ denote the skill such that $\mathrm{d}\, l_s(\tau_1)/\, \mathrm{d}\, \tau$ is greater (smaller) zero if s is greater (smaller) $\widetilde{s}$. Then replace the constant c to obtain

$$\left[w(l(\tau_1), \theta^*(l(\tau_1))) - \frac{w_{\widetilde{s}}(l(\tau_1), \theta^*(l(\tau_1)))}{w_{\widetilde{s}}(l(\tau_2), \theta^*(l(\tau_2)))} w(l(\tau_2), \theta^*(l(\tau_2))) \right] \frac{\mathrm{d}\, l(\tau_1)}{\mathrm{d}\, \tau} < 0. \qquad \text{(A.8)}$$

Now suppose, to derive a contradiction, that $w(l(\tau_2), \theta^*(l(\tau_2))) \preceq^p w(l(\tau_1), \theta^*(l(\tau_1)))$. This directly implies

$$\frac{w_s(l(\tau_1),\theta^*(l(\tau_1)))}{w_s(l(\tau_2),\theta^*(l(\tau_2)))} \gtreqless \frac{w_{\widetilde{s}}(l(\tau_1),\theta^*(l(\tau_1)))}{w_{\widetilde{s}}(l(\tau_2),\theta^*(l(\tau_2)))} \quad \text{if} \quad s \gtreqless \widetilde{s}.$$

It follows that the first vector in the inner product on the left-hand side of inequality (A.8) has positive (negative) entries for s above (below) $\widetilde{s}$. But by construction of $\widetilde{s}$, the same holds for the second vector. Their product must hence be positive, in contradiction to inequality (A.8). □

Proof of Part 2 of Theorem 3. This is the many skills analogue to Part 2 of the proof of Proposition 2.

Part 1. I first show that, if any increase in relative skill supply raises the skill premium, $\overline{F}$ must be quasiconvex along all lines in direction of $\preceq^s$. The proof is by contradiction.

Suppose that $\overline{F}$ is not quasiconvex along some line in direction of $\preceq^s$. Let $l(\tau)$ parameterize this line such that

$$\overline{F}(l(0)) = \overline{F}(l(1)) < \overline{F}(l(\overline{\tau}))$$

for some $\overline{\tau} \in (0,1)$. Suppose now without loss of generality that relative skill supply increases in direction of τ, and apply the envelope theorem in Corollary 4 of Milgrom and Segal (2002) to obtain:

$$\overline{F}(l(\overline{\tau})) - \overline{F}(l(0)) = \int_0^{\overline{\tau}} \nabla_L F(l(\tau),\theta^*(l(\tau)) \frac{\mathrm{d}\, l(\tau)}{\mathrm{d}\,\tau}\ \mathrm{d}\,\tau > 0$$
$$\overline{F}(l(1)) - \overline{F}(l(\overline{\tau})) = \int_{\overline{\tau}}^{1} \nabla_L F(l(\tau),\theta^*(l(\tau)) \frac{\mathrm{d}\, l(\tau)}{\mathrm{d}\,\tau}\ \mathrm{d}\,\tau < 0.$$

It follows that there exist $\tau_1 < \tau_2$ such that

$$\nabla_L F(l(\tau_1),\theta^*(l(\tau_1)))\frac{\mathrm{d}\, l(\tau_1)}{\mathrm{d}\,\tau} > 0 > c\nabla_L F(l(\tau_2),\theta^*(l(\tau_2)))\frac{\mathrm{d}\, l(\tau_2)}{\mathrm{d}\,\tau},$$

with $c > 0$ some real number. Since $l(\tau)$ is a line, $\mathrm{d}\, l(\tau_1)/\,\mathrm{d}\,\tau$ and $\mathrm{d}\, l(\tau_2)/\,\mathrm{d}\,\tau$ are proportional. Thus, the two inequalities imply

$$[w(l(\tau_1),\theta^*(l(\tau_1))) - cw(l(\tau_2),\theta^*(l(\tau_2)))]\,\frac{\mathrm{d}\, l(\tau_1)}{\mathrm{d}\,\tau} > 0.$$

As in the proof of Part 1 of Theorem 3 above, let $\widetilde{s}$ denote the skill such that $\mathrm{d}\, l_s(\tau_1)/\,\mathrm{d}\,\tau$ is greater (smaller) zero if s is greater (smaller) $\widetilde{s}$. Then replace the constant c to obtain

$$\left[w(l(\tau_1),\theta^*(l(\tau_1))) - \frac{w_{\widetilde{s}}(l(\tau_1),\theta^*(l(\tau_1)))}{w_{\widetilde{s}}(l(\tau_2),\theta^*(l(\tau_2)))}w(l(\tau_2),\theta^*(l(\tau_2)))\right]\frac{\mathrm{d}\, l(\tau_1)}{\mathrm{d}\,\tau} > 0. \tag{A.9}$$

Now, by hypothesis, we have $w(l(\tau_1), \theta^*(l(\tau_1))) \preceq^p w(l(\tau_2), \theta^*(l(\tau_2)))$. This directly implies

$$\frac{w_s(l(\tau_2), \theta^*(l(\tau_2)))}{w_s(l(\tau_1), \theta^*(l(\tau_1)))} \gtreqless \frac{w_{\widetilde{s}}(l(\tau_2), \theta^*(l(\tau_2)))}{w_{\widetilde{s}}(l(\tau_1), \theta^*(l(\tau_1)))} \quad \text{if} \quad s \gtreqless \widetilde{s}.$$

It follows that the first vector in the inner product on the left-hand side of inequality (A.9) has negative (positive) entries for s above (below) $\widetilde{s}$. But by construction of $\widetilde{s}$, the opposite holds for the second vector, that is, it has positive (negative) entries for s above (below) $\widetilde{s}$. Their product must hence be negative, in contradiction to inequality (A.9). It follows that the initial assumption is false and $\overline{F}$ must be quasiconvex along all lines in direction of $\preceq^s$.

Part 2. The proof is again by contradiction. Suppose that $\overline{F}$ is quasiconvex, and that there are $L \preceq^s L'$ such that $w(L', \theta^*(L')) \prec^p w(L, \theta^*(L))$.

Let $H(L) = \{l \mid \nabla_L F(L, \theta^*(L))(l - L) = 0\}$ be the hyperplane tangent to the isoquant of F at L, holding θ fixed at $\theta^*(L)$. Since the endogenous technology wages $w(L, \theta^*(L))$ are homogeneous of degree zero in L, we can restrict attention to cases where $L' \in H(L)$. Let $l(\tau)$ parameterize the line through L and L', such that $l(0) = L$ and $l(1) = L'$.

Now, since $\overline{F}$ is quasiconvex, $H(L)$ must be tangent to the (convex) lower contour set of $\overline{F}$ at L. Hence, the restriction of $\overline{F}$ to $H(L)$ must attain its minimum at L. Quasiconvexity then requires that $\overline{F}(l(\tau))$ increases in τ. As in the first part of the proof of Part 1 of Theorem 3, let $\widetilde{s} \in (0, 1)$ be the skill such that $\mathrm{d}\, l_s(1)/\mathrm{d}\,\tau > 0$ for all $s > \widetilde{s}$ and $\mathrm{d}\, l(1)/\mathrm{d}\,\tau \leq 0$ for $s \leq \widetilde{s}$. Note that there must also exist an $\widetilde{s}' \in (0, 1)$ such that $\mathrm{d}\, l_s(1)/\mathrm{d}\,\tau < 0$ if $s < \widetilde{s}'$. Moreover, by hypothesis, we have

$$\frac{w_s(l(1), \theta^*(l(1)))}{w_s(l(0), \theta^*(l(0)))} \gtreqless \frac{w_{\widetilde{s}}(l(1), \theta^*(l(1)))}{w_{\widetilde{s}}(l(0), \theta^*(l(0)))} \quad \text{if} \quad s \lesseqgtr \widetilde{s},$$

with strict inequality for a strictly positive measure of skills (see footnote 1).

Combining the information about $w(l(1), \theta^*(l(1)))$, $w(l(0), \theta^*(l(0)))$ and $\mathrm{d}\, l(1)/\mathrm{d}\,\tau$, we obtain

$$\left[w(l(1), \theta^*(l(1))) - \frac{w_{\widetilde{s}}(l(1), \theta^*(l(1)))}{w_{\widetilde{s}}(l(0), \theta^*(l(0)))} w(l(0), \theta^*(l(0))) \right] \frac{\mathrm{d}\, l(1)}{\mathrm{d}\,\tau} < 0,$$

because the left-hand side is the inner product of two vectors, one with positive (negative) entries for s above (below) $\widetilde{s}$, the other with negative (positive) entries for s above (below) $\widetilde{s}$. The inequality is strict because the vectors are simultaneously different from zero at a subset of skills of strictly positive measure. It follows that

$$w(l(1), \theta^*(l(1))) \frac{\mathrm{d}\, l(1)}{\mathrm{d}\,\tau} < \frac{w_{\widetilde{s}}(l(1), \theta^*(l(1)))}{w_{\widetilde{s}}(l(0), \theta^*(l(0)))} w(l(0), \theta^*(l(0))) \frac{\mathrm{d}\, l(1)}{\mathrm{d}\,\tau} = 0, \tag{A.10}$$

where the equality follows from the construction of $l(\tau)$. Using that wages are identical to the

L-derivative of F, inequality (A.10) yields

$$\nabla_L F(l(1), \theta^*(l(1))) \frac{\mathrm{d}\, l(1)}{\mathrm{d}\, \tau} < 0.$$

So, there exists $\tau' < 1$ such that $F(l(\tau'), \theta^*(l(1))) > F(l(1), \theta^*(l(1)))$.

Finally, because $F(l(\tau), \theta^*(l(1)))$ is a lower bound of $\overline{F}(l(\tau))$ and both are equal at $\tau = 1$, we must also have that $\overline{F}(l(\tau')) > \overline{F}(l(1))$, contradicting quasiconvexity. □

Proof of Theorem 4. Consider L and L' as in the theorem. Since $\theta^*(L)$ is homogeneous of degree zero in L, we can restrict attention to cases where L' is on the exogenous technology isoquant of $F(L, \theta^*(L))$, that is, $F(L', \theta^*(L)) = F(L, \theta^*(L))$. Let $l(\tau)$ parameterize a differentiable and monotonic path from L to L' such that $l(0) = L$, $l(1) = L'$, and $F(l(\tau), \theta^*(L)) = F(L, \theta^*(L))$ for all $\tau \in [0, 1]$. By construction, $l_1(\tau)$ and $l_3(\tau)$ are increasing, $l_2(\tau)$ is decreasing in τ.

Now suppose, to derive a contradiction, that $\theta^*(L')$ is strictly less polarizing then $\theta^*(L)$. We then have

$$F(l(0), \theta^*(L')) \leq F(l(0), \theta^*(L)) = F(l(1), \theta^*(L)) < F(l(1), \theta^*(L')), \tag{A.11}$$

where the last inequality is strict because we select θ^* as the supremum of the maximizer set in equation (2.2) (here this means that we select the most polarizing technology from the maximizer set).

Using the mean value theorem, equation (A.11) implies that there is a $\tau' \in (0, 1)$ such that

$$\nabla_L F(l(\tau'), \theta^*(L')) \frac{\mathrm{d}\, l(\tau')}{\mathrm{d}\, \tau} > 0.$$

Replacing the derivative with wages and using that $l(\tau)$ is in the isoquant of $F(L, \theta^*(L))$, we obtain

$$w(l(\tau'), \theta^*(L')) \frac{\mathrm{d}\, l(\tau')}{\mathrm{d}\, \tau} > c w(l(\tau'), \theta^*(L)) \frac{\mathrm{d}\, l(\tau')}{\mathrm{d}\, \tau} = 0,$$

for any constant $c > 0$. Rearranging yields

$$\left[w(l(\tau'), \theta^*(L')) - c w(l(\tau'), \theta^*(L))\right] \frac{\mathrm{d}\, l(\tau')}{\mathrm{d}\, \tau} > 0. \tag{A.12}$$

But now set c equal to $w_2(l(\tau'), \theta^*(L'))/w_2(l(\tau'), \theta^*(L))$. Inequality (A.12) then reduces to

$$\left[w_1(l(\tau'), \theta^*(L')) - \frac{w_2(l(\tau'), \theta^*(L'))}{w_2(l(\tau'), \theta^*(L))} w_1(l(\tau'), \theta^*(L))\right] \frac{\mathrm{d}\, l_1(\tau')}{\mathrm{d}\, \tau} + \left[w_3(l(\tau'), \theta^*(L')) - \frac{w_2(l(\tau'), \theta^*(L'))}{w_2(l(\tau'), \theta^*(L))} w_3(l(\tau'), \theta^*(L))\right] \frac{\mathrm{d}\, l_3(\tau')}{\mathrm{d}\, \tau} > 0.$$

But under the initial assumption that $\theta^*(L')$ is less polarizing than $\theta^*(L)$, both expressions in brackets are negative while the derivatives of $l_1(\tau)$ and $l_3(\tau)$ are both positive. Hence the total expression on the left-hand side must be negative, a contradiction. □

Proof of Lemma 1. First, suppose there exist $x < x'$ such that $L_{s,x} > 0$ and $K_{x'} > 0$. This requires that the cost per efficiency unit of capital is greater (smaller) than that of labor type s in task x (x'), that is,

$$\frac{w_s}{\gamma(s,x)} \leq \frac{r}{\alpha(x)} \quad \text{and} \quad \frac{w_s}{\gamma(s,x')} \geq \frac{r}{\alpha(x')}.$$

But this implies

$$\frac{\gamma(s,x')}{\gamma(s,x)} \leq \frac{\alpha(x')}{\alpha(x)},$$

which contradicts the assumed pattern of comparative advantage between labor and capital. Therefore, there exists $\widetilde{x} \in X$ such that (i) $K_x > 0$ only if $x \leq \widetilde{x}$ and (ii) $L_{s,x} > 0$ only if $x \geq \widetilde{x}$ for all s. Moreover, it is obvious that $K_x > 0$ for all $x < \widetilde{x}$, as otherwise a task would not be produced at all, increasing its relative price arbitrarily and hence violating task firms' profit maximization conditions.

Second, conditional on $\widetilde{x}$, the assignment of labor to tasks in $[\widetilde{x}, \overline{x}]$ is the same as in a model without capital and with a task set of $[\widetilde{x}, \overline{x}]$. Such a model is analyzed by Costinot and Vogel (2010), whose Lemma 1 establishes existence of a continuous and strictly increasing matching function m as proposed in the lemma.

It remains to argue why the threshold task $\widetilde{x}$ is performed by labor and not by capital. But this question turns out to be irrelevant, as we have defined an equilibrium of the model in terms of distributions of labor and capital over tasks, and a density corresponding to a distribution is only unique up to a set of measure zero. This means that whether we let $\widetilde{x}$ be performed by labor or by capital, the distributions of labor and capital over tasks, and hence the equilibrium itself, do not change. So, we can always represent the equilibrium distributions by capital and labor densities such that $K_{\widetilde{x}} = 0$ and $L_{\underline{s},\widetilde{x}} > 0$. □

Proof of Lemma 2. Consider first existence and uniqueness of the exogenous technology equilibrium. For any $\widetilde{x} \in (\underline{x}, \overline{x})$, the assignment of labor to the task set $[\widetilde{x}, \overline{x}]$ is equivalent to the labor assignment in Costinot and Vogel (2010), whose Lemma 1 establishes existence of a unique assignment function m as required by the equilibrium definition. Moreover, given m, the assignment of capital to $[\underline{x}, \widetilde{x})$ is clearly uniquely determined by the requirement that all marginal products $\alpha(x)\partial Y/\partial Y_x$ must equal r. Then, capital and labor assignment together uniquely determine task quantities, task prices, and wages via conditions (E1), (E2), and (E4).

Consider now the three properties of the exogenous technology equilibrium proposed by the lemma.

1. Consider $L' = \lambda L$. Let K_x denote the equilibrium capital density under L, let $K'_x = \lambda K_x$, and analogously for $Y'_x = \lambda Y_x$. It is then easy to check that K'_x and Y'_x, with all other equilibrium objects unchanged, form an equilibrium under the new labor supply L'. This is because final good and task production are linear homogeneous, such that scaling all inputs by a common factor does not change prices. Linear homogeneity in production also implies that final good production Y changes by the factor λ in the new equilibrium. Since aggregate capital K changes by λ as well, this must also hold for aggregate net production $Y - rK$.

2. Since all markets are competitive, the equality of wages and marginal products of labor follows from standard Walrasian equilibrium arguments.

3. According to equation (2.8), relative wages are fully determined by the matching function. Since the matching function is determined equivalently as in Costinot and Vogel (2010), an increase in $\widetilde{x}$ here has the same effects on relative wages as in increase in the lower bound of the task set in Costinot and Vogel (2010). Their Lemma 5 says that such an increase raises all skill premia.

Finally, consider the endogenous technology equilibrium. Again since all markets are competitive, standard reasoning along the lines of the first welfare theorem implies that the automation threshold $\widetilde{x}^*$ satisfies

$$\widetilde{x}^*(L) \in \operatorname*{argmax}_{\widetilde{x} \in X} F(L, \widetilde{x})$$

in any endogenous technology equilibrium (otherwise task producers could choose a different $\widetilde{x}$ and earn positive profits thereby). Moreover, any such $\widetilde{x}$ must satisfy condition (E6) and hence forms an endogenous technology equilibrium (when combined with the corresponding exogenous technology equilibrium). Existence then follows from the fact that $F(L, \widetilde{x})$ is continuous in $\widetilde{x}$ and X is compact.[2]

[2] To see that $F(L, \widetilde{x})$ is continuous in $\widetilde{x}$, note that (E1) and the final good production function (equation (2.7))

For uniqueness, suppose that there are two equilibrium technologies $\widetilde{x}_1^* < \widetilde{x}_2^*$ at some labor supply L. Then, using the assumption about comparative advantage between capital and labor (Assumption 3), (E6) implies that the least-skilled worker earns less under $\widetilde{x}_1^*$ than under $\widetilde{x}_2^*$, that is,

$$w_{\underline{s}}(L, \widetilde{x}_1^*) < w_{\underline{s}}(L, \widetilde{x}_2^*). \tag{A.13}$$

On the other hand, we know from point 3 of Lemma 2 that skill premia must also be smaller under $\widetilde{x}_1^*$, that is, $w(L, \widetilde{x}_1^*) \preceq^p w(L, \widetilde{x}_2^*)$. Finally, since $F(L, \widetilde{x}_1^*) = F(L, \widetilde{x}_2^*)$ and F is linear homogeneous in L, we have

$$\int_S w_s(L, \widetilde{x}_1^*) \,\mathrm{d}\, s = \int_S w_s(L, \widetilde{x}_2^*) \,\mathrm{d}\, s.$$

In combination with the fact that skill premia are greater under $\widetilde{x}_2^*$, this requires that the least-skilled worker earns less under $\widetilde{x}_2^*$, a contradiction to inequality (A.13). □

Proof of Lemma 3. I show that for any upper contour set of $\overline{F}$ there exists a supporting hyperplane through any point on the boundary of the upper contour set. This implies convexity of the upper contour sets and hence quasiconcavity of $\overline{F}$.

Take any L and let w^* and p_x^* denote equilibrium wages and task prices at L. Consider the hyperplane given by $\{l \mid lw^* = Lw^*\}$. Now suppose, to derive a contradiction, that there exists an L' such that $L'w^* = Lw^*$ (L' is on the mentioned hyperplane) and $\overline{F}(L') > \overline{F}(L)$ (L' is in the interior of the upper contour set bounded by L). Let Y_x^{**} and m^{**} be task quantities and matching function in equilibrium at L'. We must have that

$$\int_X Y_x^{**} p_x^* \,\mathrm{d}\, x > \int_X Y_x^* p_x^* \,\mathrm{d}\, x, \tag{A.14}$$

as otherwise final good producers would choose Y_x^{**} instead of Y_x^* in equilibrium at L (because $\overline{F}(L') > \overline{F}(L)$). But then, task producers could choose m^{**} and labor input L' to produce Y_x^{**} at labor cost $L'w^* = Lw^*$, which yields greater profits than producing Y_x^* with m^* and L (in the light of (A.14)). So, Y_x^* could not be equilibrium quantities at L. Hence, the constructed hyperplane is tangent to the upper contour set of $\overline{F}$ bounded by L. Since we can construct a supporting hyperplane in this way for any L, the proof is completed. □

Proof of Lemma 4. We can establish existence and uniqueness of the exogenous technology equi-

imply existence of an aggregate production function $F'(L, m, \{K_x\}_x, \widetilde{x})$ that is continuous in $\widetilde{x}$. Since in the exogenous technology equilibrium $\{K_x\}_x$ and m are such that they maximize aggregate production, the reduced form production function $F(L, \widetilde{x})$ is the upper envelope of $\{F'(L, m, \{K_x\}_x, \widetilde{x})\}_{m, K_x}$. As the upper envelope of a family of continuous (in $\widetilde{x}$) functions, F is itself continuous.

librium analogously to existence and uniqueness of the endogenous technology equilibrium in Lemma 2. First, note that the exogenous technology equilibrium is equivalent to the equilibrium of an otherwise identical model where the intermediate goods q_α and q_β are produced at marginal costs η_α/κ and η_β/κ and supplied under perfect competition. Call the equilibrium of this perfectly competitive model the "auxiliary equilibrium". We prove existence and uniqueness of the auxiliary equilibrium.

Suppose at first that $\widetilde{x}$ is fixed and consider conditions (E1)' to (E6)', (E8)', and (E9)'. By the same arguments as in the proof of Lemma 2, the matching function m is uniquely determined by the equilibrium conditions and $\widetilde{x}$. Similarly, the relative assignment of capital over tasks $[\underline{x}, \widetilde{x})$ is uniquely determined by the requirement that p_x is constant over these tasks (E6)'. The intermediate quantity q_β is determined by (E3)' conditional on task quantities Y_x. Solving equation (2.11) for q_β and substituting the resulting expression into final good production leads to a final good production function of the same form as in equation (2.7) in the setting without technology-embodying intermediate goods. Analogously, solving equation (2.10) for q_α and plugging the result into task production (E1)' yields a reduced form task production function that, for $x < \widetilde{x}$, only depends on capital. We can then use the derived final good production function and the reduced form task production function to uniquely determine the scale of $\{K_x\}_{x\in[\underline{x},\widetilde{x})}$ (note that the relative assignment of capital to tasks was already determined before, via condition E6'). Via the capital assignment, q_α and q_β are then determined uniquely via (E3)' and (E8)'.

Considering the determination of $\widetilde{x}$, the same arguments as in the proof of Lemma 2 imply that $\widetilde{x}$ must maximize net aggregate production in the auxiliary equilibrium, and that there is exactly one $\widetilde{x}$ consistent with the equilibrium conditions. This establishes existence and uniqueness of the auxiliary equilibrium and, by equivalence between these equilibria, of the exogenous technology equilibrium.

Consider now the properties of the modified aggregate production function $F(L, \alpha, \beta)$ and wages $w(L, \alpha, \beta)$ proposed in the lemma.

1. For linear homogeneity of F, suppose L is scaled by $\lambda > 0$. It is then easily verified that, when scaling K_x, q_α, and q_β by λ, while keeping all prices, $\widetilde{x}$, and the matching function unchanged, all equilibrium conditions are still satisfied. In this new equilibrium, all quantities are scaled by λ, so the value of F will also be scaled by λ, establishing linear homogeneity of F in L.

2. For equality of wages and the marginal product of labor in F, note that $F(L, \alpha, \beta)$ measures aggregate production in the auxiliary equilibrium. Since the auxiliary equilibrium is perfectly competitive, standard Walrasian equilibrium arguments imply equality of wages and the

marginal product of labor. Then, equivalence between auxiliary and exogenous technology equilibrium allows to transfer this result to the exogenous technology equilibrium.

3. Consider $(\alpha, \beta), (\alpha', \beta') \in \mathcal{D}$ (where $\mathcal{D}$ is the innovation possibilities frontier as given in Lemma 4) with $\alpha \leq \alpha'$, and suppose that there exists L such that $F(L, \alpha, \beta) \leq F(L, \alpha', \beta')$ and $w(L, \alpha, \beta) \npreceq^p w(L, \alpha', \beta')$. According to point 3 in Lemma 2, the proof of which holds in the present context, $w(L, \alpha, \beta) \npreceq^p w(L, \alpha', \beta')$ requires $\widetilde{x}(L, \alpha', \beta') < \widetilde{x}(L, \alpha, \beta)$. This in turn, via condition (E7)', implies

$$\frac{w_{\underline{s}}(L, \alpha, \beta)}{w_{\underline{s}}(L, \alpha', \beta')} = \frac{\alpha'}{\alpha} \frac{\gamma(\underline{s}, \widetilde{x}(L, \alpha, \beta))}{\gamma(\underline{s}, \widetilde{x}(L, \alpha', \beta'))} > 1. \tag{A.15}$$

But at the same time, the fact that $F(L, \alpha, \beta) \leq F(L, \alpha', \beta')$ and that F is linear homogeneous in L implies that

$$\int_S w_s(L, \alpha, \beta) \, \mathrm{d}\, s \leq \int_S w_s(L, \alpha', \beta') \, \mathrm{d}\, s.$$

Finally, the initial assumption that $\widetilde{x}(L, \alpha', \beta') < \widetilde{x}(L, \alpha, \beta)$ implies, by the arguments in the proof of point 3 of Lemma 2, that $w(L, \alpha', \beta') \preceq^p w(L, \alpha, \beta)$. Now, greater skill premia and a lower total wage bill under (α, β) require that the least skilled worker's wage is lower under (α, β) than under (α', β'), in contradiction to equation (A.15).

Consider now the endogenous technology equilibrium, where α and β are determined via (E10)'. Take any

$$(\alpha^*(L), \beta^*(L)) \in \underset{(\alpha,\beta)\in\mathcal{D}}{\operatorname{argmax}} F(L, \alpha, \beta),$$

and recall that

$$F(L, \alpha, \beta) = Y - rK - \frac{\eta_\alpha}{\kappa} q_\alpha - \frac{\eta_\beta}{\kappa},$$

where the quantities Y, K, q_α, and q_β take their exogenous technology equilibrium values. Moreover, let

$$F'(L, \widetilde{x}, \{K_x\}, m, q_\alpha, q_\beta, \alpha, \beta) = Y' - rK - \frac{\eta_\alpha}{\kappa} q_\alpha - \frac{\eta_\beta}{\kappa}$$

be net output at quantities $(L, \widetilde{x}, \{K_x\}, m, q_\alpha, q_\beta, \alpha, \beta)$, that is, Y' is gross output at these given quantities as derived from equation (2.9) and condition (E1)'. Since the exogenous technology values of $(\widetilde{x}, \{K_x\}, m, q_\alpha, q_\beta)$ maximize F', F is the upper envelope of the functions $F'(L, ., \alpha, \beta)$ (where the dot shall signify that the upper envelope is taken with respect to the variables $(\widetilde{x}, \{K_x\}, m, q_\alpha, q_\beta)$). Envelope arguments then imply that the technology pair $(\alpha^*(L), \beta^*(L))$

must satisfy the following Lagrange conditions:

$$\begin{aligned}
\frac{\partial F(L,\alpha^*(L),\beta^*(L))}{\partial \alpha} - \frac{\lambda}{\rho}\alpha^{\frac{1-\rho}{\rho}} &= \frac{\partial F'(L,\widetilde{x}^*,\{K_x^*\},m^*,q_\alpha^*,q_\beta^*,\alpha^*(L),\beta^*(L))}{\partial \alpha} - \frac{\lambda}{\rho}\alpha^{\frac{1-\rho}{\rho}} \\
&= q_\alpha^{*\,\kappa}\int_0^{\widetilde{x}^*} p_x K_x^{*1-\kappa}\,\mathrm{d}\,x - \frac{\lambda}{\rho}\alpha^{\frac{1-\rho}{\rho}} = 0 \\
\frac{\partial F(L,\alpha^*(L),\beta^*(L))}{\partial \beta} - \frac{\lambda}{\rho}\beta^{\frac{1-\rho}{\rho}} &= \frac{\partial F'(L,\widetilde{x}^*,\{K_x^*\},m^*,q_\alpha^*,q_\beta^*,\alpha^*(L),\beta^*(L))}{\partial \beta} - \frac{\lambda}{\rho}\beta^{\frac{1-\rho}{\rho}} \\
&= q_\beta^{*\,\kappa}\left(\int_X Y_x^{*\frac{\epsilon-1}{\epsilon}}\,\mathrm{d}\,x\right)^{\frac{(1-\kappa)\epsilon}{\epsilon-1}} - \frac{\lambda}{\rho}\beta^{\frac{1-\rho}{\rho}} = 0,
\end{aligned}$$

where the $(\widetilde{x}^*,\{K_x^*\},m^*,q_\alpha^*,q_\beta^*)$ denote the exogenous technology equilibrium quantities of the corresponding variables. Then comparison of the conditions reveals that, with $p_D = \lambda\kappa$, any (α,β) that satisfies the Lagrange conditions also satisfies the equilibrium condition (E10)' and hence forms an endogenous technology equilibrium.

Finally, note that the Lagrange conditions require

$$\frac{\partial F(L,\alpha,\beta)/\partial\alpha}{\partial F(L,\alpha,\beta)/\partial\beta} = \left(\frac{\alpha}{\beta}\right)^{\frac{1-\rho}{\rho}}.$$

When $\rho \to 0$, the right-hand side of the equation converges to a step function that is 0 for $\alpha/\beta < 1$ and jumps to infinity at $\alpha/\beta = 1$. Thus, the equation will have a unique solution when ρ is sufficiently small, as claimed in the main text below Lemma 4. □

Proof of Lemma 5. I construct a line in direction of the relative skill supply order $\preceq^s$ along which the endogenous technology function $\overline{F}$ fails to be quasiconcave. Starting from some labor supply $\overline{S}$, consider a line through $\overline{L}$ that is (i) tangent to the isoquant of $\overline{F}$ at $\overline{L}$, and (ii) such that all supply ratios within the skill sets $[\underline{s},\widetilde{s}]$ and $(\widetilde{s},\overline{s}]$ are fixed. That is, we move along this line by scaling up (or down) all supply levels above $\widetilde{s}$ by a common factor, while scaling down (or up) all supply levels below $\widetilde{s}$ by another factor.

Holding technology fixed at $(\alpha^*(\overline{L}),\beta^*(\overline{L}))$, it is easy to see that $F(L,\alpha^*(\overline{L}),\beta^*(\overline{L}))$ is linear in L on this line. In particular, suppose we scale up labor supply above $\widetilde{s}$ by the factor $\lambda > 1$ and scale down labor supply below $\widetilde{s}$ accordingly. Assume now that the assignment of skills above $\widetilde{s}$ to tasks remains constant. Then, capital adjusts such that the quantities of all tasks performed by capital and workers below $\widetilde{s}$ scale up by the factor λ as well. This holds all task prices constant, which, under the assumption of constant labor assignment, means that wages are unchanged as well. Constancy of wages in turn confirms the initial assumption of an unchanged assignment of labor with skill above $\widetilde{s}$. So, the new (exogenous technology) equilibrium features unchanged wages compared to the initial situation. Since wages correspond to marginal products in F, constancy of wages implies linearity of F in L on the constructed line. Moreover, since by construction the line is tangent to the isoquant of $\overline{F}$ (and hence of F), F is indeed constant in L on the line.

Constancy of $F(L, \alpha^*(\overline{L}), \beta^*(\overline{L}))$ on the constructed line now directly implies that the endogenous technology function $\overline{F}(L) := F(L, \alpha^*(L), \beta^*(L))$ cannot be quasiconcave along this line. This is because $F(L, \alpha^*(\overline{L}), \beta^*(\overline{L}))$ is a lower bound for $\overline{F}(L)$ that is binding at $\overline{L}$. Since $\alpha^*(L)$ changes when relative skill supply changes (see equations (2.12) and (2.13)), the lower bound does not bind at other points on the line, such that there are points $L' \preceq^s \overline{L} \preceq^s L''$, all on the constructed line, such that $\overline{F}(L'), \overline{F}(L'') > \overline{F}(\overline{L})$, which completes the proof. □

A.1.2 Complete Equilibrium Characterization for the Assignment Model with International Trade

This section provides a rigorous definition and a detailed characterization of the trade equilibrium in the two country assignment model of Section 2.5.4. The Northern economy has the same structure as the closed economy of Section 2.5.3. In particular, the final good is produced according to

$$Y^N = \int_0^1 \left({\beta_i^N}^{\frac{1}{\kappa}} q_{\beta,i}^{NN} + {\beta_i^S}^{\frac{1}{\kappa}} q_{\beta,i}^{SN} \right)^{\kappa} \mathrm{d}\, i \left[\int_{\underline{x}}^{\overline{x}} \left(Y_x^{NN} + Y_x^{SN} \right)^{\frac{\epsilon-1}{\epsilon}} \mathrm{d}\, x \right]^{\frac{\epsilon(1-\kappa)}{\epsilon-1}},$$

where Y^N denotes final good production in the North, β_i^N and β_i^S are quality levels of the technology-embodying intermediate good (β, i) in North and South, respectively, $q_{\beta,i}^{NN}$ is the quantity of this intermediate good produced and utilized in the North, $q_{\beta,i}^{SN}$ is the quantity produced in the South and utilized in the North, and analogously Y_x^{NN} (Y_x^{SN}) is the quantity of task x produced and utilized in the North (produced in the South and utilized in the North). Tasks are produced according to

$$Y_x^N = \int_0^1 \left({\alpha_i^N}^{\frac{1}{\kappa}} q_{\alpha,i}^{NN} + {\alpha_i^S}^{\frac{1}{\kappa}} q_{\alpha,i}^{SN} \right)^{\kappa} \mathrm{d}\, i\, {K_x^N}^{1-\kappa} + \gamma^N({m^N}^{-1}(x), x) L^N_{{m^N}^{-1}(x)} \frac{\mathrm{d}\, {m^N}^{-1}(x)}{\mathrm{d}\, x}.$$

Here, Y_x^N denotes production of task x in the North, α_i^N and α_i^S are quality levels of the technology-embodying intermediate good (α, i) in the North and South, respectively, $q_{\alpha,i}^{NN}$ ($q_{\alpha,i}^{SN}$) is the quantity of this good produced and utilized in the North (produced in the South and utilized in the North), K_x^N is the amount of capital employed in task x in the North, $m^N(x)$ is the skill level assigned to task x in the North, and $\gamma^N(s, x)$ is the Northern labor productivity schedule. Tasks, labor, and the final good are supplied competitively. Tasks prices in the North are denoted by p_x^N, wages by w_s^N, and the Northern final good is the numéraire.

The technology-embodying intermediate goods are produced by monopolists at marginal cost η_α and η_β, respectively, from final good. The total quantity of good (α, i) produced in the North is denoted $q_{\alpha,i}^N$, and analogously for (β, i). Prices are $p_{\alpha,i}^N$ and $p_{\beta,i}^N$. The monopolists obtain quality levels α_i^N or β_i^N at costs $p_D {\alpha_i^N}^{\frac{1}{\rho}}$ or $p_D {\beta_i^N}^{\frac{1}{\rho}}$, respectively, where p_D denotes the price for R&D resources. R&D resources are in fixed supply D in the North.

The South is symmetric to the North, with two exceptions. First, there is no R&D sector in the South. Instead, quality levels α_i^S and β_i^S are copied with some loss δ from Northern monopolists, such that $\alpha_i^S = \delta\alpha_i^N$ and $\beta_i^S = \delta\beta_i^N$ for $\delta \in (0, \kappa)$. Second, since there are no fixed R&D expenditures required to produce them, the technology-embodying intermediates $q_{\alpha,i}^S$ and $q_{\beta,i}^S$ are supplied competitively.

Final good market clearing now requires that $Y^N = Y^{NN} + Y^{NS}$, where Y^{NN} (Y^{NS}) is the amount of final good produced and consumed in the North (produced in the North and consumed in the South), and $Y^S = Y^{SS} + Y^{SN}$, where Y^{SS} and Y^{SN} are the Southern analogues of Y^{NN} and Y^{NS}. Task market clearing requires $Y_x^N = Y_x^{NN} + Y_x^{NS}$ and $Y_x^S = Y_x^{SS} + Y_x^{SN}$ for all x. The markets for technology-embodying intermediates clear if $q_{\alpha,i}^N = q_{\alpha,i}^{NN} + q_{\alpha,i}^{NS}$, $q_{\beta,i}^N = q_{\beta,i}^{NN} + q_{\beta,i}^{NS}$, and analogously for Southern intermediate good production. Finally, trade between the two countries is balanced if

$$Y^{NS} + \int_{\underline{x}}^{\overline{x}} p_x^N Y_x^{NS} \,\mathrm{d}\,x + \int_0^1 p_{\alpha,i}^N q_{\alpha,i}^{NS} \,\mathrm{d}\,i + \int_0^1 p_{\beta,i}^N q_{\beta,i}^{NS} \,\mathrm{d}\,i$$
$$= Y^{SN} + \int_{\underline{x}}^{\overline{x}} p_x^S Y_x^{SN} \,\mathrm{d}\,x + \int_0^1 p_{\alpha,i}^S q_{\alpha,i}^{SN} \,\mathrm{d}\,i + \int_0^1 p_{\beta,i}^S q_{\beta,i}^{SN} \,\mathrm{d}\,i.$$

A trade equilibrium now consists of automation thresholds $\widetilde{x}^N$ and $\widetilde{x}^S$, matching functions $m^N(s)$ and $m^S(s)$, capital assignments $\{K_x^N\}_{x\in X}$ and $\{K_x^S\}_{x\in X}$, task production $\{Y_x^N\}_{x\in X}$ and $\{Y_x^S\}_{x\in X}$, task utilization $\{Y_x^{NN}\}_{x\in X}$, $\{Y_x^{NS}\}_{x\in X}$, $\{Y_x^{SS}\}_{x\in X}$, and $\{Y_x^{SN}\}_{x\in X}$, technology intermediate production q_α^N, q_α^S, q_β^N, and q_β^S, technology intermediate utilization q_α^{NN}, q_α^{NS}, q_α^{SS}, q_α^{SN} (and analogously for β), final good quantities Y^N and Y^S, and final good consumption Y^{NN}, Y^{NS}, Y^{SS}, Y^{SN}; task prices $\{p_x^N\}_{x\in X}$ and $\{p_x^S\}_{x\in X}$, wages $\{w_s^N\}_{s\in S}$ and $\{w_s^S\}_{s\in S}$, capital prices p_c^N and p_c^S, technology intermediate prices p_α^N, p_α^S, p_β^N, and p_β^S, a price for R&D resources p_D in the North, and a price for final good in the South p^S; such that all firms maximize profits, all markets clear, and trade is balanced. Note that the definition already uses symmetry across technology intermediate producers within each country.

The remainder of the section provides a characterization of a trade equilibrium in terms of the equilibrium of an integrated economy with labor supply $L^N + \Delta L^S$, where $\Delta \in (0, 1]$ measures the difference in labor productivity between the two countries, that is, $\gamma^S(s, x) = \Delta\gamma^N(s, x)$.

We start from the observation that free trade in tasks and final good implies that the corresponding prices must be equal across countries, that is, $p_x^N = p_x^S$ for all x and $p^S = 1$ (the Northern final good is the numéraire). Since skills are assigned to those tasks in which their marginal product is greatest, and because the labor productivity difference Δ does not depend on tasks, equality of task prices implies equality of matching functions across countries. Denote the common matching function by $m^T(x)$. It follows immediately that there is also a common automation threshold $\widetilde{x}^T$ for both countries. Moreover, wages correspond to marginal products of skills in their respective tasks, so we must have $w_s^S = \Delta w_s^N$ for all skills. Finally, because final good prices are equal and the marginal cost of capital in terms of final good is r in both countries, there will be a common price of capital $p_c^T = r$.

Consider now the supply of technology-embodying intermediate goods for task production. If only Northern monopolists supplied the goods, they would again face iso-elastic demand, such that they would charge prices $p_\alpha^N = \eta_\alpha/\kappa$. This implies a price per efficiency unit of $\eta_\alpha/(\kappa\alpha^N)$. The price at which Southern producers just break even is $p_\alpha^S = \eta_\beta$, which implies a price per efficiency unit of η_α/α^S or, with $\alpha^S = \delta\alpha^N$, $\eta_\alpha/(\delta\alpha^N)$. Since it is assumed that $\delta < \kappa$, Southern producers would incur losses when producing at Northern producers' monopoly prices. Hence, in equilibrium only Northern producers produce. Moreover, they charge monopoly prices $p_\alpha^N = \eta_\alpha/\kappa$. To obtain a condition for the quantity of these goods, consider inverse demand in the North (using symmetry across α-intermediates),

$$p_\alpha^N = \kappa\alpha^N {q_\alpha^{NN}}^{\kappa-1} \int_{\underline{x}}^{\widetilde{x}^T} p_x {K_x^N}^{1-\kappa} \, \mathrm{d}\, x,$$

and in the South,

$$p_\alpha^N = \kappa\alpha^N {q_\alpha^{NS}}^{\kappa-1} \int_{\underline{x}}^{\widetilde{x}^T} p_x {K_x^S}^{1-\kappa} \, \mathrm{d}\, x.$$

Now let $q_\alpha \coloneqq q_\alpha^N + q_\alpha^S = q_\alpha^N$ denote world production, and let $s_\alpha \coloneqq q_\alpha^{NN}/q_\alpha$ be the share of world production utilized in the North. Then, linear homogeneity of task production in q_α^{NC} and K_x^C (for both countries $C = N, S$) for all tasks $x < \widetilde{x}^T$ implies that the marginal product of capital in task x will be equal in both countries if and only if $K_x^N = s_\alpha K_x$ for all $x < \widetilde{x}^T$, where $K_x \coloneqq K_x^N + K_x^S$ is the world capital stock. Note that marginal products of capital must be equal because the price of capital is the same in both countries. With this result, we can rewrite the inverse demand for the α-intermediate in the North (or, equivalently, in the South) as

$$\begin{aligned}
p_\alpha^N &= \kappa {\alpha^N} s_\alpha^{\kappa-1} q_\alpha^{\kappa-1} \int_{\underline{x}}^{\widetilde{x}^T} p_x {K_x^N}^{1-\kappa} \,\mathrm{d}\,x \\
&= \kappa {\alpha^N} q_\alpha^{\kappa-1} \int_{\underline{x}}^{\widetilde{x}^T} p_x \left(\frac{K_x^N}{s_\alpha}\right)^{1-\kappa} \mathrm{d}\,x \\
&= \kappa {\alpha^N} q_\alpha^{\kappa-1} \int_{\underline{x}}^{\widetilde{x}^T} p_x K_x^{1-\kappa} \,\mathrm{d}\,x, \qquad \text{(A.16)}
\end{aligned}$$

which has the same form as the corresponding inverse demand in the closed economy (see equation (2.10)). Profits of α-monopolists at a given α are then given by

$$\pi_{\alpha,i}(\alpha_i^N) = \max_q \left\{ \kappa \alpha_i^N q^\kappa \int_{\underline{x}}^{\widetilde{x}^T} p_x K_x^{1-\kappa} \,\mathrm{d}\,x - \eta_\alpha q - p_D {\alpha_i^N}^{1/\rho} \right\}.$$

It follows that the first order condition for a profit maximum in α_i^N also takes the same form as in the closed economy (using symmetry to drop the i):

$$\rho\kappa q_\alpha^\kappa \int_{\underline{x}}^{\widetilde{x}^T} p_x K_x^{1-\kappa} \,\mathrm{d}\,x = p_D {\alpha^N}^{\frac{1-\rho}{\rho}}. \qquad \text{(A.17)}$$

Next, consider the production of tasks. Let $Y_x := Y_x^N + Y_x^S$ denote world production of tasks. The previous results now imply that

$$\begin{aligned}
Y_x &= \alpha^N s_\alpha q_\alpha^\kappa K_x^{1-\kappa} + \alpha^N (1 - s_\alpha) q_\alpha^\kappa K_x^{1-\kappa} \\
&= \alpha^N q_\alpha^\kappa K_x^{1-\kappa}
\end{aligned}$$

for all $x < \widetilde{x}^T$, and

$$\begin{aligned}
Y_x &= \gamma^N(m^{T^{-1}}(x), x) L^N_{m^{T^{-1}}(x)} \frac{\mathrm{d}\, m^{T^{-1}}(x)}{\mathrm{d}\, x} + \gamma^S(m^{T^{-1}}(x), x) L^S_{m^{T^{-1}}(x)} \frac{\mathrm{d}\, m^{T^{-1}}(x)}{\mathrm{d}\, x} \\
&= \gamma^N(m^{T^{-1}}(x), x) \left(L^N_{m^{T^{-1}}(x)} + \Delta L^S_{m^{T^{-1}}(x)} \right) \frac{\mathrm{d}\, m^{T^{-1}}(x)}{\mathrm{d}\, x}
\end{aligned}$$

for $x \geq \widetilde{x}^T$. Again, both equations, written in terms of world quantities, take the same form as in the closed economy.

Considering the supply of intermediate goods for final good production (β-intermediates), for the same reason as in the case of α-intermediates only Northern monopolists will produce

β-intermediates. They will therefore also face iso-elastic demand and charge constant markups, $p^N_{\beta,i} = \eta_\beta/\kappa$. To derive an inverse demand equation in terms of world quantities, define $s_\beta :=$ $(Y^{NN}_x + Y^{SN}_x)/Y_x$ as the share of world task output that is utilized in the North. Note that this share is constant across tasks, since otherwise the marginal products of tasks, and hence task prices, would differ across countries. Using s_β, inverse demand for β-intermediates in the North can be written as:

$$\begin{aligned} p^N_{\beta,i} &= \kappa \beta^N_i {q^{NN}_{\beta,i}}^{\kappa-1} \left[\int_{\underline{x}}^{\overline{x}} \left(Y^{NN}_x + Y^{SN}_x \right)^{\frac{\epsilon-1}{\epsilon}} \,\mathrm{d}\, x \right]^{\frac{\epsilon(1-\kappa)}{\epsilon}} \\ &= \kappa \beta^N_i {q^{NN}_{\beta,i}}^{\kappa-1} s_\beta^{1-\kappa} \left[\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \,\mathrm{d}\, x \right]^{\frac{\epsilon(1-\kappa)}{\epsilon}}, \end{aligned}$$

and inverse demand in the South:

$$\begin{aligned} p^N_{\beta,i} &= \kappa \beta^N_i {q^{NS}_{\beta,i}}^{\kappa-1} \left[\int_{\underline{x}}^{\overline{x}} \left(Y^{SS}_x + Y^{NS}_x \right)^{\frac{\epsilon-1}{\epsilon}} \,\mathrm{d}\, x \right]^{\frac{\epsilon(1-\kappa)}{\epsilon}} \\ &= \kappa \beta^N_i {q^{NS}_{\beta,i}}^{\kappa-1} (1 - s_\beta)^{1-\kappa} \left(\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \,\mathrm{d}\, x \right)^{\frac{\epsilon(1-\kappa)}{\epsilon}}. \end{aligned}$$

The two inverse demand equations imply that the share s_β also equals the share of world output of β-intermediates that is utilized in the North: $q^{NN}_{\beta,i} = s_\beta q_{\beta,i}$, or, using symmetry across intermediate varieties, $q^{NN}_\beta = s_\beta q_\beta$, where $q_{\beta,i}$ and q_β denote world output. With this observation, and again using symmetry across i, the inverse demand equations imply

$$p_\beta = \kappa \beta^N q_\beta^{\kappa-1} \left(\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \,\mathrm{d}\, x \right)^{\frac{\epsilon(1-\kappa)}{\epsilon}}, \tag{A.18}$$

which is the desired inverse demand equation in terms of world quantities. Profits of β-monopolists are then given by

$$\pi_{\beta,i}(\beta^N_i) = \max_q \left\{ \kappa \beta^N_i q^\kappa \left(\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \,\mathrm{d}\, x \right)^{\frac{\epsilon(1-\kappa)}{\epsilon}} - \eta_\beta q \right\} - p_D {\beta^N_i}^{\frac{1}{\rho}},$$

such that the first order condition for a profit maximum in β^N_i becomes (using symmetry to drop the i):

$$\rho \kappa q_\beta^\kappa \left(\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \,\mathrm{d}\, x \right)^{\frac{(1-\kappa)\epsilon}{\epsilon-1}} = p_D {\beta^N}^{\frac{1-\rho}{\rho}}. \tag{A.19}$$

Finally, let $Y = Y^N + Y^S$ be world (gross) production of the final good. Using the share s_β, we can write

$$Y = \beta^N (s_\beta q_\beta)^\kappa \left(\int_{\underline{x}}^{\overline{x}} (s_\beta Y_x)^{\frac{\epsilon-1}{\epsilon}} \,\mathrm{d}\,x \right)^{\frac{(1-\kappa)\epsilon}{\epsilon-1}} + \beta^N ((1-s_\beta) q_\beta)^\kappa \left(\int_{\underline{x}}^{\overline{x}} ((1-s_\beta) Y_x)^{\frac{\epsilon-1}{\epsilon}} \,\mathrm{d}\,x \right)^{\frac{(1-\kappa)\epsilon}{\epsilon-1}}$$

$$= \beta^N q_\beta^\kappa \left(\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \,\mathrm{d}\,x \right)^{\frac{(1-\kappa)\epsilon}{\epsilon-1}} . \tag{A.20}$$

This gives world final good production in terms of world quantities of tasks and β-intermediates.

Collecting the derived equations yields equilibrium conditions for world quantities and prices that are common across countries. These conditions exactly replicate the equilibrium conditions for the closed economy in Section 2.5.3. In particular, in any trade equilibrium the common matching function $m^T(x)$, the common automation threshold $\widetilde{x}^T$, world quantities $\{K_x\}_{x \in X}$, $\{Y_x\}_{x \in X}$, q_α, and q_β, and productivity levels α^N and β^N; the common prices $\{p_x\}_{x \in X}$, p_c^T, p_α^N, p_β^N, the Northern R&D price p_D, and Northern wages $\{w_s^N\}_{s \in S}$ satisfy the following conditions:

(E1)" $Y_x = \alpha^N q_\alpha^\kappa K_x^{1-\kappa}$ if $x < \widetilde{x}^T$ and $Y_x = \gamma^N(m^{T^{-1}}(x), x) \left(L^N_{m^{T^{-1}}(x)} + \Delta L^S_{m^{T^{-1}}(x)} \right) \frac{\mathrm{d}\, m^{T^{-1}}(x)}{\mathrm{d}\, x}$ if $x \geq \widetilde{x}$;

(E2)" $p_x = \frac{\partial Y}{\partial Y_x}$ for all x, where Y is given by (A.20);

(E3)" q_β satisfies equation (A.18);

(E4)" $m^T(s) \in \operatorname{argmax}_{x \in X} \gamma^N(s, x) p_x$ for all s;

(E5)" $w_s^N = \gamma^N(s, m^T(s)) p_{m^T(s)}$ for all s;

(E6)" $\left(\frac{p_\alpha^N}{\kappa \alpha^N} \right)^\kappa \left(\frac{p_c^T}{(1-\kappa)\alpha^N} \right)^{1-\kappa} = p_x$ for all $x < \widetilde{x}^T$ and $p_c^T = r$;

(E7)" $\frac{w_{\underline{s}}^N}{\gamma^N(\underline{s}, \widetilde{x}^T)} = \left(\frac{p_\alpha^N}{\kappa \alpha^N} \right)^\kappa \left(\frac{r}{(1-\kappa)\alpha} \right)^{1-\kappa}$;

(E8)" q_α satisfies equation (A.16);

(E9)" $p_\alpha^N = \frac{\eta_\alpha}{\kappa}$ and $p_\beta^N = \frac{\eta_\beta}{\kappa}$;

(E10)" α^N, β^N, and p_D satisfy equations (A.17), (A.19), and $\alpha^{N\frac{1}{\rho}} + \beta^{N\frac{1}{\rho}} = D$.

These conditions are identical to conditions (E1)' to (E10)' for the closed economy. Therefore, the collection of world quantities and prices described in (E1)" to (E10)" is identical to the equilibrium of a closed economy with labor supply $L^N + \Delta L^S$ and parameter values of the North. The difference between world quantities and prices under trade and the Northern autarky equilibrium is thus identical to the difference in Northern autarky variables that arises from a hypothetical change in labor supply from L^N to $L^N + \Delta L^S$. This verifies the claim in Section 2.5.4 that the effects of trade integration on technology α^N and wages w^N are the same as the effects of a change in labor supply from L^N to $L^N + \Delta L^S$.

A.2 Supplementary Material

A.2.1 Further Models of Endogenous Technical Change

This section derives equations (2.1) and (2.2) as characterizations of wages and equilibrium technologies in two specific models of endogenous technical change, complementing the baseline model presented in the main text. The section also provides conditions for the existence of symmetric exogenous and endogenous technology equilibria in each model.

Baseline Model First, consider existence of symmetric equilibria in the baseline model. Recall that an exogenous technology equilibrium consists of wages and labor inputs for each firm such that firms maximize their profits by choosing their labor inputs, taking wages and technologies $\theta_i = \theta$ as given, and the labor market clears. An endogenous technology equilibrium additionally consists of technologies for each firm, and firms maximize their profits by choosing labor inputs and technologies jointly, again taking wages as given.

Observation 1. *In the baseline model, there exists a symmetric exogenous technology equilibrium at any pair (L, θ) if and only if $F(L, \theta)$ is concave in L at any θ. If $F(L, \theta)$ is strictly concave in L, the symmetric equilibrium is the unique exogenous technology equilibrium.*

Moreover, there exists a symmetric endogenous technology equilibrium at any L if and only if the endogenous technology production function $\overline{F}(L)$ is concave. If $\overline{F}(L)$ is strictly concave and $\text{argmax}_{\theta \in \Theta} F(L, \theta)$ is a singleton for all L, the symmetric equilibrium is the unique endogenous technology equilibrium at any L.

Proof. For a given labor supply $\overline{L}$ and technology θ, a symmetric exogenous technology equilibrium exists if and only if we can find wages w such that

$$F(L, \theta) - wL$$

is maximized with respect to L at $\overline{L}$. Let $F'_\theta(L) := F(L, \theta)$ be the production function at fixed technology θ. Then, the problem is equivalent to finding a hyperplane that is tangent to the graph of F'_θ at $\overline{L}$ and lies above $F'_\theta(L)$ at all L. Such hyperplanes exist for all $\overline{L}$ if and only if $F'_\theta(L)$ is concave.

If $F(L,\theta)$ is strictly concave in L, the profit maximization problem has a unique solution for any wage vector for which a solution exists. Hence, all firms must have the same labor input, and the symmetric equilibrium is the only one that can exist.

A symmetric endogenous technology equilibrium at a given labor supply $\overline{L}$ exists if and only if we can find wages w such that

$$\overline{F}(L) - wL$$

is maximized with respect to L at $\overline{L}$. This is because $\operatorname{argmax}_{\theta\in\Theta} F(\overline{L},\theta)$ is always non-empty by compactness of Θ and continuity of F (such that $\overline{F}(L)$ is well defined). The existence proof then proceeds as for the exogenous technology equilibrium but with $\overline{F}$ in the place of F'_θ.

If $\overline{F}(L)$ is strictly concave and $\operatorname{argmax}_{\theta\in\Theta} F(\overline{L},\theta)$ is a singleton, the profit maximization problem of firms has a unique solution for any wage vector for which a solution exists. Hence, all firms choose the same labor input and technology, and the symmetric equilibrium is the only one that can exist.

□

The important insight is that a symmetric endogenous technology equilibrium exists at all L only if the endogenous technology production function $\overline{F}$ is concave. This prevents the analysis of the phenomena discussed in Section 2.3.2 of the main text within the baseline model, at least if the analysis is restricted to symmetric equilibria. More precisely, whenever symmetric equilibria exist everywhere in the baseline model, the induced technical change effect will never dominate the direct effect of increases in relative skill supply on skill premia, because this would require that $\overline{F}$ is not quasiconcave by Theorem 3. To analyze cases where skill premia increase in relative skill supply in a symmetric equilibrium, we must therefore consider models with interdependences across firms' technology choices, as presented in the next sections.

Spillover Model The spillover model is identical to the baseline model except for that it includes cross effects between firms' technologies. In particular, the production function of firm i is now given by $F'(L_i,\theta_i,\overline{\theta})$, where $\overline{\theta}$ is the average of all firms' technology choices, $\overline{\theta} = \int_0^1 \theta_i \, \mathrm{d}\, i$. For the average to be well defined, let Θ be a convex subset of $\mathbb{R}^N$. Instead of the average, we could use any other function of all firms' technologies that is insensitive to any single firm's θ_j. Denote by $F(L,\theta) := F'(L,\theta,\theta)$ the symmetric technology production function, which gives output as a function of labor input and a common technology for all firms.

The equilibrium definitions are as in the baseline model. At fixed technology $\theta_i = \theta$ for all i, an exogenous technology equilibrium is given by wages and labor inputs for each firm, such that firms choose their labor inputs to maximize profits given wages. As in the baseline model, it is clear that wages have to satisfy

$$w(L,\theta) = \nabla_L F'(L,\theta,\theta) = \nabla_L F(L,\theta)$$

in any symmetric exogenous technology equilibrium. So, equation (2.1) holds.

An endogenous technology equilibrium is given by wages, labor inputs, and technologies for all firms, such that firms choose their labor inputs and technologies to maximize profits, taking wages and the technologies of other firms as given. Let

$$\theta^*(L) \in \underset{\theta \in \Theta}{\operatorname{argmax}}\, F(L,\theta)$$

be a common technology across firms that maximizes output at symmetric labor inputs. Moreover, suppose that the spillovers across firms' technologies are such that each firm benefits from other firms choosing similar technologies to its own.

Assumption 7. *For each firm i and any labor input L_i,*

$$F'(L_i,\theta_i,\theta_i) \geq F'(L_i,\theta_i,\overline{\theta})$$

for all feasible $\overline{\theta}$.

In words, for any individual technology θ_i firm i's productivity is maximized when the other firms choose θ_i as well on average. This captures the notion that part of the knowledge about how to work with a given technology is non-excludable, such that firms' productivity increases when other firms operate the same technology and much useful knowledge spills over. A perhaps more stringent formalization would have any firm's productivity decrease in the average distance between its own and other firms' technologies. Such a modification is straightforward and thus omitted.

Under Assumption (7) and appropriate conditions on F', any technology $\theta^*(L)$ as described above forms a symmetric endogenous technology equilibrium when combined with wages $w(L,\theta^*(L)) = \nabla_L F(L,\theta^*(L))$ and symmetric labor inputs L for all firms. Thus, equation (2.2) applies. More comprehensively, the following results hold.

Observation 2. *In the spillover model, there exists a symmetric exogenous technology equilibrium at any pair (L,θ) if and only if the symmetric technology function $F(L,\theta)$ is concave in L at any θ. If F is strictly concave in L, the symmetric exogenous technology equilibrium is the unique exogenous technology equilibrium. Wages in any symmetric technology equilibrium are given by equation* (2.1).

Moreover, suppose Assumption 7 holds and the endogenous technology function $\overline{F}'(L,\theta) := \max_{\theta_i} F'(L,\theta_i,\theta)$ *is concave in* L*. Then, for any labor supply* L*, any technology* $\theta^*(L)$ *as given by equation* (2.2) *forms a symmetric endogenous technology equilibrium in combination with wages* $w(L,\theta^*(L)) = \nabla_L F(L,\theta^*(L))$ *and symmetric labor inputs* L *for each firm.*

Proof. The exogenous technology equilibrium is equivalent to the exogenous technology equilibrium of the baseline model, so the first part follows directly from Observation 1.

For the second part, we have to show that there are wages w such that the pair $(L,\theta^*(L))$ maximizes firm profits

$$F'(L_i,\theta_i,\theta^*(L)) - wL_i$$

with respect to (L_i,θ_i). First, for any $\theta' \in \Theta$ we have

$$F'(L,\theta^*(L),\theta^*(L)) \geq F'(L,\theta',\theta') \geq F'(L,\theta',\theta^*(L)).$$

It follows that $\overline{F}'(L,\theta^*(L)) = F'(L,\theta^*(L),\theta^*(L))$. Now let $w^* = \nabla_L F'(L,\theta^*(L),\theta^*(L))$ and note that by the envelope theorem,

$$w^* = \nabla_L \overline{F}'(L,\theta^*(L)).$$

Concavity of $\overline{F}'$ in L then implies that profits are indeed maximized at $(L,\theta^*(L))$ when wages are given by w^*. □

Uniqueness of the symmetric endogenous technology equilibrium can easily be ensured by restricting spillovers to be sufficiently weak in an appropriate sense. The more interesting result, however, is that existence of a symmetric endogenous technology equilibrium as characterized by equations (2.1) and (2.2) only requires concavity of $\overline{F}'(L,\theta)$ in L, and not in L and θ jointly. In consequence, also the symmetric technology function $F(L,\theta)$ does not have to be jointly concave in L and θ. The reason is that existence of a symmetric equilibrium only requires concavity in the choice variables of an individual firm, whereas the function $F(L,\theta)$ combines an individual firm's technology and the average technology across firms in the variable θ (by restricting the two to be the same). Therefore, in a symmetric endogenous technology equilibrium of the spillover model, skill premia may increase in relative skill supply as described in Proposition 2 and Theorem 3.

Monopolistic Competition Model The distinction between concavity of individual decision problems and the aggregate production function becomes even more transparent in the monopolistic competition model. There are now two types of firms, a continuum of final good firms and a continuum of technology firms. Final good firms produce the single consumption good (the numéraire), using labor and technology-embodying intermediate goods as inputs. Their production function is $F'(L_i, Q_i)$, where L_i is firm i's labor input and $Q_i = (Q_{i,k})_{k=1,2,\dots,K}$ is a vector of aggregates of technology-embodying intermediate goods. In particular, for each k,

$$Q_{i,k} = \int_0^1 \theta_{k,x} q_{i,k,x}^{\kappa} \,\mathrm{d}\, x,$$

where (k, x) indexes technology firms, $q_{i,k,x}$ is the quantity of firm (k, x)'s intermediate good used by final good firm i, and $\theta_{k,x}$ is the intermediate's quality. Technology firms are monopolistically competitive with substitution parameter $\kappa \in (0, 1)$. They produce their intermediate goods at constant marginal cost η_k from final good, facing inverse demand

$$p_{k,x} = \frac{\partial F'(L_i, Q_i)}{\partial Q_{i,k}} \kappa \theta_{k,x} q_{i,k,x}^{\kappa-1}$$

from final good firm i. Since inverse demand is iso-elastic, all technology firms charge a price of $p_{k,x} = \eta_k/\kappa$. The symmetric price is denoted by p_k henceforth. Moreover, denote the total output of firm (k, x) by $q_{k,x}$. Then, profits of firm (k, x) are given by

$$\pi_{k,x}(\theta_{k,x}) = \max_{(q_i)_{i\in[0,1]}} \left\{ \kappa\theta_{k,x} \int_0^1 \frac{\partial F'(L_i, Q_i)}{\partial Q_{i,k}} q_i^{\kappa} \,\mathrm{d}\, i - \eta_k \int_0^1 q_i \,\mathrm{d}\, i - C_k(\theta_{k,x}) \right\}.$$

The first order condition for the firm's quality choice is

$$\kappa \int_0^1 \frac{\partial F'(L_i, Q_i)}{\partial Q_{i,k}} q_{i,k,x}^{\kappa} \,\mathrm{d}\, i = -\frac{\mathrm{d}\, C_k(\theta_{k,x})}{\mathrm{d}\, \theta_{k,x}}.$$

It can be verified that the elasticity of the optimal $q_{i,k,x}$ in $\theta_{k,x}$ is $1/(1 - \kappa)$. Then, assuming that the elasticity of $\mathrm{d}\, C_k/\,\mathrm{d}\, \theta$ is always greater than $\kappa/(1 - \kappa)$, the first order condition has a unique solution, which is necessary and sufficient for a maximum. In summary, technology firms' problem of choosing price and quality of their output has a unique solution which is necessarily symmetric across firms. The symmetric quantities and qualities are denoted by $q_{i,k}$ and $\theta = (\theta_k)_{k=1,2,\dots,K}$ henceforth.

Equilibrium conditions can now directly be stated in terms of the symmetric choices of technology firms. In particular, an exogenous technology equilibrium is a collection of labor inputs L_i, intermediate inputs $q_{i,k}$, intermediate prices p_k, and wages w, such that final good firms choose their labor and intermediate inputs to maximize profits taking prices and wages as given, technology firms choose their prices to maximize profits taking inverse demand curves from final good firms and the quality levels of their output as given, and the labor market clears. An endogenous technology equilibrium additionally consists of quality levels θ_k for technology firms, and requires technology firms to choose both their prices and quality levels to maximize profits, taking inverse demand curves from final good firms as given.

To characterize symmetric equilibria in the form of equations (2.1) and (2.2), define the following "modified production function" (see also Lemma 4):

$$F(L_i, \theta) \coloneqq \max_{(q_k)_{k=1,2,\ldots,K}} \left\{ F'(L_i, (\theta_k q_k^\kappa)_{k=1,2,\ldots,K}) - \frac{\eta_k}{\kappa} q_k \right\} - \frac{1}{\kappa} \sum_{k=1}^{K} C_k(\theta_k).$$

Then, with technology firms' decisions given by $p_k = \eta_k/\kappa$ and θ, final good firms' objective is equivalent to maximizing

$$F(L_i, \theta) - wL_i$$

with respect to L_i.[3] Therefore, by the same arguments as in the previous models, a symmetric exogenous technology equilibrium exists at all L and θ if and only if $F(L_i, \theta)$ is concave in L_i at all θ. Moreover, in such a symmetric equilibrium, wages are given by

$$w = \nabla_L F(L, \theta),$$

that is, equation (2.1) holds. If $F(L_i, \theta)$ is also strictly concave in L_i at all θ, any exogenous technology equilibrium will feature symmetric labor inputs and wages given by equation (2.1).

For a symmetric endogenous technology equilibrium, take any technology

$$\theta^*(L) \in \operatorname*{argmax}_{\theta \in \mathbb{R}_+^K} F(L, \theta).$$

Such a technology must satisfy the first order conditions

$$\frac{\partial F'(L, Q^*(L))}{\partial Q_k} \kappa q_k^*(L)^\kappa = \frac{\mathrm{d}\, C_k(\theta^*(L)_k)}{\mathrm{d}\, \theta_k}$$

for all k, where $Q^*(L) = \theta^*(L) q^*(L)^\kappa$ and $q^*(L)$ is a solution to

$$\max_{(q_k)_{k=1,2,\ldots,K}} \left\{ F'(L, (\theta_k^*(L) q_k^\kappa)_{k=1,2,\ldots,K}) - \eta_k/\kappa q_k \right\}.$$

Thereby, $\theta^*(L)$ and $q_k^*(L)$ jointly satisfy technology firms' first oder conditions and final goods' inverse demand for intermediates, when labor inputs are symmetric. For a symmetric endogenous

[3] Note that final good firms take θ as given, such that the presence of the term $\sum C_k(\theta_k)$ in $F(L_i, \theta)$ does not change the maximization problem.

technology equilibrium, it remains to find wages w such that symmetric labor inputs maximize

$$F(L_i, \theta^*(L)) - wL.$$

Such wages, again, exist at all L if and only if F is concave in L_i at $\theta^*(L)$. Moreover, they will clearly satisfy $w = \nabla_L F(L, \theta^*(L))$. We have therefore established that a symmetric endogenous technology equilibrium with equilibrium technology given by (2.2) and wages by (2.1) exists whenever F is concave in L_i.

Observation 3. *In the monopolistic competition model, there exists a symmetric exogenous technology equilibrium at any pair* (L, θ) *if and only if* $F(L, \theta)$ *is concave in* L *at any* θ*. If* F *is strictly concave in* L*, labor inputs are symmetric in any exogenous technology equilibrium. Whenever labor inputs are symmetric, wages are given by equation* (2.1)*.*

Moreover, if $F(L, \theta)$ *is concave in* L*, there exists a symmetric endogenous technology equilibrium with equilibrium technology satisfying equation* (2.2) *and wages given by* (2.1)*.*

Uniqueness of the endogenous technology equilibrium can be ensured by imposing that F is strictly pseudoconcave in θ – such that a unique technology satisfies technology firms' first order conditions at symmetric final good firm choices – and F' is strictly concave in the $q_{i,k,x}$ – such that all final good firms indeed choose the same intermediate quantities. The more important insight from Observation 3 is, however, that existence of symmetric endogenous and exogenous technology equilibria can be guaranteed without any restriction on the curvature of $F(L, \theta)$ in L and θ jointly. Only restrictions on the curvature of F in L (for existence) and in θ (for uniqueness) individually are needed. In particular, the endogenous technology function $\overline{F}(L) = F(L, \theta^*(L))$ can be quasiconvex, as required for strong bias by Theorem 3.

Finally, note that the monopolistic competition model embeds static versions of well-known models from previous work as special cases. First, when

$$F'(L, Q) = QL^{1-\kappa},$$

with L denoting labor supply of a single skill level, we obtain a static version of the standard monopolistic competition based growth models developed by Romer (1990) and Aghion and Howitt (1992). Since this model neither features wage inequality nor biased technical change, its static version is not very interesting. A more interesting case is obtained when

$$F'(L, Q) = \left[\left(Q_1 L_1^{1-\kappa}\right)^{\rho} + \left(Q_2 L_2^{1-\kappa}\right)^{\rho}\right]^{(1/\rho)}.$$

This is a static version of the seminal directed technical change model by Acemoglu (1998).

A.2.2 Generalization of the Weak Bias Theorem

This section presents a generalization of Theorem 1 on the induced technical change effect. The generalization provides a partial converse to the statement of Theorem 1, giving precise limits to the occurrence of the weak bias phenomenon.

First, note that the skill bias order $\preceq^b$ on the set of feasible technologies Θ is actually a preorder. That is, it is reflexive, transitive, but not necessarily antisymmetric. There may, for example, be two distinct technologies θ and θ' that induce the same wage distribution at any labor input, such that $\theta \preceq^b \theta'$, $\theta' \preceq^b \theta$, and $\theta \neq \theta'$. Alternatively, θ and θ' may induce the same set of relative wages but at different wage levels. In both cases, θ and θ' can be ordered by their skill bias in both directions but they are not equal. Let $\sim$ denote the equivalence relation connecting technologies with the same skill bias, that is,

$$\theta \sim \theta' \Leftrightarrow [\theta \preceq^b \theta' \wedge \theta' \preceq^b \theta].$$

Given the preorder $\preceq^b$, we can define what it means for the partially ordered set $(\Theta, \preceq^b)$ to be a prelattice.

Definition 7. The pair $(\Theta, \preceq^b)$ is a prelattice if any two elements $\theta, \theta' \in \Theta$ have a supremum and an infimum in Θ.

Note that in a prelattice, in contrast to a lattice, supremum and infimum are not necessarily unique for all pairs of elements. Moreover, whenever all elements in Θ can be ordered according to their skill bias, as demanded by Theorem 1, then $(\Theta, \preceq^b)$ will automatically be a prelattice.

Besides the prelattice structure of $(\Theta, \preceq^b)$, the generalization of Theorem 1 requires F to be prequasisupermodular in θ.

Definition 8. The function $F(L, \theta)$ is prequasisupermodular in θ if, for any L and $\theta, \theta' \in \Theta$,

$$F(L, \underline{\theta}) \leq F(L, \theta) \text{ for all } \underline{\theta} \in \inf(\theta, \theta') \Rightarrow F(L, \theta') \leq F(L, \overline{\theta}) \text{ for some } \overline{\theta} \in \sup(\theta, \theta'),$$

where $\inf(\theta, \theta')$ denotes the set of infima of θ and θ', and $\sup(\theta, \theta')$ denotes the set of suprema.

Prequasisupermodularity is therefore defined analogously to quasisupermodularity (see, for example, (Milgrom and Shannon, 1994)), adapted to the preorder environment (quasisupermodularity is defined on sets endowed with a usual, that is, antisymmetric, order relation). Again, whenever all elements in Θ can be ordered according to their skill bias, the function F is prequasisupermodular in θ without any further assumptions. This is because θ and θ' are elements of their infimum and supremum sets themselves, then.

The generalization of Theorem 1 is now stated as follows.

Theorem 5. *Suppose* $(\Theta, \preceq^b)$ *is a prelattice and* F *is prequasisupermodular in* θ. *Then,*

$$L \preceq^s L' \Rightarrow \theta^*(L) \preceq^b \theta^*(L')$$

if and only if $\theta^*(L) \sim \theta^*(\lambda L)$ *for all* L *and* $\lambda \in \mathbb{R}_{++}$ *(that is, the skill bias of the equilibrium technology is scale invariant).*

Proof. By zero homogeneity of the skill bias of θ^*, we can restrict attention to changes from L to L' such that $F(L, \theta^*(L')) = F(L', \theta^*(L'))$. Moreover, by definition of θ^*, it must hold that $F(L, \underline{\theta}) \leq F(L, \theta^*(L))$ for all $\underline{\theta} \in \inf(\theta^*(L), \theta^*(L'))$. Therefore, by prequasisupermodularity, there must exist a $\overline{\theta} \in \sup(\theta^*(L), \theta^*(L'))$ such that $F(L, \theta^*(L')) \leq F(L, \overline{\theta})$. We can now assume that $\overline{\theta} \npreceq^b \theta^*(L')$, because otherwise the statement of the theorem is immediately satisfied. Under this assumption, it must hold that

$$F(L, \overline{\theta}) \geq F(L, \theta^*(L')) = F(L', \theta^*(L')) > F(L', \overline{\theta}),$$

where the last inequality is strict because θ^* is selected as the supremum of the maximizer set in equation (2.2). From here on, the proof proceeds analogously to the proof of Theorem 1 starting from equation (A.2): $\overline{\theta}$ here takes the role of $\theta^*(L')$ in the proof of Theorem 1, and $\theta^*(L')$ here takes the role of $\theta^*(L)$ in the proof of Theorem 1. Moreover, the inequalities (A.3) and (A.4) are reversed, and the contradiction at the end is obtained by observing that $\theta^*(L') \preceq^b \overline{\theta}$ implies that the left-hand-side of inequality (A.4) must be positive (instead of strictly negative as implied by the preceding arguments). □

Theorem 5 generalizes Theorem 1 in two ways. First, it replaces the assumption that any two technologies can be ordered according to their skill bias by imposing a prelattice structure on Θ and prequasisupermodularity on F. The prelattice structure and prequasisupermodularity ensure that the set of equilibrium technologies $\theta^*(L)$ will be totally ordered along any curve in the labor supply space that is totally ordered itself under the relative skill supply (pre)order $\preceq^s$. Second, Theorem 5 replaces zero homogeneity of $\theta^*(L)$ with zero homogeneity of the skill bias of $\theta^*(L)$. That is, those components of θ^* that do not affect relative wages are allowed to change when scaling labor supply up or down. Zero homogeneity of the skill bias of θ^* is clearly necessary for weak bias, as any violation would constitute a counterexample to the weak bias phenomenon. This gives rise to the "only if" part in Theorem 5.

Finally, note that Theorem 5 is not a direct application of the main theorem of monotone comparative statics (Theorem 4 in (Milgrom and Shannon, 1994)), although the two are closely related. The relevant part of Theorem 4 from Milgrom and Shannon (1994) says the following.

Theorem 6 (cf. Milgrom and Shannon, 1994)**.** *Let* $(X, \preceq^a)$ *be a lattice and* $(P, \preceq^b)$ *a partially ordered set. Consider a family of functions* $\{f(\cdot; p)\}_{p \in P}$ *with* $f : X \times P \to \mathbb{R}$*. Let* $f(x; p)$ *be quasisupermodular in* x *and have the single crossing property in* $(x; p)$*. Then,*

$$p \preceq^b p' \Rightarrow \sup \operatorname*{argmax}_{x \in X} f(x; p) \preceq^a \sup \operatorname*{argmax}_{x \in X} f(x; p').$$

It can be shown that the theorem still holds when $\preceq^a$ and $\preceq^b$ are preorders and F is prequasisupermodular in x. The important difference between Theorem 6 and Theorem 5 is that the former imposes the single crossing property in $(x; p)$ on F.[4] The latter instead uses specifically defined (pre)order relations $\preceq^b$ and $\preceq^s$. Indeed, these specific orderings already introduce a complementarity between changes along $\preceq^s$ (increases in relative skill supply) and changes along $\preceq^b$ (skill-biased technical change). Such a complementarity is assumed via the single crossing property in Theorem 6. One can show, however, that the conditions of Theorem 5 do not imply the single crossing property in $(\theta; L)$ for F. Therefore, given the specific environment introduced in the main text (the preorder relations and the structure of the labor supply space), Theorem 5 cannot be obtained as a corollary to Theorem 6.

[4]The single crossing property in $(x; p)$ means that $F(x', p) - F(x, p) \geq (>) 0$ implies $F(x', p') - F(x, p') \geq (>) 0$ for any $x \preceq^a x'$ and $p \preceq^b p'$.

B Appendix to Chapter 3

B.1 Proofs and Derivations

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

B.1.1 Proofs and Derivations for the Setup

Here I provide proofs and derivations for Section 3.3.

Derivation of the labor demand equation (3.3)

I derive the labor demand equation (3.3) in detail to demonstrate that the functional derivative D_{l_θ} works as expected.

Final good firm profits are given by

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

Taking the derivative D_{L_θ} as defined in Section 3.3.2 and equating it with zero yields:

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

The remaining task is to show that

$$D_{L_\theta} \int_{\underline{\theta}}^{\overline{\theta}} w_{\widetilde{\theta}} L_{\widetilde{\theta}} \, d\widetilde{\theta} = w_\theta.$$

I derive this equality for interior types $\theta \in (\underline{\theta}, \overline{\theta})$ in detail to demonstrate the working of the functional derivative D_{L_θ}. The derivations for the highest and lowest types $\overline{\theta}$ and $\underline{\theta}$ proceed analogously and are therefore omitted.

By definition:

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

Moreover, by definition of $\widetilde{L}_{\Delta,\theta}$:

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

Hence:

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

Then, by L'Hôspital's rule:

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

Applying L'Hôspital's rule again, we obtain:

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

where the last equality requires continuity of w in θ, which I assume is given in equilibrium.

Labor Supply Elasticities

This section derives expressions (3.10) and (3.11) for the labor supply elasticities with respect to the marginal retention rate and the wage. The starting point is workers' first-order condition (3.2):

$$v'(l_\theta(T, w_\theta)) = R_T'(w_\theta l_\theta(T, w_\theta)) w_\theta .$$

Taking the derivative $D_{\widetilde{\tau}}$ on both sides of the equation yields:

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

and hence:

$$D_{\tilde{\tau}} l_\theta(T, w_\theta) = \frac{w_\theta}{v''(l_\theta(T, w_\theta)) + T''(w_\theta l_\theta(T, w_\theta)) w_\theta^2}.$$

By definition of ϵ_θ^R we obtain

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

Again using the first-order condition to replace $(1 - T'(wl))w$ by $v'(l)$, we obtain equation (3.10).

For equation (3.11) differentiate the first-order condition with respect to w_θ on both sides,

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

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 $(1 - T'(wl))w$ by $v'(l)$ yields equation (3.11).

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

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

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 3 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.

Derivation of Aggregate Production, Feasible Technologies, Wages, and Wage Elasticities in the CES Case

The CES case is obtained via the following assumptions on the fundamentals of the model presented in the main text.

$$\begin{aligned}\widetilde{G}(L,\widetilde{\phi},q) &= \left[\int_{\underline{\theta}}^{\overline{\theta}} \left(\widetilde{\kappa}_\theta L_\theta^{1-\alpha} \int_0^1 \widetilde{\phi}_{\theta,k} q_{\theta,k}^{\alpha}\, dk\right)^{\frac{\widetilde{\sigma}-1}{\widetilde{\sigma}}} d\theta\right]^{\frac{\widetilde{\sigma}}{\widetilde{\sigma}-1}} \\ C_\theta(\widetilde{\phi}_{\theta,k}) &= \widetilde{\phi}_{\theta,k}^{\widetilde{\delta}} \ .\end{aligned}$$

The function $\widetilde{\kappa}$ is an exogenous component of technology and assumed to be continuously differentiable; $\widetilde{\sigma} > 0$ measures the elasticity of substitution between differentially skilled workers in the production of an individual final good firm; and $\widetilde{\delta}$ determines the convexity of the research cost function. The endogenous component of technology is $\widetilde{\phi}$.[1]

Aggregate Production To derive the aggregate production function $F(l,\phi)$ as given by equation (3.12), start from its definition:

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

The first-order conditions for the maximization with respect to q are:

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

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

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

[1] Note that here the set of technology types is equated with the set of worker types, such that technology and research costs are now indexed by θ. This reflects the assumption that for every worker type θ there exists a type of technology, embodied in the intermediate goods $q_{\theta,k}$, that raises the efficiency of labor of type θ in the production process. Moreover, the set of technology types is a continuum here, in contrast to the finite set $\{1, 2, ..., J\}$ in the general model above. As mentioned in footnote 7, the case with a continuum of technology types can be treated analogously to the finite case presented above and is therefore omitted from the presentation of the general model.

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

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

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

$$\widetilde{G}(\{h_\theta l_\theta\}_{\theta\in\Theta}, \widetilde{\phi}, q^*) = \alpha^{\frac{\alpha}{1-\alpha}} \left[\int_{\underline{\theta}}^{\overline{\theta}} \left(\eta_\theta^{-\frac{\alpha}{1-\alpha}} \widetilde{\kappa}_\theta^{\frac{1}{1-\alpha}} \widetilde{\phi}_\theta^{\frac{1}{1-\alpha}} h_\theta l_\theta\right)^{\frac{(1-\alpha)\widetilde{\sigma}}{\alpha+\widetilde{\sigma}-\alpha\widetilde{\sigma}}\frac{\widetilde{\sigma}-1}{\widetilde{\sigma}}} d\theta\right]^{\frac{\alpha+\widetilde{\sigma}-\alpha\widetilde{\sigma}}{(1-\alpha)\widetilde{\sigma}}\frac{\widetilde{\sigma}}{\widetilde{\sigma}-1}}. \tag{B.2}$$

This provides an expression for gross aggregate production. Using the maximizer q^* from equation (B.1) again, the part of gross output that goes into the production of intermediate goods becomes

$$\begin{aligned}
\int_{\underline{\theta}}^{\overline{\theta}} \eta_\theta q_\theta^*\, d\theta &= \int_{\underline{\theta}}^{\overline{\theta}} \eta_\theta^{\frac{\alpha-\alpha\widetilde{\sigma}}{\alpha+\widetilde{\sigma}-\alpha\widetilde{\sigma}}} \alpha^{\frac{\widetilde{\sigma}}{\alpha+\widetilde{\sigma}-\alpha\widetilde{\sigma}}} \left(\widetilde{\kappa}_\theta \widetilde{\phi}_\theta h_\theta^{1-\alpha} l_\theta^{1-\alpha}\right)^{\frac{\widetilde{\sigma}-1}{\alpha+\widetilde{\sigma}-\alpha\widetilde{\sigma}}} \widetilde{G}^{\frac{1}{\alpha+\widetilde{\sigma}-\alpha\widetilde{\sigma}}}\, d\theta \\
&= \alpha^{\frac{\widetilde{\sigma}}{\alpha+\widetilde{\sigma}-\alpha\widetilde{\sigma}}} \widetilde{G}^{\frac{1}{\alpha+\widetilde{\sigma}-\alpha\widetilde{\sigma}}} \int_{\underline{\theta}}^{\overline{\theta}} \left(\eta_\theta^{-\frac{\alpha}{1-\alpha}} \widetilde{\kappa}_\theta^{\frac{1}{1-\alpha}} \widetilde{\phi}_\theta^{\frac{1}{1-\alpha}} h_\theta l_\theta\right)^{\frac{(1-\alpha)\widetilde{\sigma}}{\alpha+\widetilde{\sigma}-\alpha\widetilde{\sigma}}\frac{\widetilde{\sigma}-1}{\widetilde{\sigma}}} d\theta && \text{(B.3)} \\
&= \alpha\widetilde{G}\,. && \text{(B.4)}
\end{aligned}$$

Combining equations (B.2) and (B.3), we obtain net aggregate production F as follows:

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

Defining

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

net aggregate production becomes

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

which is equation (3.12) from the main text.

Set of Feasible Technologies From the R&D resource constraint and the R&D cost function, the set of feasible technologies $\widetilde{\phi}$ follows as

$$\left\{\widetilde{\phi} : \theta \mapsto \widetilde{\phi}_\theta \in \mathbb{R}_+ \mid \int_{\underline{\theta}}^{\overline{\theta}} \widetilde{\phi}^{\widetilde{\delta}} \, d\theta \le \overline{C}\right\} \; .$$

Using the substitution

$$\phi_\theta \coloneqq \widetilde{\phi}_\theta^{\frac{1}{1-\alpha}} \quad \forall\theta \; ,$$

the set of feasible ϕ becomes

$$\Phi = \left\{\phi : \theta \mapsto \phi_\theta \in \mathbb{R}_+ \mid \int_{\underline{\theta}}^{\overline{\theta}} \phi_\theta^\delta \, d\theta \le \overline{C}\right\} \; ,$$

where $\delta \coloneqq (1-\alpha)\widetilde{\delta}$, as given in the main text.

Wages I derive expression (3.14) for interior types $\theta \in (\underline{\theta}, \overline{\theta})$. For the boundary types $\underline{\theta}$ and $\overline{\theta}$ the derivations proceed analogously and yield the same result.

Consider first the derivative

$$D_{l_\theta} F(l,\phi) = \lim_{\Delta\to 0} \frac{1}{\Delta} \left. \frac{dF(l+\mu\widetilde{l}_{\Delta,\theta},\phi)}{d\mu} \right|_{\mu=0} \; .$$

Using the definition of $\widetilde{l}_{\Delta,\theta}$ this derivative becomes

$$\left. \frac{dF(l+\mu\widetilde{l}_{\Delta,\theta},\phi)}{d\mu} \right|_{\mu=0}$$
$$= F(l,\phi)^{\frac{1}{\sigma}} \left[\int_{\theta-\Delta}^{\theta} \left(\kappa_{\widetilde{\theta}} \phi_{\widetilde{\theta}} h_{\widetilde{\theta}}\right)^{\frac{\sigma-1}{\sigma}} l_{\widetilde{\theta}}^{-\frac{1}{\sigma}} \frac{\widetilde{\theta}-\theta+\Delta}{\Delta} \, d\widetilde{\theta} + \int_{\theta}^{\theta+\Delta} \left(\kappa_{\widetilde{\theta}} \phi_{\widetilde{\theta}} h_{\widetilde{\theta}}\right)^{\frac{\sigma-1}{\sigma}} l_{\widetilde{\theta}}^{-\frac{1}{\sigma}} \frac{\theta-\widetilde{\theta}+\Delta}{\Delta} \, d\widetilde{\theta}\right] \; .$$

Taking limits and applying L'Hôspital's rule yields:

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

and

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

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 (B.6) below; and continuity of l is presumed in all equilibria under consideration.

Finally, combine the two previous expressions to obtain

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

and therewith

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

Wage Elasticities Again I focus on the derivations for interior types. Given expression (3.14), the own-wage substitution elasticity is simply the elasticity of w_θ with respect to l_θ:

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

The cross-wage substitution elasticity $\gamma_{\theta,\widetilde{\theta}}$ is

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

For the technical change elasticities, consider first the determination of equilibrium technology described by equation (3.8). First-order conditions for the maximization problem in equation (3.8) are

$$\delta \phi_\theta^{*\delta-1} \lambda = \left(\kappa_\theta h_\theta l_\theta\right)^{\frac{\sigma-1}{\sigma}} \phi_\theta^{*-\frac{1}{\sigma}} F(l, \phi)^{\frac{1}{\delta}} \quad \forall \theta , \tag{B.5}$$

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 via 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

$$\phi_\theta^* = (\delta\lambda)^{\frac{-\sigma}{(\delta-1)\sigma+1}} (\kappa_\theta h_\theta l_\theta)^{\frac{\sigma-1}{(\delta-1)\sigma+1}} F(l,\phi)^{\frac{1}{(\delta-1)\sigma+1}} \quad \forall\theta \; . \tag{B.6}$$

Then, we can use the R&D resource constraint

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

to solve for the Lagrange multiplier:

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

Plugging this into equation (B.6), we obtain the following expression for the equilibrium technology ϕ^*:

$$\phi_\theta^* = \overline{C}^{\frac{1}{\delta}} (\kappa_\theta h_\theta l_\theta)^{\frac{\sigma-1}{(\delta-1)\sigma+1}} \left[\int_{\underline{\theta}}^{\overline{\theta}} (\kappa_\theta h_\theta l_\theta)^{\frac{(\sigma-1)\delta}{(\delta-1)\sigma+1}} \, d\theta \right]^{-\frac{1}{\delta}} \; . \tag{B.7}$$

We can now use equation (B.7) to derive the technical change elasticities. The own-wage technical change elasticity is simply derived from equations (3.14) and (B.7) as

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

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

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

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

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

Analogously to the computation of the derivative $D_{l_\theta} F(l,\phi)$ above, we can compute $D_{\phi,l_{\widetilde{\theta}}}\phi^*_\theta(l)$ using equation (B.7):

$$
\begin{aligned}
D_{\phi,l_{\widetilde{\theta}}}\phi^*_\theta(l) = &-\frac{\sigma-1}{(\delta-1)\sigma+1}\overline{C}^{-\frac{1}{\delta}}\left(\kappa_\theta h_\theta l_\theta\right)^{\frac{\sigma-1}{(\delta-1)\sigma+1}}\left[\int_{\underline{\theta}}^{\overline{\theta}}\left(\kappa_\theta h_\theta l_\theta\right)^{\frac{(\sigma-1)\delta}{(\delta-1)\sigma+1}}\,d\theta\right]^{-\frac{1}{\delta}} \\
&\times\left[\int_{\underline{\theta}}^{\overline{\theta}}\left(\kappa_\theta h_\theta l_\theta\right)^{\frac{(\sigma-1)\delta}{(\delta-1)\sigma+1}}\,d\theta\right]^{-1}\left(\kappa_{\widetilde{\theta}} h_{\widetilde{\theta}} l_{\widetilde{\theta}}\right)^{\frac{(\sigma-1)\delta}{(\delta-1)\sigma+1}} l_{\widetilde{\theta}}^{-1} \\
= &-\frac{\sigma-1}{(\delta-1)\sigma+1}\frac{\phi^*_\theta(l)}{l_{\widetilde{\theta}}}\frac{\phi^*_{\widetilde{\theta}}(l)}{\overline{C}}\ .
\end{aligned}
$$

Thereby,

$$
\rho_{\theta,\widetilde{\theta}} = -\frac{\sigma-1}{\sigma}\frac{\sigma-1}{(\delta-1)\sigma+1}\frac{\phi^*_{\widetilde{\theta}}(l)}{\overline{C}}\ . \tag{B.8}
$$

To derive the expression from the main text, note that we can rewrite the first-order condition (B.5) as

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

which implies

$$
\phi_\theta^{*\delta} = \frac{1}{\lambda\delta} w_\theta h_\theta l_\theta\ . \tag{B.9}
$$

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

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

Using this to eliminate λ in equation (B.9), we obtain

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

Finally, combining this with equation (B.8) yields

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

which is the expression given in the main text.

Linear Homogeneity of Aggregate Production

In the main text I assume that final good firms' production function $\widetilde{G}$ 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.

Lemma 12. *The aggregate production function F defined in* (3.6) *is linear homogeneous in l.*

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

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

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 $\widetilde{G}$ 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 $\widetilde{G}$ is concave in l and q, first-order conditions are sufficient for a maximum, and $\lambda q^*(l, \phi)$ is a maximizer of $\widetilde{G}$ at λl and ϕ. So, using linear homogeneity of $\widetilde{G}$ again,

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

□

Consider next the equilibrium aggregate production function

$$F^*(l) := F(l, \phi^*(l)) . \tag{B.10}$$

Lemma 13. *The equilibrium aggregate production function F^* defined in* (B.10) *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 12):

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

for any $\lambda > 0$. □

B.1.2 Proofs for Directed Technical Change

Lemma 6 is a local version of Theorem 5 in Loebbing (2018). Yet it is not strictly covered by the theorem, because, as described in the main text and footnote 16, I use a slightly unusual definition of quasisupermodularity, which allows me to dispense with the lattice structure of Φ. So, I provide a proof for Lemma 6 here. The proof follows closely the proof of Theorem 5 in Loebbing (2018).

Proof of Lemma 6. Take any two labor inputs l and $\widetilde{l}$ such that $\widetilde{l}$ has greater relative skill supply, that is, $\widetilde{l}_\theta/\widetilde{l}_{\widetilde{\theta}} \geq l_\theta/l_{\widetilde{\theta}}$. Since F is linear homogeneous in labor (Lemma 12), wages are independent of the scale of the labor input. So, for the purpose of Lemma 6, we can always scale l up or down such that $F(l, \phi^*(\widetilde{l})) = F(\widetilde{l}, \phi^*(\widetilde{l}))$. In words, we scale l such that it is contained in the (exogenous technology) isoquant of F through $(\widetilde{l}, \phi^*(\widetilde{l}))$.

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

Now assume, to derive a contradiction, that $\phi^*(\widetilde{l}) \not\succeq^{sb} \phi^*(l)$. Then, $\overline{\phi} \neq \phi^*(\widetilde{l})$ and, by uniqueness of $\operatorname{argmax}_{\phi\in\Phi} F(\widetilde{l}, \phi)$ (Assumption 3), we must have $F(\widetilde{l}, \phi^*(\widetilde{l})) > F(\widetilde{l}, \overline{\phi})$.

Combining the previous results, we obtain

$$F(l, \overline{\phi}) \geq F(l, \phi^*(\widetilde{l})) = F(\widetilde{l}, \phi^*(\widetilde{l})) > F(\widetilde{l}, \overline{\phi}) \tag{B.11}$$

for some $\overline{\phi} \succeq^{sb} \phi^*(l), \phi^*(\widetilde{l})$ and $\underline{\phi} \preceq^{sb} \phi^*(l), \phi^*(\widetilde{l})$. In words, increasing relative skill supply by moving from l to $\widetilde{l}$ leaves output unchanged at $\phi^*(\widetilde{l})$ but reduces output at $\overline{\phi}$. Intuitively, this is incompatible with $\overline{\phi}$ being more skill-complementary than $\phi^*(\widetilde{l})$, which is what we show formally in the following.

To that end, consider a monotonic and differentiable path $l(\tau)$ from l to $\widetilde{l}$ such that $l(0) = l$,

$l(1) = \widetilde{l}$ and $F(l(\tau), \phi^*(\widetilde{l})) = F(l, \phi^*(\widetilde{l}))$ for all $\tau \in [0,1]$. Here I mean by monotonicity that each entry $l_\theta(\tau)$ (in the vector $l(\tau)$) is monotonic in τ. Applying the mean value theorem, the inequalities in (B.11) imply that there is a $\widetilde{\tau} \in (0,1)$ such that

$$\int_{\underline{\theta}}^{\overline{\theta}} D_{l_\theta} F(l(\widetilde{\tau}), \overline{\phi}) \frac{dl_\theta(\widetilde{\tau})}{d\tau}\, d\theta < 0\ . \tag{B.12}$$

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

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

we can now extend inequality (B.12) to

$$\int_{\underline{\theta}}^{\overline{\theta}} \left(D_{l_\theta} F(l(\widetilde{\tau}), \overline{\phi}) - \frac{D_{l_{\widetilde{\theta}}} F(l(\widetilde{\tau}), \overline{\phi})}{D_{l_{\widetilde{\theta}}} F(l(\widetilde{\tau}), \phi^*(\widetilde{l}))} D_{l_\theta} F(l(\widetilde{\tau}), \phi^*(\widetilde{l})) \right) \frac{dl_\theta(\widetilde{\tau})}{d\tau}\, d\theta < 0\ . \tag{B.13}$$

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

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

is also positive for $\theta > \widetilde{\theta}$ and negative for $\theta \leq \widetilde{\theta}$. This implies that the right-hand side of (B.13) must be (weakly) positive, a contradiction.

We have hence shown that $\phi^*(\widetilde{l}) \succeq^{sb} \phi^*(l)$, that is, the equilibrium technology is more skill-biased under $\widetilde{l}$ than under l. So, the increase in relative skill supply induces skill-biased technical change. The local implication of this global result is Lemma 6. □

B.1.3 Proofs and Derivations for the Tax Reform Analysis

In this section I provide all proofs and omitted derivations for the tax reform analysis in Section 3.5. I start with the proof of Lemma 7, which provides alternative characterizations of progressive tax reforms as defined by Definition 6.

Proof of Lemma 7. The strategy of the proof is to show that statement 1 implies statement 2, statement 2 implies statement 3, and statement 3 implies statement 1.

(1 ⇒ 2) Take a function r such that $R_{\widetilde{T}}(y) = r(R_T(y))$ for all y. Differentiating both sides with respect to y yields

$$r'(R_T(y))R'_T(y) = R'_{\widetilde{T}}(y) \quad \forall y \ ,$$

and after taking logs and rearranging:

$$\log r'(R_T(y)) = \log R'_{\widetilde{T}}(y) - \log R'_T(y) \quad \forall y \ .$$

Differentiating again with respect to y and multiplying through by y gives

$$\frac{r''(R_T(y))}{r'(R_T(y))} R'_T(y) y = -(P_{\widetilde{T}}(y) - P_T(y)) < 0 \quad \forall y \ ,$$

where the inequality is Definition 6. The assumption that $T'(y) < 1$ and $T'(y) + \mu\tau'(y) < 1$ for all y (see Assumption 3) implies that $r'(R_T(y)) > 0$ and $R'_T(y) > 0$, such that $r''(R_T(y)) < 0$ for all y, which is statement 2.

(2 ⇒ 3) Statement 2 implies

$$\frac{R'_{\widetilde{T}}(y)}{R'_{\widetilde{T}}(\widetilde{y})} = \frac{r'(R_T(y))}{r'(R_T(\widetilde{y}))} \frac{R'_T(y)}{R'_T(\widetilde{y})} \leq \frac{R'_T(y)}{R'_T(\widetilde{y})} \quad \forall y \geq \widetilde{Y} \ ,$$

because r is concave and R_T strictly increasing. Replacing $R'_{\widetilde{T}}(y)$ by $R'_T(y) - \mu\tau'(y)$ and rearranging yields

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

and hence:

$$\frac{\tau'(y)}{R'_T(y)} \geq \frac{\tau'(\widetilde{y})}{R'_T(\widetilde{y})} \quad \forall y \geq \widetilde{y} \ ,$$

which is statement 3.

(3 ⇒ 1) We can transform statement 3 into

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

Taking logs and rearranging yields

$$\log R'_{\widetilde{T}}(y) - \log R'_{\widetilde{T}}(\widetilde{y}) \leq \log R'_T(y) - \log R'_T(\widetilde{y}) \quad \forall y \geq \widetilde{y} \ .$$

Setting $\widetilde{y} = y - d$, dividing both sides of the equation by d, and taking the limit as $d \to 0$, we obtain

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

and hence

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

□

Labor Input Effects

First, note that equation (3.22) is easily derived by applying the derivative D_τ to labor inputs. In particular, accounting for the general equilibrium contingencies between labor supply and wages, we can write labor supply as $l_\theta(T, w_\theta)$ and wages as $w_\theta(l(T, w_\theta), \phi^*(l(T, w_\theta)))$. Then, using derivatives and elasticities as defined in the main text, it is straightforward to derive equation (3.22).

Starting from equation (3.22) I prove Lemma 8.

Proof of Lemma 8. **Step 1.** It is easy to see that

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

for all $n \geq 1$. Hence, the two expressions (3.25) and (3.26) are equal.

Step 2. Suppose for now that all the series in expressions (3.25) and (3.26) converge. Then, take expression (3.25) and insert it into the fixed point equation (3.22):

$$\begin{aligned}
\sum_{n=1}^{\infty} \widehat{l}^{(n)}_{\theta,\tau} &= -\,\epsilon^R_\theta \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} + \epsilon^w_\theta \zeta_{\theta,\theta} \sum_{n=1}^{\infty} \widehat{l}^{(n)}_{\theta,\tau} + \epsilon^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \zeta_{\theta,\widetilde{\theta}} \sum_{n=1}^{\infty} \widehat{l}^{(n)}_{\widetilde{\theta},\tau} \, d\widetilde{\theta} \\
&= -\,\epsilon^R_\theta \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} + \sum_{n=1}^{\infty} \left[\epsilon^w_\theta \zeta_{\theta,\theta} \widehat{l}^{(n)}_{\theta,\tau} + \epsilon^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \zeta_{\theta,\widetilde{\theta}} \widehat{l}^{(n)}_{\widetilde{\theta},\tau} \, d\widetilde{\theta} \right] \\
&= -\,\epsilon^R_\theta \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} + \sum_{n=1}^{\infty} \widehat{l}^{(n+1)}_{\theta,\tau} \\
&= -\,\epsilon^R_\theta \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} + \sum_{n=2}^{\infty} \widehat{l}^{(n)}_{\theta,\tau} \\
&= \sum_{n=1}^{\infty} \widehat{l}^{(n)}_{\theta,\tau} \; .
\end{aligned}$$

This proves that, conditional upon convergence of the series, expression (3.25) solves the fixed point equation (3.22). Then, by Step 1, expression (3.26) also solves the fixed point equation conditional upon convergence.

Step 3. Regarding convergence, consider expression (3.25) first. Start from the definition of $\widehat{l}^{(n)}_{\theta,\tau}$ and take the square of both sides of the equation:

$$\left(\widehat{l}^{(n)}_{\theta,\tau}\right)^2 = (\epsilon^w_\theta \zeta_{\theta,\theta})^2 \left(\widehat{l}^{(n-1)}_{\theta,\tau}\right)^2 + (\epsilon^w_\theta)^2 \left(\int_{\underline{\theta}}^{\overline{\theta}} \zeta_{\theta,\widetilde{\theta}} \widehat{l}^{(n-1)}_{\widetilde{\theta},\tau}\, d\widetilde{\theta}\right)^2 + 2\epsilon^w_\theta \zeta_{\theta,\theta} \widehat{l}^{(n-1)}_{\theta,\tau} \epsilon^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \zeta_{\theta,\widetilde{\theta}} \widehat{l}^{(n-1)}_{\widetilde{\theta},\tau}\, d\widetilde{\theta}\,.$$

By the Cauchy-Schwarz inequality,

$$\left(\widehat{l}^{(n)}_{\theta,\tau}\right)^2 \leq (\epsilon^w_\theta \zeta_{\theta,\theta})^2 \left(\widehat{l}^{(n-1)}_{\theta,\tau}\right)^2 + (\epsilon^w_\theta)^2 \int_{\underline{\theta}}^{\overline{\theta}} \zeta^2_{\theta,\widetilde{\theta}}\, d\widetilde{\theta} \int_{\underline{\theta}}^{\overline{\theta}} \left(\widehat{l}^{(n-1)}_{\widetilde{\theta},\tau}\right)^2 d\widetilde{\theta}$$
$$+ 2\epsilon^w_\theta \zeta_{\theta,\theta} \widehat{l}^{(n-1)}_{\theta,\tau} \epsilon^w_\theta \sqrt{\int_{\underline{\theta}}^{\overline{\theta}} \zeta^2_{\theta,\widetilde{\theta}}\, d\widetilde{\theta}} \sqrt{\int_{\underline{\theta}}^{\overline{\theta}} \left(\widehat{l}^{(n-1)}_{\widetilde{\theta},\tau}\right)^2 d\widetilde{\theta}}\,,$$

and after integrating over θ,

$$\int_{\underline{\theta}}^{\overline{\theta}} \left(\widehat{l}^{(n)}_{\theta,\tau}\right)^2 d\theta \leq \int_{\underline{\theta}}^{\overline{\theta}} (\epsilon^w_\theta \zeta_{\theta,\theta})^2 \left(\widehat{l}^{(n-1)}_{\theta,\tau}\right)^2 d\theta + \int_{\underline{\theta}}^{\overline{\theta}} (\epsilon^w_\theta)^2 \int_{\underline{\theta}}^{\overline{\theta}} \zeta^2_{\theta,\widetilde{\theta}}\, d\widetilde{\theta}\, d\theta \int_{\underline{\theta}}^{\overline{\theta}} \left(\widehat{l}^{(n-1)}_{\widetilde{\theta},\tau}\right)^2 d\widetilde{\theta}$$
$$+ 2\int_{\underline{\theta}}^{\overline{\theta}} \epsilon^w_\theta \zeta_{\theta,\theta} \widehat{l}^{(n-1)}_{\theta,\tau} \epsilon^w_\theta \sqrt{\int_{\underline{\theta}}^{\overline{\theta}} \zeta^2_{\theta,\widetilde{\theta}}\, d\widetilde{\theta}\, d\theta} \sqrt{\int_{\underline{\theta}}^{\overline{\theta}} \left(\widehat{l}^{(n-1)}_{\widetilde{\theta},\tau}\right)^2 d\widetilde{\theta}}\,.$$

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

$$\int_{\underline{\theta}}^{\overline{\theta}} \left(\widehat{l}^{(n)}_{\theta,\tau}\right)^2 d\theta \leq \sup_{\theta\in\Theta} \left[(\epsilon^w_\theta \zeta_{\theta,\theta})^2\right] \int_{\underline{\theta}}^{\overline{\theta}} \left(\widehat{l}^{(n-1)}_{\theta,\tau}\right)^2 d\theta + \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} \left(\epsilon^w_\theta \zeta_{\theta,\widetilde{\theta}}\right)^2 d\widetilde{\theta}\, d\theta \int_{\underline{\theta}}^{\overline{\theta}} \left(\widehat{l}^{(n-1)}_{\widetilde{\theta},\tau}\right)^2 d\widetilde{\theta}$$
$$+ 2\sqrt{\int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} \left(\epsilon^w_\theta \zeta_{\theta,\theta} \epsilon^w_\theta \zeta_{\theta,\widetilde{\theta}}\right)^2 d\widetilde{\theta}\, d\theta} \int_{\underline{\theta}}^{\overline{\theta}} \left(\widehat{l}^{(n-1)}_{\widetilde{\theta},\tau}\right)^2 d\widetilde{\theta}\,.$$

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

$$\sup_{\theta\in\Theta} (\epsilon^w_\theta \zeta_{\theta,\theta})^2 + \int_{\underline{\theta}}^{\overline{\theta}} (\epsilon^w_\theta \zeta_{\theta,\widetilde{\theta}})^2\, d\widetilde{\theta}\, d\theta + 2\sqrt{\int_{\underline{\theta}}^{\overline{\theta}} (\epsilon^w_\theta \zeta_{\theta,\theta})^2 (\epsilon^w_\theta \zeta_{\theta,\widetilde{\theta}})^2\, d\widetilde{\theta}\, d\theta}\,,$$

which is strictly smaller than one by condition (3.24). Hence, the term $\int_{\underline{\theta}}^{\overline{\theta}} \left(\widehat{l}_{\theta,\tau}^{(n-1)}\right)^2 d\theta$ is dominated by a geometric sequence converging to zero.

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

$$\widehat{l}_{\theta,\tau}^{(n)} \leq \epsilon_{\theta}^{w} \zeta_{\theta,\theta} \widehat{l}_{\theta,\tau}^{(n-1)} + \epsilon_{\theta,\tau}^{w} \sqrt{\int_{\underline{\theta}}^{\overline{\theta}} \zeta_{\theta,\widetilde{\theta}}^{2} \, d\widetilde{\theta}} \sqrt{\int_{\underline{\theta}}^{\overline{\theta}} \left(\widehat{l}_{\widetilde{\theta}}^{(n-1)}\right)^2 d\widetilde{\theta}} \; . \tag{B.14}$$

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 $c \in (\epsilon_{\theta}^{w} \zeta_{\theta,\theta}, 1)$ and for any $N \in \mathbb{N}$, there must exist $\overline{N}_N > N$ such that

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

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

$$\frac{|\epsilon_{\theta,\tau}^{w} \sqrt{\int_{\underline{\theta}}^{\overline{\theta}} \zeta_{\theta,\widetilde{\theta}}^{2} \, d\widetilde{\theta}} \sqrt{\int_{\underline{\theta}}^{\overline{\theta}} \left(\widehat{l}_{\widetilde{\theta}}^{(n-1)}\right)^2 d\widetilde{\theta}}|}{\widehat{l}_{\theta,\tau}^{(\overline{N}_N)}} \to 0 \quad \text{as } N \to \infty \; .$$

But with equation (B.14) this implies, as $N \to \infty$,

$$\frac{|\widehat{l}_{\theta,\tau}^{(\overline{N}_N)}|}{|\widehat{l}_{\theta,\tau}^{(\overline{N}_N - 1)}|} \to |\epsilon_{\theta}^{w} \zeta_{\theta,\theta}| < c \; ,$$

a contradiction.

So, $\widehat{l}_{\theta,\tau}^{(n)}$ is dominated by a geometric sequence converging to zero and the series $\sum_{n=1}^{\infty} \widehat{l}_{\theta,\tau}^{(n)}$ 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 $\zeta_{\theta,\widetilde{\theta}}$ by $\rho_{\theta,\widetilde{\theta}}$, the reasoning in step 3 implies that $\widetilde{TE}_{\theta,\tau}$ converges. Second, note that

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

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, if ϵ_θ^w is constant in θ and $\epsilon_\theta^R \tau'(y_\theta(T))/(1 - T'(y_\theta(T)))$ increases in θ, the component $\widetilde{TE}_{\theta,\tau}$ decreases in θ as well. The proof is by induction.

If $\epsilon_\theta^R \tau'(y_\theta(T))/(1 - T'(y_\theta(T)))$ increases in θ, then by Lemma 6 we have that the term

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

is decreasing in θ.

Suppose now that $\widetilde{TE}_{\theta,\tau}^{(n)}$ decreases in θ. Then, again by Lemma 6,

$$\rho_{\theta,\theta} \widetilde{TE}_{\theta,\tau}^{(n)} + \int_{\underline{\theta}}^{\overline{\theta}} \rho_{\theta,\widetilde{\theta}} \widetilde{TE}_{\theta,\tau}^{(n)} \, d\widetilde{\theta}$$

decreases in θ. If ϵ_θ^w is constant in θ, the same holds for

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

But this is equal to $\widetilde{TE}_{\theta,\tau}^{(n+1)}$. Hence inductively, $\widetilde{TE}_{\theta,\tau}^{(n)}$ decreases in θ for all $n \geq 1$. So the sum $\sum_{n=1} \widetilde{TE}_{\theta,\tau}^{(n)}$ decreases in θ as well, which yields the desired result. □

The proof of Corollary 9 requires specific results for the labor response to tax reforms that hold in the CES case. The following lemma provides these results.

Lemma 14. *Fix an initial tax T and assume that F and Φ are CES as introduced in Section 3.3.4. Moreover, let the elasticity ϵ_θ^w be constant in θ, that is, $\epsilon_\theta^w = \epsilon^w$ for all $\theta \in \Theta$. Then, the effect of tax reform τ on labor inputs can be written as*

$$\widehat{l}_{\theta,\tau}(T) = -\bar{\epsilon}_\theta^R \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} - (\gamma^{CES} + \rho^{CES}) \epsilon^w \int_{\underline{\theta}}^{\overline{\theta}} \frac{w_{\widetilde{\theta}} l_{\widetilde{\theta}} h_{\widetilde{\theta}}}{F(l(T), \phi^*(T))} (-\bar{\epsilon}_{\widetilde{\theta}}^R) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1 - T'(y_{\widetilde{\theta}}(T))} \, d\widetilde{\theta} \, , \quad \text{(B.15)}$$

where

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

So, $\widehat{l}_{\theta,\tau}(T)$ decreases in θ if and only if

$$\bar{\epsilon}^R_{\widetilde{\theta}} \frac{\tau'(y_{\widetilde{\theta}}(T))}{1 - T'(y_{\widetilde{\theta}}(T))}$$

increases in θ.

If in addition ϵ^R_θ is constant in θ, $\widehat{l}_{\theta,\tau}(T)$ decreases in θ if and only if τ is progressive.

Proof. The fastest way to prove equation (B.15) is to check that it satisfies the fixed point equation (3.22). In the CES case and with ϵ^w_θ constant, this equation becomes

$$\widehat{l}_{\theta,\tau}(T) = -\epsilon^R_\theta \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} + \epsilon^w(\gamma^{CES} + \rho^{CES})\widehat{l}_{\theta,\tau}(T) - \epsilon^w(\gamma^{CES} + \rho^{CES}) \int_{\underline{\theta}}^{\overline{\theta}} \frac{w_{\widetilde{\theta}} l_{\widetilde{\theta}} h_{\widetilde{\theta}}}{F(l(T), \phi^*(T))} \widehat{l}_{\widetilde{\theta},\tau} \, d\widetilde{\theta} \, .$$

Inserting equation (B.15) yields:

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

where the last equality follows from the fact that F is linear homogeneous in l (see Lemma 12) and Euler's homogeneous function theorem. So, equation (B.15) solves the fixed point equation (3.22).

The remainder of Lemma 14 then follows from the observation that the second term on the right-hand side of equation (B.15) is independent of θ and that, in the CES case, $\bar{\epsilon}_\theta^R$ is constant in θ if ϵ_θ^R and ϵ_θ^w are constant in θ. □

Directed Technical Change Effects

Using the labor input responses from Lemma 8, I prove the results from Section 3.5.3 on the directed technical change effects of tax reforms.

First, I obtain equation (3.27) by applying the derivative $D_{\phi,\tau}$ to wages. In particular, accounting for the general equilibrium contingencies between wages and labor supply, we can write wages as $w_\theta(l(T, w_\theta), \phi^*(l(T, w_\theta)))$ and labor supply as $l_\theta(T, w_\theta)$. Then, using derivatives and elasticities as defined in the main text, it is straightforward to derive equation (3.27).

Combining equation (3.27) with Lemma 8, we can prove Proposition 3.

Proof of Proposition 3. Equation (3.28) is obtained immediately by inserting equation (3.26) from Lemma 8 into (3.27).

To sign the slopes of $DE_{\theta,\tau}$ and $TE_{\theta,\tau}$, note that by Lemma 6 the induced technical change effect

$$\rho_{\theta,\theta}\widehat{l}_{\theta,\tau}(T) + \int_{\underline{\theta}}^{\overline{\theta}} \rho_{\theta,\widetilde{\theta}}\widehat{l}_{\widetilde{\theta},\tau}(T)\, d\widetilde{\theta}$$

decreases in θ if $\widehat{l}_{\theta,\tau}(T)$ decreases in θ. This immediately implies that $DE_{\theta,\tau}$ decreases in θ if $\epsilon_\theta^R \tau'(y_\theta(T))/(1 - T'(y_\theta(T)))$ increases in θ. Moreover, Lemma 8 says that $\widetilde{TE}_{\theta,\tau}$ decreases in θ if ϵ_θ^w is constant in θ and $\epsilon_\theta^R \tau'(y_\theta(T))/(1 - T'(y_\theta(T)))$ increases in θ. So, under these conditions also $TE_{\theta,\tau}$ decreases in θ. □

Corollary 9 gives the directed technical change effects of a reform τ in the CES case. I prove this by applying Lemma 14.

Proof of Corollary 9. Since aggregate production F and the equilibrium aggregate production function F^* are linear homogeneous in l (see Lemmas 12 and 13), the induced technical change effects of a proportional change in all types' labor inputs are zero:

$$\rho_{\theta,\theta}\widehat{l}_{\theta,\tau}(T) + \int_{\underline{\theta}}^{\overline{\theta}} \rho_{\theta,\widetilde{\theta}}\widehat{l}_{\widetilde{\theta},\tau}(T)\, d\widetilde{\theta} = 0 \quad \text{for all } \theta$$

if $\widehat{l}_{\theta,\tau}(T)$ is constant in θ. Hence, inserting equation (B.15) into equation (3.27), the second term of equation (B.15) vanishes. This leaves

$$\frac{1}{w_\theta} D_{\phi,\tau} w_\theta(T, \phi^*(T)) = \rho^{CES}(-\bar{\epsilon}_\theta^R) \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} - \rho^{CES} \int_{\underline{\theta}}^{\overline{\theta}} \frac{w_{\widetilde{\theta}} l_{\widetilde{\theta}} h_{\widetilde{\theta}}}{F(l(T), \phi^*(T))} (-\bar{\epsilon}_{\widetilde{\theta}}^R) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1 - T'(y_{\widetilde{\theta}}(T))} \, d\widetilde{\theta} \,,$$

which is equation (3.29).

The second term on the right-hand side of equation (3.29) is independent of θ. This immediately implies that the relative wage change is decreasing in θ if $\bar{\epsilon}_\theta^R \tau'(y_\theta(T))/(1 - T'(y_\theta(T)))$ increases in θ.

Finally, if ϵ_θ^R is constant in θ (in addition to ϵ_θ^w, which is required by Corollary 9 anyway), $\bar{\epsilon}_\theta^R$ is constant in θ as well, and the relative wage change decreases in θ for any progressive reform. □

Within-Technology Substitution Effects

Proposition 3 provides a general formula for the directed technical change effects of tax reforms on wages. Here, I state its counterpart for within-technology substitution effects. I use this when computing the total wage effects of tax reforms in the quantitative analysis in Section 3.7.

Proposition 11. *Fix an initial tax T and let conditions* (3.23) *and* (3.24) *be satisfied.*

Then, the relative effect of the within-technology factor substitution induced by tax reform τ on wages can be written as

$$\begin{aligned} \frac{1}{w_\theta} D_\tau w_\theta(T, \phi^*(T)) =& \gamma_{\theta,\theta}(-\epsilon_\theta^R) \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} + \int_{\underline{\theta}}^{\overline{\theta}} \gamma_{\theta,\widetilde{\theta}}(-\epsilon_{\widetilde{\theta}}^R) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1 - T'(y_{\widetilde{\theta}}(T))} \, d\widetilde{\theta} \\ &+ \gamma_{\theta,\theta} \widetilde{TE}_{\theta,\tau}(T) + \int_{\underline{\theta}}^{\overline{\theta}} \gamma_{\theta,\widetilde{\theta}} \widetilde{TE}_{\widetilde{\theta},\tau}(T) \, d\widetilde{\theta} \\ &+ \gamma_{\theta,\theta} \widetilde{SE}_{\theta,\tau}(T) + \int_{\underline{\theta}}^{\overline{\theta}} \gamma_{\theta,\widetilde{\theta}} \widetilde{SE}_{\widetilde{\theta},\tau}(T) \, d\widetilde{\theta} \,, \end{aligned} \tag{B.16}$$

for all $\theta \in \Theta$, where $\widetilde{TE}_{\theta,\tau}(T)$ and $\widetilde{SE}_{\theta,\tau}(T)$ are defined in Lemma 8.

Proof. Analogously to the induced technical change effects in equation (3.27), the substitution effects of tax reform τ on wages can be written as

$$\frac{1}{w_\theta} D_\tau w_\theta(T, \phi^*(T)) = \gamma_{\theta,\theta} \widehat{l}_{\theta,\tau}(T) + \int_{\underline{\theta}}^{\overline{\theta}} \gamma_{\theta,\widetilde{\theta}} \widehat{l}_{\widetilde{\theta},\tau}(T) \, d\widetilde{\theta} \,. \tag{B.17}$$

Replacing equation (3.27) by equation (B.17), the proof proceeds analogously to the proof of equation (3.28) in Proposition 3 and is therefore omitted. □

In the CES case, the within-technology substitution effects can be expressed as follows.

Corollary 12. *Fix an initial tax T and assume that F and Φ are CES as introduced in Section 3.3.4. Moreover, let the elasticity ϵ_θ^w be constant in θ, that is, $\epsilon_\theta^w = \epsilon^w$ for all $\theta \in \Theta$. Then the relative wage effect of the within-technology factor substitution induced by tax reform τ satisfies*

$$\frac{1}{w_\theta} D_\tau w_\theta(T, \phi^*(T)) = \gamma^{CES}(-\bar{\epsilon}_\theta^R)\frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} - \gamma^{CES}\int_{\underline{\theta}}^{\overline{\theta}} \frac{w_{\widetilde{\theta}} l_{\widetilde{\theta}} h_{\widetilde{\theta}}}{F(l(T), \phi^*(T))}(-\bar{\epsilon}_{\widetilde{\theta}}^R)\frac{\tau'(y_{\widetilde{\theta}}(T))}{1 - T'(y_{\widetilde{\theta}}(T))}\, d\widetilde{\theta} \quad \text{(B.18)}$$

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

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

Proof. Again replacing equation (3.27) by equation (B.17), the proof of Corollary 12 is analogous to the proof of its counterpart for induced technical change effects, Corollary 9. □

B.1.4 Proofs and Derivations for Optimal Taxes

This section contains all proofs and omitted derivations for the analysis of optimal taxes in Section 3.6 of the main text.

General Case

In the general case, optimal taxes are obtained by maximizing wefare $W(c, l)$ subject to the resource constraint (3.30) and the incentive compatibility constraint (3.31). 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 tax rates into the equations, and finally prove the sign conditions for the term TE_θ^* at the bottom and the top of the type distribution using directed technical change theory.

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

Lemma 15. *The pair of consumption and labor inputs (c, l) satisfies the resource and incentive compatibility constraints* (3.30) *and* (3.31) *if and only if $c = c^*(l)$, where $c^*(l) = \{c_\theta^*(l)\}_{\theta\in\Theta}$ is* *determined by*

$$c^*_\theta(l) = F(l,\phi^*(l)) - \int_{\underline{\theta}}^{\overline{\theta}} v(l_{\tilde\theta}) h_{\tilde\theta}\, d\tilde\theta - \int_{\underline{\theta}}^{\overline{\theta}} v'(l_{\tilde\theta}) l_{\tilde\theta} (1 - H_{\tilde\theta}) \widehat{w}_{\tilde\theta}\, d\tilde\theta + \int_{\underline{\theta}}^{\theta} v'(l_{\tilde\theta}) l_{\tilde\theta} \widehat{w}_{\tilde\theta}\, d\tilde\theta + v(l_\theta) \tag{B.19}$$

for all θ.

Proof. (⇒) I first show that constraints (3.30) and (3.31) imply equation (B.19). 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 (3.31), this implies for all θ:

$$\begin{aligned} c_\theta &= c_{\underline{\theta}} + \int_{\underline{\theta}}^{\theta} v'(l_{\tilde\theta})(w'_{\tilde\theta} l_{\tilde\theta} + w_{\tilde\theta} l'_{\tilde\theta}) \frac{1}{w_{\tilde\theta}}\, d\tilde\theta \\ &= c_{\underline{\theta}} + \int_{\underline{\theta}}^{\theta} v'(l_{\tilde\theta}) \widehat{w}_{\tilde\theta} l_{\tilde\theta}\, d\tilde\theta + v(l_\theta) - v(l_{\underline{\theta}})\ . \end{aligned} \tag{B.20}$$

Combining this with the resource constraint (3.30), we obtain the following expression for $c_{\underline{\theta}}$:

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

Using integration by parts to solve the double integral yields:

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

Inserting this back into equation (B.20), we obtain:

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

which is equation (B.19) defining the function c^* above.

(⇐) Differentiating c^* with respect to θ shows immediately that equation (B.19) implies the incentive compatibility constraint (3.31). Similarly, after multiplying c^*_θ by h_θ and integrating over $[\underline{\theta}, \overline{\theta}]$, standard computations show that

$$\int_{\underline{\theta}}^{\overline{\theta}} c^*_\theta(l)\, d\theta = F(l,\phi^*(l))\ ,$$

which proves that equation (B.19) also implies the resource constraint (3.30). □

Proof of Proposition 4. Lemma 15 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 4.

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, the derivative can be written as

$$\begin{aligned} D_{l_\theta}\widehat{W}(l) = {} & w_\theta h_\theta - v'(l_\theta)h_\theta \\ & - \left(v''(l_\theta)l_\theta + v'(l_\theta)\right)(1 - H_\theta)\widehat{w}_\theta + \widetilde{g}_\theta(1 - H_\theta)\left(v''(l_\theta)l_\theta + v'(l_\theta)\right) \\ & - \lim_{\Delta\to 0}\frac{1}{\Delta}\int_{\underline{\theta}}^{\overline{\theta}} v'(l_{\widetilde{\theta}})l_{\widetilde{\theta}}(1 - H_{\widetilde{\theta}})\left.\frac{d\widehat{w}^*_{\widetilde{\theta}}(l + \mu\widetilde{l}_{\Delta,\theta})}{d\mu}\right|_{\mu=0} d\widetilde{\theta} \\ & + \lim_{\Delta\to 0}\frac{1}{\Delta}\int_{\underline{\theta}}^{\overline{\theta}} g_{\widetilde{\theta}}h_{\widetilde{\theta}}\int_{\underline{\theta}}^{\widetilde{\theta}} v'(l_{\widehat{\theta}})l_{\widehat{\theta}}\left.\frac{d\widehat{w}^*_{\widehat{\theta}}(l + \mu\widetilde{l}_{\Delta,\theta})}{d\mu}\right|_{\mu=0} d\widehat{\theta}\, d\widetilde{\theta} \end{aligned} \tag{B.21}$$

for all θ, where the terms in the first two lines were derived following the procedure detailed in Sections B.1.1 and B.1.1, which uses continuity of l and $\widehat{w}$ in θ.[2] Following the notation introduced in Section 3.3.2, the expression $d\widehat{w}^*_{\widetilde{\theta}}(l + \mu\widetilde{l}_{\Delta,\theta})/d\mu\big|_{\mu=0}$ denotes the total derivative of $\widehat{w}$ in the direction of l_θ, accounting both for the substitution and the induced technical change effects:

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

Using integration by parts to solve the double integral in equation (B.21), the derivative of the

[2] The wage growth function $\widehat{w}$ is continuous in θ because l is C^1 by hypothesis of Proposition 4.

welfare function becomes

$$\begin{aligned}
D_{l_\theta}\widehat{W}(l) &= w_\theta h_\theta - v'(l_\theta)h_\theta - (1-\widetilde{g}_\theta)\left(v''(l_\theta)l_\theta + v'(l_\theta)\right)(1-H_\theta)\widehat{w}_\theta \\
&\quad - \lim_{\Delta\to 0}\frac{1}{\Delta}\int_{\underline{\theta}}^{\overline{\theta}} v'(l_{\widetilde{\theta}})l_{\widetilde{\theta}}(1-H_{\widetilde{\theta}})\left.\frac{d\widehat{w}^*_{\widetilde{\theta}}(l+\mu\widetilde{l}_{\Delta,\theta})}{d\mu}\right|_{\mu=0} d\widetilde{\theta} \\
&\quad + \lim_{\Delta\to 0}\frac{1}{\Delta}\int_{\underline{\theta}}^{\overline{\theta}} \widetilde{g}_{\widetilde{\theta}}(1-H_{\widetilde{\theta}})v'(l_{\widetilde{\theta}})l_{\widetilde{\theta}}\left.\frac{d\widehat{w}^*_{\widetilde{\theta}}(l+\mu\widetilde{l}_{\Delta,\theta})}{d\mu}\right|_{\mu=0} d\widetilde{\theta} \\
&= w_\theta h_\theta - v'(l_\theta)h_\theta - (1-\widetilde{g}_\theta)\left(v''(l_\theta)l_\theta + v'(l_\theta)\right)(1-H_\theta)\widehat{w}_\theta \\
&\quad - \lim_{\Delta\to 0}\frac{1}{\Delta}\int_{\underline{\theta}}^{\overline{\theta}} (1-\widetilde{g}_{\widetilde{\theta}})v'(l_{\widetilde{\theta}})l_{\widetilde{\theta}}(1-H_{\widetilde{\theta}})\left.\frac{d\widehat{w}^*_{\widetilde{\theta}}(l+\mu\widetilde{l}_{\Delta,\theta})}{d\mu}\right|_{\mu=0} d\widetilde{\theta}\,. \qquad \text{(B.22)}
\end{aligned}$$

We now use workers' first-order condition (3.2) to introduce marginal tax rates into the equation. In particular condition (3.2) implies

$$v'(l_\theta)l_\theta = (1-T'(y_\theta))y_\theta \qquad \text{(B.23)}$$

and

$$v''(l_\theta)l_\theta\widehat{w}_\theta + v'(l_\theta)\widehat{w}_\theta = \left(1+\frac{1}{e_\theta}\right)(1-T'(y_\theta))w'_\theta\,. \qquad \text{(B.24)}$$

Using equations (3.2), (B.23), and (B.24) in equation (B.22), we obtain

$$\begin{aligned}
D_{l_\theta}\widehat{W}(l) &= T'(y_\theta)y_\theta h_\theta - (1-T'(y_\theta))\left(1+\frac{1}{e_\theta}\right)(1-\widetilde{g}_\theta)(1-H_\theta)w'_\theta l_\theta \\
&\quad - \lim_{\Delta\to 0}\frac{1}{\Delta}\int_{\underline{\theta}}^{\overline{\theta}} (1-T'(y_{\widetilde{\theta}}))y_{\widetilde{\theta}}(1-\widetilde{g}_{\widetilde{\theta}})(1-H_{\widetilde{\theta}})\left.\frac{d\widehat{w}^*_{\widetilde{\theta}}(l+\mu\widetilde{l}_{\Delta,\theta})}{d\mu}\right|_{\mu=0} d\widetilde{\theta}\,.
\end{aligned}$$

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

$$\begin{aligned}
D_{l_\theta}\widehat{W}(l) &= T'(y_\theta)w_\theta h_\theta - (1-T'(y_\theta))\left(1+\frac{1}{e_\theta}\right)(1-\widetilde{g}_\theta)(1-H_\theta)w'_\theta \\
&\quad - \lim_{\Delta\to 0}\frac{1}{\Delta}\int_{\underline{\theta}}^{\overline{\theta}} (1-T'(y_{\widetilde{\theta}}))y_{\widetilde{\theta}}(1-\widetilde{g}_{\widetilde{\theta}})(1-H_{\widetilde{\theta}})\left.\frac{d\widehat{w}_{\widetilde{\theta}}(l+\mu\widetilde{l}_{\Delta,\theta},\phi^*(l))}{d\mu}\right|_{\mu=0} d\widetilde{\theta} \\
&\quad - \lim_{\Delta\to 0}\frac{1}{\Delta}\int_{\underline{\theta}}^{\overline{\theta}} (1-T'(y_{\widetilde{\theta}}))y_{\widetilde{\theta}}(1-\widetilde{g}_{\widetilde{\theta}})(1-H_{\widetilde{\theta}})\left.\frac{d\widehat{w}_{\widetilde{\theta}}(l,\phi^*(l+\mu\widetilde{l}_{\Delta,\theta}))}{d\mu}\right|_{\mu=0} d\widetilde{\theta}\,.
\end{aligned}$$

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

$$\begin{aligned}\frac{T'(y_\theta)}{1-T'(y_\theta)} &= \left(1+\frac{1}{e_\theta}\right)(1-\widetilde{g}_\theta)\frac{1-H_\theta}{h_\theta}\widehat{w}_\theta \\ &- \lim_{\Delta\to 0}\frac{1}{\Delta}\int_{\underline{\theta}}^{\overline{\theta}} \frac{1-T'(y_{\widetilde{\theta}})}{1-T'(y_\theta)}\frac{1-H_{\widetilde{\theta}}}{h_\theta w_\theta}(1-\widetilde{g}_{\widetilde{\theta}})y_{\widetilde{\theta}} \left.\frac{d\widehat{w}_{\widetilde{\theta}}(l+\mu\widetilde{l}_{\Delta,\theta},\phi^*(l))}{d\mu}\right|_{\mu=0} d\widetilde{\theta} \\ &- \lim_{\Delta\to 0}\frac{1}{\Delta}\int_{\underline{\theta}}^{\overline{\theta}} \frac{1-T'(y_{\widetilde{\theta}})}{1-T'(y_\theta)}\frac{1-H_{\widetilde{\theta}}}{h_\theta w_\theta}(1-\widetilde{g}_{\widetilde{\theta}})y_{\widetilde{\theta}} \left.\frac{d\widehat{w}_{\widetilde{\theta}}(l,\phi^*(l+\mu\widetilde{l}_{\Delta,\theta}))}{d\mu}\right|_{\mu=0} d\widetilde{\theta}\,.\end{aligned} \tag{B.25}$$

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 4:

$$\begin{aligned}\frac{T'(y_\theta)}{1-T'(y_\theta)} &= \left(1+\frac{1}{e_\theta}\right)(1-\widetilde{g}_\theta)\frac{1-N_{w_\theta}}{n_{w_\theta} w_\theta} \\ &- \lim_{\Delta\to 0}\frac{1}{\Delta}\int_{\underline{\theta}}^{\overline{\theta}} \frac{1-T'(y_{\widetilde{\theta}})}{1-T'(y_\theta)}\frac{1-H_{\widetilde{\theta}}}{h_\theta w_\theta}(1-\widetilde{g}_{\widetilde{\theta}})y_{\widetilde{\theta}} \left.\frac{d\widehat{w}_{\widetilde{\theta}}(l+\mu\widetilde{l}_{\Delta,\theta},\phi^*(l))}{d\mu}\right|_{\mu=0} d\widetilde{\theta} \\ &- \lim_{\Delta\to 0}\frac{1}{\Delta}\int_{\underline{\theta}}^{\overline{\theta}} \frac{1-T'(y_{\widetilde{\theta}})}{1-T'(y_\theta)}\frac{1-H_{\widetilde{\theta}}}{h_\theta w_\theta}(1-\widetilde{g}_{\widetilde{\theta}})y_{\widetilde{\theta}} \left.\frac{d\widehat{w}_{\widetilde{\theta}}(l,\phi^*(l+\mu\widetilde{l}_{\Delta,\theta}))}{d\mu}\right|_{\mu=0} d\widetilde{\theta}\,.\end{aligned}$$

Part 2. The second part of the proof is to show that $TE^*_{\underline{\theta}} \leq 0$ and $TE^*_{\overline{\theta}} \geq 0$. We only consider $TE^*_{\overline{\theta}}$ because the proof for $TE^*_{\underline{\theta}}$ works analogously.

Consider first the derivative $d\widehat{w}_{\widetilde{\theta}}(l,\phi^*(l+\mu\widetilde{l}_{\Delta,\overline{\theta}}))/d\mu\Big|_{\mu=0}$. It measures the local induced technical change effect of the labor input change $\widetilde{l}_{\Delta,\overline{\theta}}$ (defined in Section 3.3.2) on relative wages.

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

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

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

$$\frac{1}{\widetilde{l}_{\Delta,\overline{\theta},\theta}}\frac{d\widetilde{l}_{\Delta,\overline{\theta},\theta}}{d\theta} \geq \frac{1}{l_\theta}\frac{dl_\theta}{d\theta}\,.^{3}$$

[3] Here we use that, by hypothesis, $\limsup_{\theta\to\overline{\theta}} l'_\theta < \infty$.

So, for $\Delta < \epsilon$, the relative labor input change $\widetilde{l}_{\Delta,\overline{\theta},\theta}/l_\theta$ increases in θ. Thus, by Lemma 6, we obtain that

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

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

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

and therefore

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

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

Note first that $SE^*_{\overline{\theta}} + TE^*_{\overline{\theta}}$ can be written as

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

We have already shown in part 2 that, for sufficiently small Δ, the relative labor input change $\widetilde{l}_{\Delta,\overline{\theta},\theta}/l_\theta$ increases in θ. By definition of strong relative bias of technology (see equation (3.20)), this implies

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

for all θ. Analogously to part 2, it follows that $SE^*_{\overline{\theta}} + TE^*_{\overline{\theta}} \geq 0$. □

Proof of Corollary 10. The limit expression for the optimal marginal tax rate follows immediately from the convergence assumptions in the corollary. 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

$$\begin{aligned}
&\int_{\theta}^{\overline{\theta}} (1 - T'(y_{\widetilde{\theta}}))h_{\widetilde{\theta}} w_{\widetilde{\theta}} l_{\widetilde{\theta}} TE^*_{\widetilde{\theta}} \, d\widetilde{\theta} \\
&= \lim_{\Delta \to 0} \frac{1}{\Delta} \int_{\underline{\theta}}^{\overline{\theta}} (1 - T'(y_{\widehat{\theta}}))(1 - H_{\widehat{\theta}})(1 - \widetilde{g}_{\widehat{\theta}}) y_{\widehat{\theta}} \int_{\theta}^{\overline{\theta}} l_{\widetilde{\theta}} \left. \frac{d\widehat{w}_{\widehat{\theta}}(l, \phi^*(l + \mu \widetilde{l}_{\Delta, \widetilde{\theta}}))}{d\mu} \right|_{\mu=0} d\widetilde{\theta} \, d\widehat{\theta} \\
&= \int_{\underline{\theta}}^{\overline{\theta}} (1 - T'(y_{\widehat{\theta}}))(1 - H_{\widehat{\theta}})(1 - \widetilde{g}_{\widehat{\theta}}) y_{\widehat{\theta}} \left. \frac{d\widehat{w}_{\widehat{\theta}}(l, \phi^*(l + \mu dl))}{d\mu} \right|_{\mu=0} d\widehat{\theta} \, ,
\end{aligned}$$

where the labor input change dl is given by

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

Since $dl_{\widetilde{\theta}}/l_{\widetilde{\theta}}$ increases in $\widetilde{\theta}$, that is, dl is an increase in relative skill supply, we have by Lemma 6:

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

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

$$\int_{\theta}^{\overline{\theta}} (1 - T'(y_{\widetilde{\theta}}))h_{\widetilde{\theta}} w_{\widetilde{\theta}} l_{\widetilde{\theta}} TE^*_{\widetilde{\theta}} \, d\widetilde{\theta} \geq 0$$

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

$$\begin{aligned}
0 &\leq \lim_{\theta \to \overline{\theta}} \frac{\int_{\theta}^{\overline{\theta}} (1 - T'(y_{\widetilde{\theta}}))h_{\widetilde{\theta}} w_{\widetilde{\theta}} l_{\widetilde{\theta}} TE^*_{\widetilde{\theta}} \, d\widetilde{\theta}}{\int_{\theta}^{\overline{\theta}} (1 - T'(y_{\widetilde{\theta}}))h_{\widetilde{\theta}} w_{\widetilde{\theta}} l_{\widetilde{\theta}} \, d\widetilde{\theta}} \\
&= \lim_{\theta \to \overline{\theta}} \frac{-(1 - T'(y_{\theta}))h_{\theta} w_{\theta} l_{\theta} TE^*_{\theta}}{-(1 - T'(y_{\theta}))h_{\theta} w_{\theta} l_{\theta}} \\
&= \lim_{\theta \to \overline{\theta}} TE^*_{\theta} \, ,
\end{aligned}$$

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

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

$$\int_{\theta}^{\overline{\theta}} (1 - T'(y_{\widetilde{\theta}}))h_{\widetilde{\theta}} w_{\widetilde{\theta}} l_{\widetilde{\theta}} \left(TE^*_{\widetilde{\theta}} + SE^*_{\widetilde{\theta}} \right) \, d\widetilde{\theta} \geq 0$$

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

$$\lim_{\theta \to \bar{\theta}} (TE^*_\theta + SE^*_\theta) \geq 0 \ .$$

□

CES Case

To derive expression (3.33) for optimal tax rates in Proposition 5, I start by specializing the terms TE^*_θ and SE^*_θ to the CES case.

Lemma 16. *Suppose the conditions of Proposition 4 are satisfied and F and Φ take the CES form introduced in Section 3.3.4. Then, the terms TE^*_θ and SE^*_θ take the following form for every θ:*

$$SE^*_\theta = (1 - g_\theta)\gamma^{CES} - \left(1 + \frac{1}{e_\theta}\right) \frac{1 - H_\theta}{h_\theta}(1 - \widetilde{g}_\theta)\widehat{l}_\theta \gamma^{CES} \tag{B.26}$$

$$TE^*_\theta = (1 - g_\theta)\rho^{CES} - \left(1 + \frac{1}{e_\theta}\right) \frac{1 - H_\theta}{h_\theta}(1 - \widetilde{g}_\theta)\widehat{l}_\theta \rho^{CES} \ . \tag{B.27}$$

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 derivative of $\widehat{w}_{\widetilde{\theta}}$ in TE^*_θ. From equation (3.14) we obtain

$$\widehat{w}_{\widetilde{\theta}} = \frac{\sigma - 1}{\sigma}\widehat{\kappa}_{\widetilde{\theta}} + \frac{\sigma - 1}{\sigma}\widehat{\phi}_{\widetilde{\theta}} - \frac{1}{\sigma}\widehat{l}_{\widetilde{\theta}} - \frac{1}{\sigma}\widehat{h}_{\widetilde{\theta}} \tag{B.28}$$

and from equation (B.6):

$$\widehat{\phi}^*_{\widetilde{\theta}} = \frac{\sigma - 1}{(\delta - 1)\sigma + 1}\left(\widehat{\kappa}_{\widetilde{\theta}} + \widehat{l}_{\widetilde{\theta}} + \widehat{h}_{\widetilde{\theta}}\right) \ . \tag{B.29}$$

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

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

Analogously, the induced technical change effect is given by

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

Using the last expression and the definition of $\widetilde{l}_{\Delta,\theta}$, the term TE^*_θ becomes

$$TE^*_\theta = \lim_{\Delta\to 0}\frac{1}{\Delta^2}\int_{\theta-\Delta}^{\theta}\frac{1-T'(y_{\widetilde{\theta}})}{1-T'(y_\theta)}\frac{1-H_{\widetilde{\theta}}}{h_\theta w_\theta}(1-\widetilde{g}_{\widetilde{\theta}})w_{\widetilde{\theta}}\rho^{CES}\left(1-\widehat{l}_{\widetilde{\theta}}(\widetilde{\theta}-\theta+\Delta)\right)d\widetilde{\theta}$$
$$+\frac{1}{\Delta^2}\int_{\theta}^{\theta+\Delta}\frac{1-T'(y_{\widetilde{\theta}})}{1-T'(y_\theta)}\frac{1-H_{\widetilde{\theta}}}{h_\theta w_\theta}(1-\widetilde{g}_{\widetilde{\theta}})w_{\widetilde{\theta}}\rho^{CES}\left(-1-\widehat{l}_{\widetilde{\theta}}(\theta-\widetilde{\theta}+\Delta)\right)d\widetilde{\theta} .$$

Applying L'Hôspital's rule yields:

$$TE^*_\theta = \lim_{\Delta\to 0}\frac{1}{2\Delta}\frac{1-T'(y_{\theta-\Delta})}{1-T'(y_\theta)}\frac{1-H_{\theta-\Delta}}{h_\theta w_\theta}(1-\widetilde{g}_{\theta-\Delta})w_{\theta-\Delta}\rho^{CES}$$
$$-\frac{1}{2\Delta}\int_{\theta-\Delta}^{\theta}\frac{1-T'(y_{\widetilde{\theta}})}{1-T'(y_\theta)}\frac{1-H_{\widetilde{\theta}}}{h_\theta w_\theta}(1-\widetilde{g}_{\widetilde{\theta}})w_{\widetilde{\theta}}\rho^{CES}\widehat{l}_{\widetilde{\theta}}\,d\widetilde{\theta}$$
$$-\frac{1}{2\Delta}\frac{1-T'(y_{\theta+\Delta})}{1-T'(y_\theta)}\frac{1-H_{\theta+\Delta}}{h_\theta w_\theta}(1-\widetilde{g}_{\theta+\Delta})w_{\theta+\Delta}\rho^{CES}$$
$$-\frac{1}{2\Delta}\int_{\theta}^{\theta+\Delta}\frac{1-T'(y_{\widetilde{\theta}})}{1-T'(y_\theta)}\frac{1-H_{\widetilde{\theta}}}{h_\theta w_\theta}(1-\widetilde{g}_{\widetilde{\theta}})w_{\widetilde{\theta}}\rho^{CES}\widehat{l}_{\widetilde{\theta}}\,d\widetilde{\theta} .$$

Rearranging and replacing marginal retention rates by workers' first-order condition (3.2), we obtain:

$$TE^*_\theta = \lim_{\Delta\to 0}\frac{1}{2\Delta}\rho^{CES}\left[\frac{v'(l_{\theta-\Delta}}{v'(l_\theta)}\frac{1-H_{\theta-\Delta}}{h_\theta}(1-\widetilde{g}_{\theta-\Delta}) - \frac{v'(l_{\theta+\Delta}}{v'(l_\theta)}\frac{1-H_{\theta+\Delta}}{h_\theta}(1-\widetilde{g}_{\theta+\Delta})\right]$$
$$-\frac{1}{2\Delta}\int_{\theta-\Delta}^{\theta}\frac{v'(l_{\widetilde{\theta}}}{v'(l_\theta)}\frac{1-H_{\widetilde{\theta}}}{h_\theta}(1-\widetilde{g}_{\widetilde{\theta}})\rho^{CES}\widehat{l}_{\widetilde{\theta}}\,d\widetilde{\theta}$$
$$-\frac{1}{2\Delta}\int_{\theta}^{\theta+\Delta}\frac{v'(l_{\widetilde{\theta}}}{v'(l_\theta)}\frac{1-H_{\widetilde{\theta}}}{h_\theta}(1-\widetilde{g}_{\widetilde{\theta}})\rho^{CES}\widehat{l}_{\widetilde{\theta}}\,d\widetilde{\theta} .$$

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

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

Using the definition of the elasticity e_θ yields:

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

Finally, it is straightforward to show that

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

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

$$TE^*_\theta = -\frac{1}{e_\theta}\frac{1-H_\theta}{h_\theta}(1-\widetilde{g}_\theta)\widehat{l}_\theta\rho^{CES} + (1-\widetilde{g}_\theta)\rho^{CES} + (\widetilde{g}_\theta - g_\theta)\rho^{CES} - \frac{1-H_\theta}{h_\theta}(1-\widetilde{g}_\theta)\widehat{l}_\theta\rho^{CES}$$

and after rearranging:

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

which is the desired expression. □

Besides providing an important step in the derivation of equation (3.33), Lemma 16 allows to revisit the sign of TE^*_θ at the bottom and the top of the type distribution. In the general case, Proposition 4 shows that the directed technical change term is weakly positive at the top and weakly negative at the bottom. For the CES case, Lemma 16 implies[4]

$$TE^*_{\overline{\theta}} = (1-\widetilde{g}_\theta)\rho^{CES} > 0$$

and

$$TE^*_{\underline{\theta}} = (1-g_{\underline{\theta}})\rho^{CES} < 0 \ .$$

Hence, in the CES case the sign restrictions on the directed technical change term hold strictly. Moreover, Lemma 16 implies the opposite signs for the substitution term at the top and bottom types:

$$SE^*_{\overline{\theta}} = (1-\widetilde{g}_\theta)\gamma^{CES} > 0$$

and

$$SE^*_{\underline{\theta}} = (1-g_{\underline{\theta}})\gamma^{CES} < 0 \ .$$

[4]This again assumes that $\limsup_{\theta\to\overline{\theta}} l'_\theta < \infty$ and $\liminf_{\theta\to\underline{\theta}} l'_\theta > -\infty$ under the optimal tax, as in the second part of Proposition 4. Moreover, the strict inequalities below require that marginal welfare weights are strictly decreasing at the optimum over parts of type distribution.

Hence, at the highest and lowest income levels, directed technical change and within-technology substitution effects push optimal marginal tax rates in opposing directions.

With the expressions from Lemma 16, we are now in a position to derive equation (3.33).

Proof of Proposition 5. We start with equation (B.25) from the proof of Proposition 4 and replace SE^*_θ and TE^*_θ by the expressions from Lemma 16. This yields:

$$\frac{T'(y_\theta)}{1-T'(y_\theta)} = \left(1+\frac{1}{e_\theta}\right)\frac{1-H_\theta}{h_\theta}(1-\widetilde{g}_\theta)\widehat{w}_\theta \tag{B.30}$$

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

$$= \left(1+\frac{1}{e_\theta}\right)\frac{1-H_\theta}{h_\theta}(1-\widetilde{g}_\theta)\left(\widehat{w}_\theta - (\gamma^{CES}+\rho^{CES})\widehat{l}_\theta\right) \tag{B.31}$$

$$+ \gamma^{CES}(1-g_\theta) + \rho^{CES}(1-g_\theta)\ . \tag{B.32}$$

The wage growth rate $\widehat{w}_\theta$ can be computed from equations (B.28) and (B.29) as

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

Using this in the previous expression for marginal tax rates yields

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

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

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

and hence

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

Finally, with the change-of-variable $h_\theta = b_{\beta_\theta}\beta'_\theta$, we obtain equation (3.33):

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

□

Finally, we prove Corollary 11, which characterizes the optimal asymptotic tax in the Pareto tail of the income distribution for the CES case.

Proof of Corollary 11. The corollary starts from the assumption that, under some initial tax $\overline{T}$ with constant marginal top tax rate, the income distribution has the Pareto property (i.e., its inverse hazard ratio is 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 (3.33) from Proposition 5 then yields the desired result.

First, by two changes-of-variable we obtain

$$\frac{1-B_{\beta_\theta}}{b_{\beta_\theta}\beta_\theta} = \frac{1-H_\theta}{h_\theta}\widehat{\beta}_\theta = \frac{1-M_{y_\theta}}{m_{y_\theta}y_\theta}\frac{\widehat{\beta}_\theta}{\widehat{y}_\theta}\ , \tag{B.33}$$

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

$$\widehat{y}_\theta = \widehat{w}_\theta + \widehat{l}_\theta = (1+\epsilon^w_\theta)\widehat{w}_\theta$$

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

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

It follows that

$$\frac{\widehat{\beta}_\theta}{\widehat{y}_\theta} = \frac{1-(\gamma^{CES}+\rho^{CES})\epsilon^w_\theta}{1+\epsilon^w_\theta}$$

and, with equation (B.33),

$$\frac{1-B_{\beta_\theta}}{b_{\beta_\theta}\beta_\theta} = \frac{1-M_{y_\theta}}{m_{y_\theta}y_\theta}\frac{1-(\gamma^{CES}+\rho^{CES})\epsilon^w_\theta}{1+\epsilon^w_\theta}\ , \tag{B.34}$$

where incomes and the labor supply elasticity ϵ_θ^w are assessed under the tax $\overline{T}$. In particular,

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

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

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

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

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

Combining these limits, we obtain

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

In words, from the observed Pareto tail of the income distribution under tax $\overline{T}$ 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 (3.33) from Proposition 5, we obtain the following expression for the optimal marginal tax rate in the upper tail of the income distribution:

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

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

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 3.3.1 with the exception of the condition for equilibrium technology (3.8). Instead of following equation (3.8), the exogenous technology planner believes that technology remains fixed at its equilibrium value under a given tax $\overline{T}$, $\phi^*(l(\overline{T}))$. The idea is that the planner observes the economy under the tax $\overline{T}$ when computing optimal taxes and believes technology to be exogenous.

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

Proposition 12. *Suppose the conditions of Proposition 4 are satisfied and F and Φ take the CES form introduced in Section 3.3.4. Suppose equilibrium variables are determined according to conditions* (3.2), (3.1), (3.7), *and* (3.9), *plus the (exogenous) technology equation*

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

where $\overline{T}$ is a given initial tax function.

Then, at every type θ, the exogenous technology planner's preferred tax $T^{ex}_{\overline{T}}$ satisfies the following conditions.

$$\frac{T^{ex\prime}_{\overline{T}}(y_\theta)}{1 - T^{ex\prime}_{\overline{T}}(y_\theta)} = \left(1 + \frac{1}{e_\theta}\right) \frac{1 - \overline{B}_{\overline{\beta}_\theta}}{\overline{b}_{\overline{\beta}_\theta} \overline{\beta}_\theta} (1 - \widetilde{g}_\theta) + \gamma^{CES}(1 - g_\theta) \ ,$$

where all variables satisfy the equations listed above under the tax $T^{ex}_{\overline{T}}$; the function $\overline{\beta} : \theta \mapsto \overline{\beta}_\theta$ is given by

$$\overline{\beta}_\theta \coloneqq \kappa_\theta^{1+\gamma^{CES}} h_\theta^{\gamma^{CES}} \left(\phi^*(\overline{T})\right)^{1+\gamma^{CES}} \qquad \forall \, \theta \ ;$$

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

Proof. It can be verified that all steps in the proof of Proposition 4 hold for the case of the exogenous technology planner when imposing

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

With this constraint, we can derive a counterpart to equation (B.25) for the exogenous technology planner:

$$\frac{T'(y_\theta)}{1 - T'(y_\theta)} = \left(1 + \frac{1}{e_\theta}\right)(1 - \widetilde{g}_\theta) \frac{1 - H_\theta}{h_\theta} \widehat{w}_\theta - SE^*_\theta \ .$$

Using Lemma 16 to replace SE^*_θ, we obtain:

$$\begin{aligned} \frac{T'(y_\theta)}{1 - T'(y_\theta)} &= \left(1 + \frac{1}{e_\theta}\right) \frac{1 - H_\theta}{h_\theta} (1 - \widetilde{g}_\theta) \widehat{w}_\theta \\ &\quad + \gamma^{CES} \left[(1 - g_\theta) - \left(1 + \frac{1}{e_\theta}\right) \frac{1 - H_\theta}{h_\theta} (1 - \widetilde{g}_\theta) \widehat{l}_\theta \right] \\ &= \left(1 + \frac{1}{e_\theta}\right) \frac{1 - H_\theta}{h_\theta} (1 - \widetilde{g}_\theta) \left(\widehat{w}_\theta - \gamma^{CES} \widehat{l}_\theta \right) + \gamma^{CES}(1 - g_\theta) \ . \end{aligned}$$

From the perspective of the exogenous technology planner, the wage growth rate is now given by

$$\widehat{w}_\theta = (1+\gamma^{CES})\widehat{\kappa}_\theta + (1+\gamma^{CES})\widehat{\phi}^*_\theta(\overline{T}) + \gamma^{CES}\widehat{h}_\theta + \gamma^{CES}\widehat{l}_\theta \ ,$$

where $\widehat{\phi}^*_\theta(\overline{T})$ denotes the growth rate of technology that prevails in equilibrium under the initial tax system $\overline{T}$. Using this in the previous expression for the exogenous technology planner's optimal tax rates and applying the definition of $\overline{\beta}$ yields:

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

With the change of variable $h_\theta = \bar{b}_{\overline{\beta}_\theta}\overline{\beta}'_\theta$, we obtain equation (3.35),

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

which completes the proof. □

For the optimal tax in the upper Pareto tail of the income distribution, the exogenous technology planner obtains the following characterization.

Corollary 13. *Suppose the conditions of Proposition 4 are satisfied and F and Φ take the CES form introduced in Section 3.3.4. Suppose equilibrium variables were determined by conditions* (3.2), (3.1), (3.7), *and* (3.9), *and by the (exogenous) technology equation*

$$\phi^*(l) = \phi^*(l(\overline{T})) = \underset{\phi\in\Phi}{\operatorname{argmax}}\, F(l(\overline{T}),\phi) \quad \forall l \ ,$$

where $\overline{T}$ is a given initial tax function with $\overline{T}'(y) = \tau^{top}$ for all $y \geq \widetilde{y}$ and some threshold $\widetilde{y}$.

Moreover, assume that under the tax $\overline{T}$ the income distribution satisfies

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

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

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

at the exogenous technology planner's preferred tax.

Then, the exogenous technology planner's preferred tax $T^{ex}_{\overline{T}}$ satisfies

$$\lim_{\theta \to \bar{\theta}} \frac{T^{ex\prime}_{\overline{T}}(y_\theta)}{1 - T^{ex\prime}_{\overline{T}}(y_\theta)} = \frac{1 - g^{top}}{ae} + \frac{a-1}{a} \gamma^{CES} (1 - g^{top}) \ .$$

Proof. The corollary can be derived from Proposition 12 in the same way as Corollary 11 is derived from Proposition 5. In particular, consider first the implications of the Pareto shape of the income distribution under tax $\overline{T}$ for the exogenous inequality measure $\overline{\beta}$. Using changes-of-variable, we obtain

$$\frac{1 - \overline{B}_{\overline{\beta}_\theta}}{\bar{b}_{\overline{\beta}_\theta} \overline{\beta}_\theta} = \frac{1 - H_\theta}{h_\theta} \widehat{\overline{\beta}}_\theta = \frac{1 - M_{y_\theta}}{m_{y_\theta} y_\theta} \frac{\widehat{\overline{\beta}}_\theta}{\widehat{y}_\theta} \ , \tag{B.35}$$

where all incomes are assessed at the given tax $\overline{T}$. The exogenous technology planner's measure of exogenous inequality $\overline{\beta}$ now evolves over the type space according to

$$\widehat{\overline{\beta}}_\theta = \widehat{w}_\theta - \gamma^{CES} \widehat{l}_\theta = \left(1 - \gamma^{CES} \epsilon^w_\theta\right) \widehat{w}_\theta \ ,$$

while

$$\widehat{y}_\theta = \widehat{w}_\theta + \widehat{l}_\theta = (1 + \epsilon^w_\theta) \widehat{w}_\theta \ ,$$

where all endogenous variables are assessed at equilibrium under the given tax $\overline{T}$.[5] Combining the previous expressions, we find that

$$\frac{\widehat{\overline{\beta}}_\theta}{\widehat{y}_\theta} = \frac{1 - \gamma^{CES} \epsilon^w_\theta}{1 + \epsilon^w_\theta}$$

and, with equation (B.35),

$$\frac{1 - \overline{B}_{\overline{\beta}_\theta}}{\bar{b}_{\overline{\beta}_\theta} \overline{\beta}_\theta} = \frac{1 - M_{y_\theta}}{m_{y_\theta} y_\theta} \frac{1 - \gamma^{CES} \epsilon^w_\theta}{1 + \epsilon^w_\theta} \ , \tag{B.36}$$

where incomes and the labor supply elasticity ϵ^w_θ are assessed under the tax $\overline{T}$. Inserting this expression into equation (3.35) for optimal marginal tax rates computed by the exogenous technology

[5]Note that under the initial tax $\overline{T}$ 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 $\overline{T}$ as perceived by the exogenous technology planner and their true equilibrium values.

planner and taking limits yields:[6]

$$\lim_{\theta\to\bar\theta}\frac{T^{ex\prime}_{\overline{T}}(y_\theta)}{1-T^{ex\prime}_{\overline{T}}(y_\theta)} = \left(1+\frac{1}{e}\right)\frac{1-\gamma^{CES}e}{a(1+e)}(1-g^{top})+\gamma^{CES}(1-g^{top})$$
$$=\frac{1-g^{top}}{ae}+\frac{a-1}{a}\gamma^{CES}(1-g^{top})\ .$$

□

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 $\overline{\beta}$ instead of β to measure the degree of exogenous inequality in the economy. Here I show that the exogenous technology planner thereby overestimates exogenous inequality: the function $\overline{\beta}$ progresses at a higher rate in θ than the function β, such that equation (3.36) holds.

First, let $\bar{l}$ be the equilibrium labor input under the initial tax system $\overline{T}$. Then, by construction the growth rates of β and $\overline{\beta}$ must satisfy

$$\widehat{\beta}_\theta = \widehat{w}_\theta(\bar{l},\phi^*(\bar{l}))-(\gamma^{CES}+\rho^{CES})\widehat{\bar{l}}_\theta$$

and

$$\widehat{\overline{\beta}}_\theta = \widehat{w}_\theta(\bar{l},\phi^*(\bar{l}))-\gamma^{CES}\widehat{\bar{l}}_\theta\ .$$

Moreover, by Assumption 3, the marginal rate of tax $\overline{T}$ is strictly below 1 everywhere and hence its rate of progressivity is below 1 as well. Then, equation (3.11) implies that $\epsilon^w_\theta>0$ and hence $\widehat{\bar{l}}_\theta=\epsilon^w_\theta\widehat{w}_\theta(\bar{l},\phi^*(\bar{l}))>0$ for all θ. Combining this with the expressions for $\widehat{\beta}_\theta$ and $\widehat{\overline{\beta}}_\theta$, we find that $\widehat{\beta}_\theta<\widehat{\overline{\beta}}_\theta$. Now use a change-of-variable to obtain

$$b_{\beta_\theta}\beta_\theta=\frac{f_\theta}{\widehat{\beta}_\theta}>\frac{f_\theta}{\widehat{\overline{\beta}}_\theta}=\bar{b}_{\overline{\beta}_\theta}\overline{\beta}_\theta\ ,$$

and hence

$$\frac{1-B_{\beta_\theta}}{b_{\beta_\theta}\beta_\theta}<\frac{1-\overline{B}_{\overline{\beta}_\theta}}{\bar{b}_{\overline{\beta}_\theta}\overline{\beta}_\theta}\ ,$$

which proves equation (3.36) in the main text.

[6] The limit computations are analogous to those in the proof of Corollary 11, so I omit the details here.

Intuitively, since labor supply increases in skill θ under tax $\overline{T}$, technology under this tax must be skill-biased. The exogenous technology planner falsely believes this skill bias to be exogenous and to persist irrespective of changes in labor supply. Thereby, he overestimates the degree of exogenous inequality in the economy.

B.2 Calibration Details

This Appendix provides more detailed information on two steps of the calibration procedure described in Section 3.7.1 of the main text: the calibration of the directed technical change elasticity ρ^{CES} and of the exogenous technology parameter $\kappa : \theta \mapsto \kappa_\theta$.

B.2.1 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 3.1. The long-run estimates in Table 3.1 (10 years or more) are equated with the sum of within-technology substitution and directed technical change elasticities $\gamma^{CES} + \rho^{CES}$. The short-run estimate (2 years, from Carneiro et al. (2019)) is equated with γ^{CES}, in line with other short-run estimates as discussed in the main text.

Here, I give a brief overview over each of the studies listed in Table 3.1 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 3.1 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 3.1 in the present paper.

Second, Figures 4 and 5 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 3.1 in the present paper.

Finally, the third row of Table 3.1 is obtained from the plots presented in Figure 4 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

$$\frac{\log\left(\frac{w_c^{17}}{w_{nc}^{17}}\right) - \log\left(\frac{w_c^{0}}{w_{nc}^{0}}\right)}{\log\left(\frac{l_c^{17}}{l_{nc}^{17}}\right) - \log\left(\frac{l_c^{0}}{l_{nc}^{0}}\right)} \approx 0.05 \ ,$$

which is the estimate used in Table 3.1.

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 relative supply increase 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 3.1.

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 3.1.

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 (B.6). 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, 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 3.1. 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 3.1.

Relative to the other studies listed in Table 3.1, 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.2.2 Calibration of the Exogenous Technology Parameter

To calibrate the exogenous technology parameter κ, an estimate of the earnings distribution under the initial tax system $\overline{T}$ 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 $\overline{T}$ 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%.[7]

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

$$\frac{ym_y}{1-M_y} = \frac{ym_y^{lognormal}}{1-M_y^{lognormal}}\left(1-\Phi\left(\frac{y-200000}{\sigma^{normal}}\right)\right) + \frac{ym_y^{Pareto}}{1-M_y^{Pareto}}\Phi\left(\frac{y-200000}{\sigma}\right) \ ,$$

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 depicted in Figure B.1 (see again Figure 2 in Diamond and Saez 2011 for comparison with the empirical US hazard ratio in 2005).

Given the hazard ratio of incomes, 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 $\overline{T}$.

[7] The main results are robust to other values of the mass point between 0% and 10%.

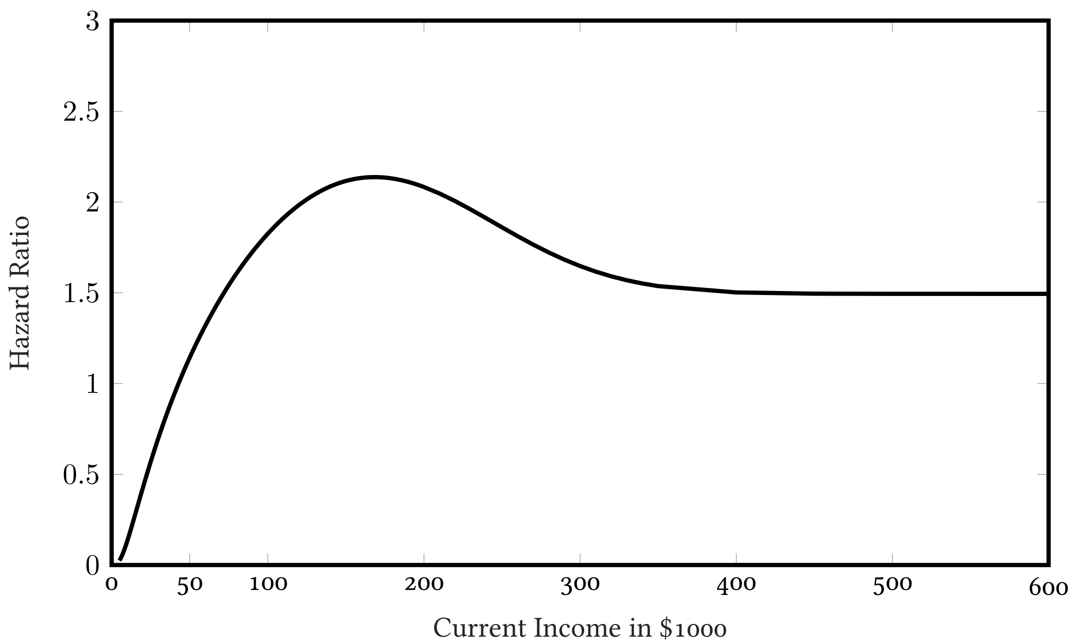

Figure B.1. The figure shows the hazard ratio of the income distribution under the initial tax $\overline{T}$ used to calibrate the exogenous technology parameter κ. The construction of the hazard ratio follows the description in the text. The hazard ratio approximates the empirical hazard ratio of the US earnings distribution in 2005, as depicted, for example, in Figure 2 in Diamond and Saez (2011).

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 (3.2) by l_θ and solving for it yields

$$l_\theta = \left(R'_{\overline{T}}(y_\theta)y_\theta\right)^{\frac{e}{1+e}} ,$$

where I used that the disutility of labor is isoelastic in the quantitative analysis. With the estimate of $\overline{T}$ described in the main text, the previous equation allows to compute labor inputs under $\overline{T}$.

For the second step, start from equations (3.14) and (B.7), copied here for convenience:

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

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

The total amount of R&D resources $\overline{C}$ is not identified separately from κ, so I normalize it to satisfy

$$\overline{C} = \int_{\underline{\theta}}^{\overline{\theta}} (\kappa_\theta h_\theta l_\theta)^{\frac{(\sigma-1)\delta}{(\delta-1)\sigma+1}} \, d\theta$$

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 $F = \int_{\underline{\theta}}^{\overline{\theta}} y_\theta m_y \, dy$ by Euler's theorem, this allows to compute κ.

Finally, the optimal tax formulae in Proposition 5 and Proposition 12 require the inverse hazard ratio of the exogenous parameters β and $\overline{\beta}$, respectively. In principle, β and $\overline{\beta}$ 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 rounds of approximations, I choose a more direct way and compute the inverse hazard ratios of β and $\overline{\beta}$ directly from the hazard ratio of the income distribution, using equations (B.34) and (B.36). 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.

B.3 Supplementary Material

This appendix contains several results complementary to those presented in the main text. Section B.3.1 presents an alternative representation of the labor inputs effects of tax reforms, which clarifies the relationship between the corresponding results in the main text and those presented in Sachs et al. (2020). Section B.3.2 provides results for the welfare effects of tax reforms, extending the analysis of the tax reform effects on wage inequality in the main text. Section B.3.3 derives an alternative optimal tax formula for the CES case, which gives rise to an alternative intuition behind the impact of directed technical change on optimal taxes. Sections B.3.4 and B.3.5 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 B.3.6 contains all proofs of the results presented in this appendix.

B.3.1 Alternative Representation of Labor Input Responses to Tax Reforms

In this section, I compare equation (3.26) for the effects of a tax reform on labor inputs (Lemma 8) with an alternative expression for these effects obtained by following the iteration approach of Sachs et al. (2020).

For that, I first define elasticities of aggregate labor supply of a given type with respect to the marginal retention rate and the wage. In particular, note that, if all workers of type θ change their labor supply jointly, the wage w_θ will react. The wage response then induces a change in labor supply in addition to the direct response described by the individual labor supply elasticities ϵ_θ^R and ϵ_θ^w. Starting from the individual labor supply elasticities, we can construct elasticities that account for this feedback effect as follows:

$$\bar{\epsilon}_\theta^R(T,l,w) := \frac{\epsilon_\theta^R(T,l,w)}{1-(\gamma_{\theta,\theta}+\rho_{\theta,\theta})\epsilon_\theta^w(T,l,w)} \tag{B.37}$$

and

$$\bar{\epsilon}_\theta^w(T,l,w) := \frac{\epsilon_\theta^w(T,l,w)}{1-(\gamma_{\theta,\theta}+\rho_{\theta,\theta})\epsilon_\theta^w(T,l,w)}\ . \tag{B.38}$$

Relative to the individual elasticities, the aggregate elasticities are scaled by the feedback from the wage to labor supply. If the own-wage effect $\gamma_{\theta,\theta}+\rho_{\theta,\theta}$ is negative (positive), the individual elasticity is scaled down (up), as an increase in labor supply depresses (raises) the wage, which then counteracts (amplifies) the initial labor supply change.

We can now use the aggregate labor supply elasticities to rearrange equation (3.22) as follows:

$$\widehat{l}_{\theta,\tau} = -\bar{\epsilon}_\theta^R \frac{\tau'(y_\theta(T))}{1-T'(y_\theta(T))} + \bar{\epsilon}_\theta^w \int_{\underline{\theta}}^{\overline{\theta}} (\gamma_{\theta,\widetilde{\theta}}+\rho_{\theta,\widetilde{\theta}})\widehat{l}_{\widetilde{\theta},\tau}(T)\, d\widetilde{\theta}\ . \tag{B.39}$$

This fixed point equation can be solved by iteratively inserting the right-hand side of the equation into itself (see the proof below for details).

Lemma 17. *Fix an initial tax T and let $(\gamma^{CES}+\rho^{CES})\epsilon_\theta^w < 1$ under T such that the aggregate elasticities $\bar{\epsilon}_\theta^R$ and $\bar{\epsilon}_\theta^R$ are well defined. Moreover, suppose that under T,*[8]

$$\int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} (\bar{\epsilon}_\theta^w(\gamma_{\theta,\widetilde{\theta}}+\rho_{\theta,\widetilde{\theta}}))^2\, d\widetilde{\theta} d\theta < 1 \quad \textit{and} \quad \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} (\bar{\epsilon}_\theta^w \rho_{\theta,\widetilde{\theta}})^2\, d\widetilde{\theta} d\theta < 1\ .$$

[8] These conditions serve the same purpose as conditions (3.23) and (3.24) in Lemma (8).

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

$$\widehat{l}_{\theta,\tau} = -\overline{\epsilon}_{\theta}^{R} \frac{\tau'(y_\theta(T))}{1 - T'(y_\theta(T))} + \overline{\epsilon}_{\theta}^{w} \int_{\underline{\theta}}^{\overline{\theta}} \overline{\rho}_{\theta,\widetilde{\theta}} (-\overline{\epsilon}_{\widetilde{\theta}}^{R}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1 - T'(y_{\widetilde{\theta}}(T))} \, d\widetilde{\theta}$$
$$+ \overline{\epsilon}_{\theta}^{w} \int_{\underline{\theta}}^{\overline{\theta}} \overline{\gamma}_{\theta,\widetilde{\theta}} (-\overline{\epsilon}_{\widetilde{\theta}}^{R}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1 - T'(y_{\widetilde{\theta}}(T))} \, d\widetilde{\theta} \,, \quad \text{(B.40)}$$

where

$$\overline{\rho}_{\theta,\widetilde{\theta}} = \sum_{n=1}^{\infty} \rho_{\theta,\widetilde{\theta}}^{(n)}$$
$$\rho_{\theta,\widetilde{\theta}}^{(1)} = \rho_{\theta,\widetilde{\theta}}$$
$$\rho_{\theta,\widetilde{\theta}}^{(n)} = \int_{\underline{\theta}}^{\overline{\theta}} \rho_{\theta,\widehat{\theta}}^{(n-1)} \overline{\epsilon}_{\widehat{\theta}}^{w} \rho_{\widehat{\theta},\widetilde{\theta}} \, d\widehat{\theta} \quad \forall \, n > 1$$

and

$$\overline{\gamma}_{\theta,\widetilde{\theta}} = \sum_{n=1}^{\infty} \gamma_{\theta,\widetilde{\theta}}^{(n)}$$
$$\gamma_{\theta,\widetilde{\theta}}^{(1)} = \gamma_{\theta,\widetilde{\theta}}$$
$$\gamma_{\theta,\widetilde{\theta}}^{(n)} = \int_{\underline{\theta}}^{\overline{\theta}} \gamma_{\theta,\widehat{\theta}}^{(n-1)} \overline{\epsilon}_{\widehat{\theta}}^{w} (\gamma_{\widehat{\theta},\widetilde{\theta}} + \rho_{\widehat{\theta},\widetilde{\theta}}) \, d\widehat{\theta} + \int_{\underline{\theta}}^{\overline{\theta}} \rho_{\theta,\widehat{\theta}}^{(n-1)} \overline{\epsilon}_{\widehat{\theta}}^{w} \gamma_{\widehat{\theta},\widetilde{\theta}} \, d\widehat{\theta} \quad \forall \, n > 1.$$

Proof. See Appendix B.3.6. □

The difference between the expressions provided by Lemma 8 and those given by Lemma 17 is that Lemma 17 uses the aggregate labor supply elasticities $\overline{\epsilon}_{\theta}^{R}$ and $\overline{\epsilon}_{\theta}^{w}$ whereas Lemma 8 uses individual labor supply elasticities. As can be seen from equations (B.37) and (B.38), the aggregate elasticities already contain the own-wage elasticities $\gamma_{\theta,\theta}$ and $\rho_{\theta,\theta}$. This makes it difficult to disentangle directed technical change from within-technology substitution effects on the basis of Lemma 17. For example, when decomposing the total labor supply effect given by Lemma 17 along the lines of the decomposition provided in equation (3.26) of Lemma 8, we would end up with a directed technical change component that still contains within-technology substitution effects via the aggregate labor supply elasticities. Signing the impact of the directed technical change component on relative wages would then require much more demanding restrictions. Specifically, we would have to restrict heterogeneity in own-wage elasticities, which would require restrictions on the aggregate production F, beyond the restrictions already necessitated by the application of directed technical change theory.

B.3.2 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 expected 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:

$$\widetilde{W}(T) := V\left(\{u_\theta(c_\theta(T), l_\theta(T))\}_{\theta\in\Theta}\right) \ .$$

Given the welfare function $\widetilde{W}(T)$, the welfare effect of a tax reform can now be decomposed as follows.

Proposition 13. *Fix an initial tax T. The welfare effect of a tax reform τ can be written as*

$$\begin{aligned} D_\tau \widetilde{W}(T) = & \underbrace{\int_{\underline{\theta}}^{\overline{\theta}} (1-g_\theta)\tau(y_\theta(T))h_\theta \, d\theta}_{:=ME_\tau(T)} + \underbrace{\int_{\underline{\theta}}^{\overline{\theta}} T'(y_\theta(T))y_\theta(T)(-\epsilon_\theta^R)\frac{\tau'(y_\theta(T))}{1-T'(y_\theta(T))}h_\theta \, d\theta}_{:=BE_\tau(T)} \\ & + \underbrace{\int_{\underline{\theta}}^{\overline{\theta}} \left[g_\theta(1-T'(y_\theta(T))) + T'(y_\theta(T))(1+\epsilon_\theta^w)\right] y_\theta(T)\frac{1}{w_\theta(T)}D_{\phi,\tau}w_\theta(T,\phi^*(T))h_\theta \, d\theta}_{:=TE_\tau^W(T)} \\ & + \underbrace{\int_{\underline{\theta}}^{\overline{\theta}} \left[g_\theta(1-T'(y_\theta(T))) + T'(y_\theta(T))(1+\epsilon_\theta^w)\right] y_\theta(T)\frac{1}{w_\theta(T)}D_\tau w_\theta(T,\phi^*(T))h_\theta \, d\theta}_{:=SE_\tau^W(T)} \ . \end{aligned}$$

Proof. See Appendix B.3.6. □

Proposition 13 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,

$$\int_{\underline{\theta}}^{\overline{\theta}} g_\theta(1-T'(y_\theta(T)))y_\theta(T)\frac{1}{w_\theta(T)}D_{\phi,\tau}w_\theta(T,\phi^*(T))h_\theta \, d\theta \ ,$$

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,

$$\int_{\underline{\theta}}^{\overline{\theta}} T'(y_\theta(T))(1+\epsilon_\theta^w) y_\theta(T) \frac{1}{w_\theta(T)} D_{\phi,\tau} w_\theta(T, \phi^*(T)) h_\theta \, d\theta \ ,$$

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.[9]

Importantly, even if the induced technical change reduces the skill premium (e.g., because τ is progressive and the conditions of Corollary 9 are satisfied), we cannot sign the directed technical change effects on welfare unambiguously. This is because, when for example starting 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.[10]

The final term in Proposition 13, $SE_\tau^W(T)$ 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 $SE_\tau^W(T)$ can generally not be signed.

Importantly, however, Proposition 13 can be combined with equations (3.28) and (B.16) for the relative wage effects of tax reforms from Propositions 3 and 11 (in Appendix B.1.3). This yields a formula for the welfare effects of tax reforms in terms of empirically observable quantities and welfare weights. I use this combination of expressions to quantify the welfare effects of tax reforms and the contribution of directed technical change in Appendix B.3.4.

The implications of Proposition 13 may be somewhat unexpected in light of the previous section's result. 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

[9] The 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.

[10] This is similar to 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.

adjust the question posed by Proposition 13.

Concretely, 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 which welfare can be improved by some progressive reform. In particular, let

$$\mathcal{T} := \Big\{ T \mid T \text{ is CRP, } \exists \tau \text{ progressive s.t. } D_\tau \widetilde{W}(T) > 0 \Big\}$$

denote the set of CRP tax schedules that can be improved in a welfare sense by a progressive tax reform. The restriction to CRP taxes is imposed to invoke Corollary 9. Specifically, combining the CRP restriction with isoelastic disutility of labor and the CES production structure from Section 3.3.4 ensures, according to Corollary 9, that any progressive tax reform induces equalizing technical change.

Note at this point that CRP tax schedules provide good approximations to the actual income tax in the US at several points in time over the last five decades, including the current US tax system (Heathcote et al., 2017). Hence, the restriction to initial taxes that take the CRP form seems without much loss for application to empirical tax schedules.

As a benchmark for comparison that does not include directed technical change effects, let

$$D_\tau^{ex} \widetilde{W}(T) := D_\tau \widetilde{W}(T) \Big|_{\rho_{\theta,\widetilde{\theta}}=0 \; \forall \, \theta,\widetilde{\theta}} \tag{B.41}$$

denote the welfare effect of a reform τ when counterfactually setting all technical change elasticities to zero (or, put differently, when holding technology fixed). Then, we can define

$$\mathcal{T}^{ex} := \Big\{ T \mid T \text{ is CRP, } \exists \tau \text{ progressive s.t. } D_\tau^{ex} \widetilde{W}(T) > 0 \Big\}$$

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 14. *Suppose F and Φ are CES as introduced in Section 3.3.4 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 B.3.6. □

This result 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. It aligns neatly with Corollary 9, whereby progressive reforms induce equalizing technical change.

The idea behind Proposition 14 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 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.

B.3.3 Alternative Optimal Tax Formula

In the following, I provide an alternative expression for optimal marginal tax rates in the CES case, which allows for an interpretation of directed technical change effects via their effect on the aggregate labor supply elasticity $\bar{\epsilon}_\theta^R$ (see Appendix B.3.1).

Proposition 15. *Suppose the conditions of Proposition 4 are satisfied and F and Φ take the CES form introduced in Section 3.3.4. Additionally, let $(\gamma^{CES} + \rho^{CES})\epsilon_\theta^w < 1$ at the optimal tax T. Then, the conditions for optimal marginal tax rates can be written as*

$$\frac{T'(y)}{1 - T'(y)} = \frac{1}{\bar{\epsilon}_{\theta_y}^R} \frac{1 - M_y}{m_y y}(1 - \widetilde{g}_{\theta_y}) + \gamma^{CES}(1 - g_{\theta_y}) + \rho^{CES}(1 - g_{\theta_y}) \;, \tag{B.42}$$

where

$$\bar{\epsilon}_{\theta_y}^R(T, l, w) := \frac{\epsilon_{\theta_y}^R(T, l, w)}{1 - (\gamma_{\theta_y,\theta_y} + \rho_{\theta_y,\theta_y})\epsilon_{\theta_y}^w(T, l, w)}$$

denotes the elasticity of aggregate labor supply of type θ with respect to the marginal retention rate (see Appendix B.3.1); all variables are evaluated at equilibrium under the optimal tax T; M and m denote the cumulative distribution and the density function of y at the optimum; and θ_y denotes the type of workers who earn income y at the optimum.

Proof. See Appendix B.3.6. □

Equation (B.42) offers an alternative perspective on the role of directed technical change in the CES case. It extends the expression in Proposition 3 of Sachs et al. (2020) to account for directed technical change.

The adjustment term $\rho^{CES}(1 - g_{\theta_y})$ is the same as in equation (3.33) and does not provide any new insights. As described in the main text, it calls for a progressive adjustment of marginal tax rates.

The second place in equation (B.42) where directed technical change elasticities appear is the aggregate labor supply elasticity $\bar{\epsilon}_\theta^R$. This elasticity measures the response of labor supply of type θ to an increase in the marginal retention rate when taking into account the change in type θ's wage induced by the labor supply response. The change in type θ's wage is modified by directed technical change effects. Specifically, when labor supply of type θ falls due to an increase in the marginal tax rate, directed technical change reduces the wage of type θ and thereby amplifies the fall in labor supply. In this way, directed technical change magnifies the labor supply response to marginal tax changes, which leads to a downwards adjustment of optimal marginal tax rates.

B.3.4 Quantitative Analysis: Welfare Effects of Tax Reforms

Here, I complement the quantitative analysis of the wage effects of progressive tax reforms in Section 3.7 by assessing the welfare effects of these reforms. In particular, I consider the same progressive tax reform as described in Section3.7 and compute its welfare effects based on Proposition 13 from Appendix B.3.6. The calibration is the same as described in the main text.

Figure B.2 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 13 and the subsequent discussion), the influence of directed 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

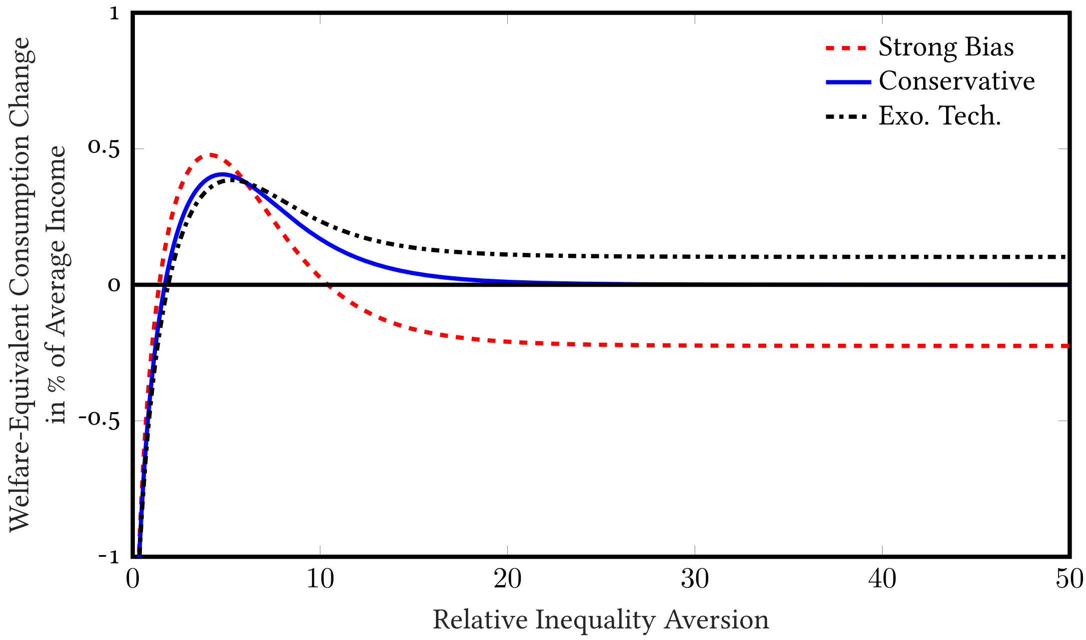

Figure B.2. The figure displays changes in the lump-sum payment that are equivalent to the welfare effect of the progressive tax reform described in the text. These changes are measured in % of the pre-reform average income and displayed for different values of the relative inequality aversion parameter r. The baseline calibration described in the main text applies.

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; accounting for directed technical change makes the reform undesirable.

Notwithstanding these results, Proposition 14 implies that even for such very high values of inequality aversion, there must be a different progressive reform that raises welfare when accounting for directed technical change.

B.3.5 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 B.3 also displays the corresponding marginal tax rates preferred by the exogenous technology planner.

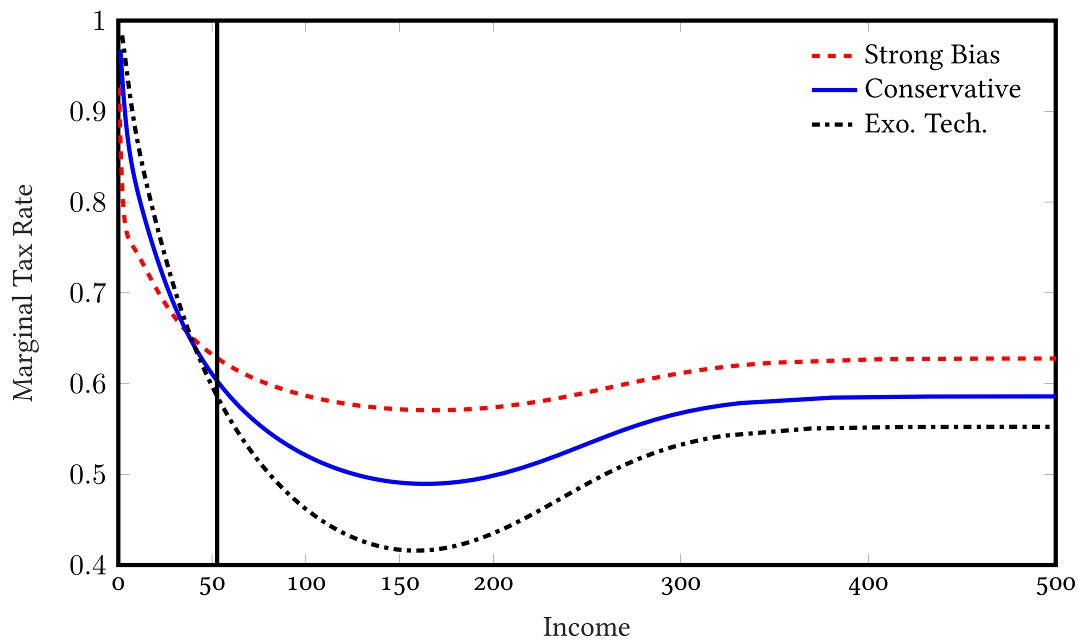

Figure B.3. The figure displays optimal marginal tax rates by income level for a relative inequality aversion parameter of $r = 50$. Otherwise, the baseline calibration described in the main text applies.

The qualitative insights from Section 3.7 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 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 disappears. 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.

B.3.6 Proofs for the Supplementary Material

This section contains all proofs for the results presented in Appendix B.3.

Alternative Representation of Labor Input Responses to Tax Reforms

Proof. Following Sachs et al. (2020), I solve the fixed point equation (B.39) by iteration. Within the iteration steps, I separate the directed technical change from the (within-technology) substitution effects to obtain a decomposition of the total labor input response along the lines of Lemma 8.

Step 1. The first part of the proof proceeds by induction. We start by substituting equation (B.39) into itself:

$$\begin{aligned}
\widehat{l}_{\theta,\tau} =& -\bar{\epsilon}^R_\theta \frac{\tau'(y_\theta(T))}{1-T'(y_\theta(T))} + \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} (\gamma_{\theta,\widetilde{\theta}} + \rho_{\theta,\widetilde{\theta}}) \left[-\bar{\epsilon}^R_{\widetilde{\theta}} \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))} + \bar{\epsilon}^w_{\widetilde{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} (\gamma_{\widetilde{\theta},\widehat{\theta}} + \rho_{\widetilde{\theta},\widehat{\theta}}) \widehat{l}_{\widehat{\theta},\tau}(T)\, d\widehat{\theta} \right] d\widetilde{\theta} \\
=& -\bar{\epsilon}^R_\theta \frac{\tau'(y_\theta(T))}{1-T'(y_\theta(T))} + \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \rho_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))}\, d\widetilde{\theta} + \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \rho_{\theta,\widetilde{\theta}} \bar{\epsilon}^w_{\widetilde{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} \rho_{\widetilde{\theta},\widehat{\theta}} \widehat{l}_{\widehat{\theta},\tau}(T)\, d\widehat{\theta}\, d\widetilde{\theta} \\
&+ \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \gamma_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))}\, d\widetilde{\theta} + \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \gamma_{\theta,\widetilde{\theta}} \bar{\epsilon}^w_{\widetilde{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} (\gamma_{\widetilde{\theta},\widehat{\theta}} + \rho_{\widetilde{\theta},\widehat{\theta}}) \widehat{l}_{\widehat{\theta},\tau}(T)\, d\widehat{\theta}\, d\widetilde{\theta} \\
&+ \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \rho_{\theta,\widetilde{\theta}} \bar{\epsilon}^w_{\widetilde{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} \gamma_{\widetilde{\theta},\widehat{\theta}} \widehat{l}_{\widehat{\theta},\tau}(T)\, d\widehat{\theta}\, d\widetilde{\theta}\,.
\end{aligned}$$

Changing the order of integration in the terms containing $\widehat{l}_{\widehat{\theta},\tau}$ and summarizing the last two terms, we obtain

$$\begin{aligned}
\widehat{l}_{\theta,\tau} =& -\bar{\epsilon}^R_\theta \frac{\tau'(y_\theta(T))}{1-T'(y_\theta(T))} + \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \rho_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))}\, d\widetilde{\theta} + \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} \rho_{\theta,\widetilde{\theta}} \bar{\epsilon}^w_{\widetilde{\theta}} \rho_{\widetilde{\theta},\widehat{\theta}}\, d\widetilde{\theta}\, \widehat{l}_{\widehat{\theta},\tau}(T)\, d\widehat{\theta} \\
&+ \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \gamma_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))}\, d\widetilde{\theta} \\
&+ \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \left[\int_{\underline{\theta}}^{\overline{\theta}} \gamma_{\theta,\widetilde{\theta}} \bar{\epsilon}^w_{\widetilde{\theta}} (\gamma_{\widetilde{\theta},\widehat{\theta}} + \rho_{\widetilde{\theta},\widehat{\theta}})\, d\widetilde{\theta} + \int_{\underline{\theta}}^{\overline{\theta}} \rho_{\theta,\widetilde{\theta}} \bar{\epsilon}^w_{\widetilde{\theta}} \gamma_{\widetilde{\theta},\widehat{\theta}}\, d\widetilde{\theta} \right] \widehat{l}_{\widehat{\theta},\tau}(T)\, d\widehat{\theta}\,.
\end{aligned}$$

Using the definitions of $\rho^{(n)}_{\theta,\widetilde{\theta}}$ and $\gamma^{(n)}_{\theta,\widetilde{\theta}}$ given in the lemma, this expression can be rewritten as

$$\widehat{l}_{\theta,\tau} = -\,\bar{\epsilon}^R_\theta \frac{\tau'(y_\theta(T))}{1-T'(y_\theta(T))} + \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \rho^{(1)}_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))}\, d\widetilde{\theta} + \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \rho^{(2)}_{\theta,\widehat{\theta}} \widehat{l}_{\widehat{\theta},\tau}(T)\, d\widehat{\theta}$$
$$+\,\bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \gamma^{(1)}_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))}\, d\widetilde{\theta} + \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \gamma^{(2)}_{\theta,\widehat{\theta}} \widehat{l}_{\widehat{\theta},\tau}(T)\, d\widehat{\theta}\,.$$

This constitutes the base case for induction. As an induction hypothesis, suppose now that for any $N \geq 1$ the following holds:

$$\widehat{l}_{\theta,\tau} = -\,\bar{\epsilon}^R_\theta \frac{\tau'(y_\theta(T))}{1-T'(y_\theta(T))} + \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \sum_{n=1}^{N} \rho^{(n)}_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))}\, d\widetilde{\theta} + \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \rho^{(N+1)}_{\theta,\widehat{\theta}} \widehat{l}_{\widehat{\theta},\tau}(T)\, d\widehat{\theta}$$
$$+\,\bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \sum_{n=1}^{N} \gamma^{(n)}_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))}\, d\widetilde{\theta} + \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \gamma^{(N+1)}_{\theta,\widehat{\theta}} \widehat{l}_{\widehat{\theta},\tau}(T)\, d\widehat{\theta}\,. \tag{B.43}$$

Then, using equation (B.39) to substitute for $\widehat{l}_{\widehat{\theta},\tau}$ on the right-hand-side yields

$$\widehat{l}_{\theta,\tau} = -\,\bar{\epsilon}^R_\theta \frac{\tau'(y_\theta(T))}{1-T'(y_\theta(T))} + \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \sum_{n=1}^{N} \rho^{(n)}_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))}\, d\widetilde{\theta}$$
$$+\,\bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \rho^{(N+1)}_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))}\, d\widetilde{\theta} + \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \rho^{(N+1)}_{\theta,\widetilde{\theta}} \bar{\epsilon}^w_{\widetilde{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} \rho_{\widetilde{\theta},\widehat{\theta}} \widehat{l}_{\widehat{\theta},\tau}(T)\, d\widehat{\theta}\, d\widetilde{\theta}$$
$$+\,\bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \sum_{n=1}^{N} \gamma^{(n)}_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))}\, d\widetilde{\theta} + \bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \gamma^{(N+1)}_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))}\, d\widetilde{\theta}$$
$$+\,\bar{\epsilon}^w_\theta \int_{\underline{\theta}}^{\overline{\theta}} \left[\gamma^{(N+1)}_{\theta,\widetilde{\theta}} \bar{\epsilon}^w_{\widetilde{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} (\gamma_{\widetilde{\theta},\widehat{\theta}} + \rho_{\widetilde{\theta},\widehat{\theta}}) \widehat{l}_{\widehat{\theta},\tau}(T)\, d\widehat{\theta} + \rho^{(N+1)}_{\theta,\widetilde{\theta}} \bar{\epsilon}^w_{\widetilde{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} \gamma_{\widetilde{\theta},\widehat{\theta}} \widehat{l}_{\widehat{\theta},\tau}(T)\, d\widehat{\theta} \right] d\widetilde{\theta}\,.$$

Changing again the order of integration in the terms containing $\widehat{l}_{\widehat{\theta},\tau}$ yields

$$\begin{aligned}
\widehat{l}_{\theta,\tau} = &- \bar{\epsilon}^R_{\theta} \frac{\tau'(y_\theta(T))}{1-T'(y_\theta(T))} + \bar{\epsilon}^w_{\theta} \int_{\underline{\theta}}^{\overline{\theta}} \sum_{n=1}^{N} \rho^{(n)}_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))} \, d\widetilde{\theta} \\
&+ \bar{\epsilon}^w_{\theta} \int_{\underline{\theta}}^{\overline{\theta}} \rho^{(N+1)}_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))} \, d\widetilde{\theta} + \bar{\epsilon}^w_{\theta} \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} \rho^{(N+1)}_{\theta,\widetilde{\theta}} \bar{\epsilon}^w_{\widetilde{\theta}} \rho_{\widetilde{\theta},\widehat{\theta}} \, d\widetilde{\theta} \widehat{l}_{\widehat{\theta},\tau}(T) \, d\widehat{\theta} \\
&+ \bar{\epsilon}^w_{\theta} \int_{\underline{\theta}}^{\overline{\theta}} \sum_{n=1}^{N} \gamma^{(n)}_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))} \, d\widetilde{\theta} + \bar{\epsilon}^w_{\theta} \int_{\underline{\theta}}^{\overline{\theta}} \gamma^{(N+1)}_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))} \, d\widetilde{\theta} \\
&+ \bar{\epsilon}^w_{\theta} \int_{\underline{\theta}}^{\overline{\theta}} \left[\int_{\underline{\theta}}^{\overline{\theta}} \gamma^{(N+1)}_{\theta,\widetilde{\theta}} \bar{\epsilon}^w_{\widetilde{\theta}} (\gamma_{\widetilde{\theta},\widehat{\theta}} + \rho_{\widetilde{\theta},\widehat{\theta}}) \, d\widetilde{\theta} + \int_{\underline{\theta}}^{\overline{\theta}} \rho^{(N+1)}_{\theta,\widetilde{\theta}} \bar{\epsilon}^w_{\widetilde{\theta}} \gamma_{\widetilde{\theta},\widehat{\theta}} \, d\widetilde{\theta} \right] \widehat{l}_{\widehat{\theta},\tau}(T) \, d\widehat{\theta} \\
= &- \bar{\epsilon}^R_{\theta} \frac{\tau'(y_\theta(T))}{1-T'(y_\theta(T))} + \bar{\epsilon}^w_{\theta} \int_{\underline{\theta}}^{\overline{\theta}} \sum_{n=1}^{N+1} \rho^{(n)}_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))} \, d\widetilde{\theta} + \bar{\epsilon}^w_{\theta} \int_{\underline{\theta}}^{\overline{\theta}} \rho^{(N+2)}_{\theta,\widehat{\theta}} \widehat{l}_{\widehat{\theta},\tau}(T) \, d\widehat{\theta} \\
&+ \bar{\epsilon}^w_{\theta} \int_{\underline{\theta}}^{\overline{\theta}} \sum_{n=1}^{N+1} \gamma^{(n)}_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))} \, d\widetilde{\theta} + \bar{\epsilon}^w_{\theta} \int_{\underline{\theta}}^{\overline{\theta}} \gamma^{(N+2)}_{\theta,\widehat{\theta}} \widehat{l}_{\widehat{\theta},\tau}(T) \, d\widehat{\theta} \, .
\end{aligned}$$

Hence the induction hypothesis (B.43) holds for any $N \geq 1$.

Step 2. I take the induction hypothesis from step 1 and let N go to infinity:

$$\begin{aligned}
\widehat{l}_{\theta,\tau} = &- \bar{\epsilon}^R_{\theta} \frac{\tau'(y_\theta(T))}{1-T'(y_\theta(T))} + \bar{\epsilon}^w_{\theta} \int_{\underline{\theta}}^{\overline{\theta}} \sum_{n=1}^{\infty} \rho^{(n)}_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))} \, d\widetilde{\theta} + \bar{\epsilon}^w_{\theta} \lim_{N\to\infty} \int_{\underline{\theta}}^{\overline{\theta}} \rho^{(N)}_{\theta,\widehat{\theta}} \widehat{l}_{\widehat{\theta},\tau}(T) \, d\widehat{\theta} \\
&+ \bar{\epsilon}^w_{\theta} \int_{\underline{\theta}}^{\overline{\theta}} \sum_{n=1}^{\infty} \gamma^{(n)}_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))} \, d\widetilde{\theta} + \bar{\epsilon}^w_{\theta} \lim_{N\to\infty} \int_{\underline{\theta}}^{\overline{\theta}} \gamma^{(N)}_{\theta,\widehat{\theta}} \widehat{l}_{\widehat{\theta},\tau}(T) \, d\widehat{\theta} \, .
\end{aligned}$$

The goal is to prove that the infinite series are convergent while the limit expressions containing $\widehat{l}_{\widehat{\theta},\tau}$ vanish on the right-hand side. Let

$$\begin{aligned}
A_n &:= \int_{\underline{\theta}}^{\overline{\theta}} \rho^{(n)}_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))} \, d\widetilde{\theta} \\
B_n &:= \int_{\underline{\theta}}^{\overline{\theta}} \gamma^{(n)}_{\theta,\widetilde{\theta}} (-\bar{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))} \, d\widetilde{\theta} \, .
\end{aligned}$$

I start with the series $\sum_{n=1}^{\infty} A_n$. First, using the definition of $\rho^{(n)}_{\theta,\widetilde{\theta}}$, the Cauchy-Schwarz inequality implies

$$\left(\rho^{(n)}_{\theta,\widetilde{\theta}}\right)^2 = \left(\int_{\underline{\theta}}^{\overline{\theta}} \rho^{(n-1)}_{\theta,\widehat{\theta}} \overline{\epsilon}^w_{\widehat{\theta}} \rho_{\widehat{\theta},\widetilde{\theta}} \, d\widehat{\theta}\right)^2 \leq \int_{\underline{\theta}}^{\overline{\theta}} \left(\rho^{(n-1)}_{\theta,\widehat{\theta}}\right)^2 d\widehat{\theta} \int_{\underline{\theta}}^{\overline{\theta}} \left(\overline{\epsilon}^w_{\widehat{\theta}} \rho_{\widehat{\theta},\widetilde{\theta}}\right)^2 d\widehat{\theta} \,.$$

Integrating over $\widetilde{\theta}$ yields

$$\int_{\underline{\theta}}^{\overline{\theta}} \left(\rho^{(n)}_{\theta,\widetilde{\theta}}\right)^2 d\widetilde{\theta} \leq \int_{\underline{\theta}}^{\overline{\theta}} \left(\rho^{(n-1)}_{\theta,\widehat{\theta}}\right)^2 d\widehat{\theta} \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} \left(\overline{\epsilon}^w_{\widehat{\theta}} \rho_{\widehat{\theta},\widetilde{\theta}}\right)^2 d\widehat{\theta} \, d\widetilde{\theta} \,.$$

Then, applying the inequality iteratively $n-2$ times, we obtain

$$\int_{\underline{\theta}}^{\overline{\theta}} \left(\rho^{(n)}_{\theta,\widetilde{\theta}}\right)^2 d\widetilde{\theta} \leq \int_{\underline{\theta}}^{\overline{\theta}} \left(\rho_{\theta,\widehat{\theta}}\right)^2 d\widehat{\theta} \left[\int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} \left(\overline{\epsilon}^w_{\widehat{\theta}} \rho_{\widehat{\theta},\widetilde{\theta}}\right)^2 d\widehat{\theta} \, d\widetilde{\theta}\right]^{n-1} \,.$$

Moreover, again applying the Cauchy-Schwarz inequality, we have

$$A_n^2 \leq \int_{\underline{\theta}}^{\overline{\theta}} \left(\rho^{(n)}_{\theta,\widetilde{\theta}}\right)^2 d\widetilde{\theta} \int_{\underline{\theta}}^{\overline{\theta}} \left(-\overline{\epsilon}^R_{\widetilde{\theta}} \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))}\right)^2 d\widetilde{\theta}$$

and hence,

$$\begin{aligned} |A_n| \leq & \left[\sqrt{\int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} \left(\overline{\epsilon}^w_{\widehat{\theta}} \rho_{\widehat{\theta},\widetilde{\theta}}\right)^2 d\widehat{\theta} \, d\widetilde{\theta}}\right]^{n-1} \sqrt{\int_{\underline{\theta}}^{\overline{\theta}} \left(\rho_{\theta,\widehat{\theta}}\right)^2 d\widehat{\theta}} \sqrt{\int_{\underline{\theta}}^{\overline{\theta}} \left(-\overline{\epsilon}^R_{\widetilde{\theta}} \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))}\right)^2 d\widetilde{\theta}} \\ =: & \overline{A}_n \,. \end{aligned}$$

The sequence $\{\overline{A}_n\}_{n\in\mathbb{N}}$ is geometric. Moreover, since

$$\int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} \left(\overline{\epsilon}^w_{\widehat{\theta}} \rho_{\widehat{\theta},\widetilde{\theta}}\right)^2 d\widehat{\theta} \, d\widetilde{\theta} < 1 \,,$$

it converges to zero. Hence, $\{A_n\}_{n\in\mathbb{N}}$ is dominated in absolute value by a geometric sequence converging to zero. The series $\sum_{n=1}^{\infty} A_n$ is therefore convergent.

For convergence of $\sum_{n=1}^{\infty} B_n$, I show that $\sum_{n=1}(A_n + B_n)$ converges. Convergence of $\sum_{n=1}^{\infty} A_n$ and $\sum_{n=1}^{\infty}(A_n + B_n)$ then immediately implies convergence of $\sum_{n=1}^{\infty} B_n$. By definition we have

$$A_n + B_n = \int_{\underline{\theta}}^{\overline{\theta}} (\gamma^{(n)}_{\theta,\widetilde{\theta}} + \rho^{(n)}_{\theta,\widetilde{\theta}})(-\overline{\epsilon}^R_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))} \, d\widetilde{\theta}$$

and

$$\gamma^{(n)}_{\theta,\widetilde{\theta}} + \rho^{(n)}_{\theta,\widetilde{\theta}} = \int_{\underline{\theta}}^{\overline{\theta}} (\gamma^{(n-1)}_{\theta,\widehat{\theta}} + \rho^{(n-1)}_{\theta,\widehat{\theta}}) \overline{\epsilon}^{w}_{\widehat{\theta}} (\gamma_{\widehat{\theta},\widetilde{\theta}} + \rho_{\widehat{\theta},\widetilde{\theta}})\, d\widehat{\theta}\ .$$

Convergence of $\sum_{n=1}^{\infty}(A_n + B_n)$ now follows from exactly the same steps as convergence of $\sum_{n=1}^{\infty} A_n$, with the only difference that $\rho_{\theta,\widetilde{\theta}}$ is replaced by $\gamma_{\theta,\widetilde{\theta}} + \rho_{\theta,\widetilde{\theta}}$ in every step. Following these steps, the condition

$$\int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} \left(\overline{\epsilon}^{w}_{\widehat{\theta}} (\gamma_{\widehat{\theta},\widetilde{\theta}} + \rho_{\widehat{\theta},\widetilde{\theta}})\right)^2 \, d\widehat{\theta}\, d\widetilde{\theta} < 1$$

implies that the sequence $\{A_n + B_n\}_{n\in\mathbb{N}}$ is dominated in absolute value by a geometric sequence converging to zero, which establishes convergence of $\sum_{n=1}^{\infty}(A_n + B_n)$.

Finally, consider the limits

$$\lim_{N\to\infty} \int_{\underline{\theta}}^{\overline{\theta}} \rho^{(N)}_{\theta,\widehat{\theta}} \widehat{l}_{\widehat{\theta},\tau}(T)\, d\widehat{\theta} \quad \text{and} \quad \lim_{N\to\infty} \int_{\underline{\theta}}^{\overline{\theta}} \gamma^{(N)}_{\theta,\widehat{\theta}} \widehat{l}_{\widehat{\theta},\tau}(T)\, d\widehat{\theta}\ .$$

We have already shown that

$$\sum_{n=1}^{\infty} \int_{\underline{\theta}}^{\overline{\theta}} \rho^{(n)}_{\theta,\widetilde{\theta}} (-\overline{\epsilon}^{R}_{\widetilde{\theta}}) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1 - T'(y_{\widetilde{\theta}}(T))}\, d\widetilde{\theta}$$

converges independently of the specific values of $\overline{\epsilon}^{R}_{\widetilde{\theta}}$ and $\tau'(y_{\widetilde{\theta}}(T))/(1 - T'(y_{\widetilde{\theta}}(T)))$. Thus, by the same reasoning the series

$$\sum_{n=1}^{\infty} \int_{\underline{\theta}}^{\overline{\theta}} \rho^{(n)}_{\theta,\widetilde{\theta}} (-\overline{\epsilon}^{R}_{\widetilde{\theta}}) \widehat{l}_{\widetilde{\theta},\tau}\, d\widetilde{\theta}$$

converges. We must therefore have

$$\lim_{N\to\infty} \int_{\underline{\theta}}^{\overline{\theta}} \rho^{(N)}_{\theta,\widehat{\theta}} \widehat{l}_{\widehat{\theta},\tau}(T)\, d\widehat{\theta} = 0\ .$$

Analogous reasoning shows that

$$\lim_{N\to\infty} \int_{\underline{\theta}}^{\overline{\theta}} \gamma^{(N)}_{\theta,\widehat{\theta}} \widehat{l}_{\widehat{\theta},\tau}(T)\, d\widehat{\theta} = 0\ .$$

So, we have shown that, as $N \to \infty$, the induction hypothesis of step 1 becomes

$$\begin{aligned}\widehat{l}_{\theta,\tau} = & -\bar{\epsilon}_\theta^R \frac{\tau'(y_\theta(T))}{1-T'(y_\theta(T))} + \bar{\epsilon}_\theta^w \int_{\underline{\theta}}^{\overline{\theta}} \sum_{n=1}^{\infty} \rho_{\theta,\widetilde{\theta}}^{(n)} (-\bar{\epsilon}_{\widetilde{\theta}}^R) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))}\, d\widetilde{\theta} \\ & + \bar{\epsilon}_\theta^w \int_{\underline{\theta}}^{\overline{\theta}} \sum_{n=1}^{\infty} \gamma_{\theta,\widetilde{\theta}}^{(n)} (-\bar{\epsilon}_{\widetilde{\theta}}^R) \frac{\tau'(y_{\widetilde{\theta}}(T))}{1-T'(y_{\widetilde{\theta}}(T))}\, d\widetilde{\theta}\ ,\end{aligned}$$

which proves Lemma 17. □

Welfare Effects of Tax Reforms

Proof of Proposition 13. Welfare can be written as

$$\widetilde{W}(T) = V\left(\{R_T\left(w_\theta(T,\phi^*(T))l_\theta(T,w_\theta)\right) + S\left(l(T,w),w(T,\phi^*(T)),T\right) - v\left(l_\theta(T,w_\theta)\right)\}_{\theta\in\Theta}\right)\ ,$$

where

$$S\left(l(T,w),w(T,\phi^*(T)),T\right) = \int_{\underline{\theta}}^{\overline{\theta}} T\left(w_\theta(T,\phi^*(T))l_\theta(T,w_\theta)\right) h_\theta\, d\theta\ .$$

Taking the derivative D_τ yields:

$$\begin{aligned}D_\tau \widetilde{W}(T) = & \int_{\underline{\theta}}^{\overline{\theta}} \left(-g_\theta h_\theta \tau(w_\theta l_\theta) + g_\theta h_\theta \int_{\underline{\theta}}^{\overline{\theta}} h_{\widetilde{\theta}} \tau(w_{\widetilde{\theta}} l_{\widetilde{\theta}})\, d\widetilde{\theta}\right) d\theta \\ & + \int_{\underline{\theta}}^{\overline{\theta}} g_\theta h_\theta \int_{\underline{\theta}}^{\overline{\theta}} h_{\widetilde{\theta}} T'(w_{\widetilde{\theta}} l_{\widetilde{\theta}}) w_{\widetilde{\theta}} (-\epsilon_{\widetilde{\theta}}^R) \frac{l_{\widetilde{\theta}}}{1-T'(w_{\widetilde{\theta}} l_{\widetilde{\theta}})} \tau'(w_{\widetilde{\theta}} l_{\widetilde{\theta}})\, d\widetilde{\theta}\, d\theta \\ & + \int_{\underline{\theta}}^{\overline{\theta}} g_\theta h_\theta (1-T'(w_\theta l_\theta)) l_\theta D_{\phi,\tau} w_\theta(T,\phi^*(T))\, d\theta \\ & + \int_{\underline{\theta}}^{\overline{\theta}} g_\theta h_\theta \int_{\underline{\theta}}^{\overline{\theta}} h_{\widetilde{\theta}} T'(w_{\widetilde{\theta}} l_{\widetilde{\theta}}) \left(l_{\widetilde{\theta}} D_{\phi,\tau} w_{\widetilde{\theta}}(T,\phi^*(T)) + w_{\widetilde{\theta}} \epsilon_{\widetilde{\theta}}^w \frac{l_{\widetilde{\theta}}}{w_{\widetilde{\theta}}} D_{\phi,\tau} w_{\widetilde{\theta}}(T,\phi^*(T))\right) d\widetilde{\theta}\, d\theta \\ & + \int_{\underline{\theta}}^{\overline{\theta}} g_\theta h_\theta (1-T'(w_\theta l_\theta)) l_\theta D_\tau w_\theta(T,\phi^*(T))\, d\theta \\ & + \int_{\underline{\theta}}^{\overline{\theta}} g_\theta h_\theta \int_{\underline{\theta}}^{\overline{\theta}} h_{\widetilde{\theta}} T'(w_{\widetilde{\theta}} l_{\widetilde{\theta}}) \left(l_{\widetilde{\theta}} D_\tau w_{\widetilde{\theta}}(T,\phi^*(T)) + w_{\widetilde{\theta}} \epsilon_{\widetilde{\theta}}^w \frac{l_{\widetilde{\theta}}}{w_{\widetilde{\theta}}} D_\tau w_{\widetilde{\theta}}(T,\phi^*(T))\right) d\widetilde{\theta}\, d\theta\ .\end{aligned}$$

Using $\int_{\underline{\theta}}^{\overline{\theta}} h_\theta g_\theta \, d\theta = 1$, we can rearrange this expression to obtain

$$\begin{aligned} D_\tau \widetilde{W}(T) = &\int_{\underline{\theta}}^{\overline{\theta}} (1-g_\theta)\tau(w_\theta l_\theta) h_\theta \, d\theta + \int_{\underline{\theta}}^{\overline{\theta}} T'(w_\theta l_\theta) w_\theta l_\theta (-\epsilon_\theta^R) \frac{\tau'(w_\theta l_\theta)}{1-T'(w_\theta l_\theta)} h_\theta \, d\theta \\ &+ \int_{\underline{\theta}}^{\overline{\theta}} \left[g_\theta (1-T'(w_\theta l_\theta)) + T'(w_\theta l_\theta)(1+\epsilon_\theta^w) \right] l_\theta D_{\phi,\tau} w_\theta(T, \phi^*(T)) h_\theta \, d\theta \\ &+ \int_{\underline{\theta}}^{\overline{\theta}} \left[g_\theta (1-T'(w_\theta l_\theta)) + T'(w_\theta l_\theta)(1+\epsilon_\theta^w) \right] l_\theta D_\tau w_\theta(T, \phi^*(T)) h_\theta \, d\theta \ . \end{aligned}$$

□

Proof of Proposition 14. Take any initial tax $T \in \mathcal{T}^{ex}$ and any reform τ^{ex} that is progressive and raises welfare when neglecting induced technical change effects, that is, $D_{\tau^{ex}}^{ex} \widetilde{W}(T) > 0$ (see equation (B.41)). The strategy of the proof is to construct another progressive reform τ^{en} that raises welfare when accounting for induced technical change effects, that is, $D_{\tau^{en}} \widetilde{W}(T) > 0$. Constructing such a reform proves that $T \in \mathcal{T}$ and hence $\mathcal{T}^{ex} \subseteq \mathcal{T}$.

I will construct the reform τ^{en} such that it exactly replicates the labor input changes that τ^{ex} would induce if there were no induced technical change effects. To move back and forth between induced labor input changes and progressive reforms, I use Lemma 14. In particular, note that Proposition 14 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 14 then says that any progressive reform τ induces labor input responses $\widehat{l}_{\theta,\tau}$ that decrease in θ and, conversely, any reform that induces labor input changes that decrease in θ is progressive.

After these preparations, take any $T \in \mathcal{T}^{ex}$ and a progressive reform τ^{ex} that would raise welfare if technology remained constant, that is, $D_{\tau^{ex}}^{ex} \widetilde{W}(T) > 0$. We can write welfare as a function of consumption and labor inputs only, that is,

$$W(c,l) \coloneqq V\left(\{u_\theta(c_\theta, l_\theta)\}_{\theta \in \Theta}\right) \ .$$

Then, the effect of reform τ^{ex} on welfare, ignoring induced 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 induced technical change effects by

$$D_{\tau^{ex}}^{ex} l_\theta(T) \coloneqq D_{\tau^{ex}} l_\theta(T) \big|_{\rho_{\theta,\widetilde{\theta}} = 0 \ \forall \theta, \widetilde{\theta}}$$

and similarly the consumption response that ignores induced 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 $\widetilde{T}$, consumption and labor allocations must satisfy

$$c_{\widetilde{\theta}}(\widetilde{T}) - v\left(\frac{w_{\widetilde{\theta}}(\widetilde{T}, \phi^*(\widetilde{T}))l_{\widetilde{\theta}}(\widetilde{T})}{w_\theta(\widetilde{T}, \phi^*(\widetilde{T}))}\right) \leq c_\theta(\widetilde{T}) - v(l_\theta(\widetilde{T})) \quad \text{for all } \theta, \widetilde{\theta} \ .$$

Via an envelope argument this implies

$$c'_\theta(\widetilde{T}) = v'(l_\theta(\widetilde{T}))\left[l'_\theta(\widetilde{T}) + \widehat{w}_\theta(\widetilde{T}, \phi^*(\widetilde{T}))l_\theta(\widetilde{T})\right] \quad \text{for all } \theta \ .$$

Here, $\widehat{w}_\theta = w'_\theta / w_\theta$ 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:

$$c_\theta(\widetilde{T}) = c_{\underline{\theta}}(\widetilde{T}) + \int_{\underline{\theta}}^{\theta} v'(l_{\widetilde{\theta}}(\widetilde{T}))\left[l'_{\widetilde{\theta}}(\widetilde{T}) + \widehat{w}_{\widetilde{\theta}}(\widetilde{T}, \phi^*(\widetilde{T}))l_{\widetilde{\theta}}(\widetilde{T})\right] d\widetilde{\theta} \quad \text{for all } \theta \ . \tag{B.44}$$

The level $c_{\underline{\theta}}$ is determined via the resource constraint:

$$\int_{\underline{\theta}}^{\overline{\theta}} c_\theta(\widetilde{T})h_\theta \, d\theta = F(l(\widetilde{T}), \phi^*(l(\widetilde{T}))) \ . \tag{B.45}$$

Using equation (B.44), the response of consumption to tax reform τ^{ex}, ignoring induced technical change effects, can be expressed as

$$\begin{aligned} D^{ex}_{\tau^{ex}} c_\theta(T) = D^{ex}_{\tau^{ex}} c_{\underline{\theta}}(T) &+ \int_{\underline{\theta}}^{\theta} v''(l_{\widetilde{\theta}}(T))\left(D^{ex}_{\tau^{ex}} l_{\widetilde{\theta}}(T)\right)\left[l'_{\widetilde{\theta}}(T) + \widehat{w}_{\widetilde{\theta}}(T, \phi^*(T))l_{\widetilde{\theta}}(T)\right] d\widetilde{\theta} \\ &+ \int_{\underline{\theta}}^{\theta} v'(l_{\widetilde{\theta}}(T))\left[D^{ex}_{\tau^{ex}} l'_{\widetilde{\theta}}(T) + \widehat{w}_{\widetilde{\theta}}(T, \phi^*(T))D^{ex}_{\tau^{ex}} l_{\widetilde{\theta}}(T)\right] d\widetilde{\theta} \\ &+ \int_{\underline{\theta}}^{\theta} v'(l_{\widetilde{\theta}}(T))l_{\widetilde{\theta}}(T) D^{ex}_{\tau^{ex}} \widehat{w}_{\widetilde{\theta}}(T, \phi^*(T)) \, d\widetilde{\theta} \ . \end{aligned} \tag{B.46}$$

Note that the last line contains only the constant-technology effect $D^{ex}_{\tau^{ex}} \widehat{w}_{\widetilde{\theta}}(T, \phi^*(T))$ but not the induced technical change effect $D^{ex}_{\phi, \tau^{ex}} \widehat{w}_{\widetilde{\theta}}(T, \phi^*(T))$. The resource constraint (B.45) implies

$$\int_{\underline{\theta}}^{\overline{\theta}} h_\theta D^{ex}_{\tau^{ex}} c_\theta(T) \, d\theta = \frac{d}{d\mu} F(l(T + \mu\tau^{ex}), \phi^*(l(T)))\bigg|_{\mu=0} \ . \tag{B.47}$$

Step 2. Suppose now that we can find a reform τ^{en} that replicates the labor input change $D^{ex}_{\tau^{ex}} l_\theta(T)$ while accounting for induced technical change effects. That is, take τ^{en} such that

$$D_{\tau^{en}} l_\theta(T) = D^{ex}_{\tau^{ex}} l_\theta(T) \quad \text{for all } \theta \ .$$

I verify below that such a reform exists. Again using equation (B.44), the consumption response to τ^{en}, also accounting for induced technical change effects, can be expressed as

$$\begin{aligned} D_{\tau^{en}} c_\theta(T) = D_{\tau^{en}} c_{\underline{\theta}}(T) &+ \int_{\underline{\theta}}^{\theta} v''(l_{\widetilde{\theta}}(T)) \left(D_{\tau^{en}} l_{\widetilde{\theta}}(T)\right) \left[l'_{\widetilde{\theta}}(T) + \widehat{w}_{\widetilde{\theta}}(T, \phi^*(T)) l_{\widetilde{\theta}}(T)\right] d\widetilde{\theta} \\ &+ \int_{\underline{\theta}}^{\theta} v'(l_{\widetilde{\theta}}(T)) \left[D_{\tau^{en}} l'_{\widetilde{\theta}}(T) + \widehat{w}_{\widetilde{\theta}}(T, \phi^*(T)) D_{\tau^{en}} l_{\widetilde{\theta}}(T)\right] d\widetilde{\theta} \\ &+ \int_{\underline{\theta}}^{\theta} v'(l_{\widetilde{\theta}}(T)) l_{\widetilde{\theta}}(T) \left[D_{\tau^{en}} \widehat{w}_{\widetilde{\theta}}(T, \phi^*(T)) + D_{\phi, \tau^{en}} \widehat{w}_{\widetilde{\theta}}(T, \phi^*(T))\right] d\widetilde{\theta} \ . \end{aligned} \tag{B.48}$$

Note that here the last line contains the total effect of τ^{en} on the wage growth rate $\widehat{w}_{\widetilde{\theta}}$, that is, the sum of the direct and the induced technical change effect. The resource constraint (B.45) now implies

$$\int_{\underline{\theta}}^{\overline{\theta}} h_\theta D_{\tau^{en}} c_\theta(T) \, d\theta = \frac{d}{d\mu} F(l(T + \mu\tau^{en}), \phi^*(l(T + \mu\tau^{en}))) \bigg|_{\mu=0} \ . \tag{B.49}$$

The principle of taxation says that every incentive compatible and resource feasible consumption-labor allocation can be implemented by some tax $\widetilde{T}$. 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 $\widetilde{\tau}$. 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 induced technical change effects with the welfare effect of reform τ^{en} while accounting for induced technical change effects. Since the labor input changes are identical in both scenarios, the only difference in the two welfare effects stems from the different consumption responses:

$$D_{\tau^{en}} \widetilde{\widetilde{W}}(T) - D^{ex}_{\tau^{ex}} \widetilde{\widetilde{W}}(T) = \int_{\underline{\theta}}^{\overline{\theta}} g_\theta h_\theta \left(D_{\tau^{en}} c_\theta(T) - D^{ex}_{\tau^{ex}} c_\theta(T)\right) \, d\theta \ . \tag{B.50}$$

From equations (B.46) and (B.48), the difference in consumption responses can be expressed as, for every θ,

$$D_{\tau^{en}} c_\theta(T) - D^{ex}_{\tau^{ex}} c_\theta(T) = \int_{\underline{\theta}}^{\theta} v'(l_{\widetilde{\theta}}(T)) l_{\widetilde{\theta}}(T) D_{\phi,\tau^{en}} \widehat{w}_{\widetilde{\theta}}(T, \phi^*(T))\, d\widetilde{\theta} \ . \tag{B.51}$$

Here I used that the labor response is the same in both scenarios, such that the constant-technology effect on the wage growth rate is the same as well, that is,

$$D^{ex}_{\tau^{ex}} \widehat{w}_{\widetilde{\theta}}(T, \phi^*(T)) = D_{\tau^{en}} \widehat{w}_{\widetilde{\theta}}(T, \phi^*(T)) \ .$$

Next, consider the induced technology effect on the wage growth rate:

$$\begin{aligned} D_{\phi,\tau^{en}} \widehat{w}_{\widetilde{\theta}}(T, \phi^*(T)) =& D_{\phi,\tau^{en}} \left[\frac{d}{d\theta} \log\left(w_{\widetilde{\theta}}(T, \phi^*(T)) \right) \right] \\ =& \frac{d}{d\theta} \left[D_{\phi,\tau^{en}} \log\left(w_{\widetilde{\theta}}(T, \phi^*(T)) \right) \right] \\ =& \frac{d}{d\theta} \left[\frac{1}{w_{\widetilde{\theta}}(T, \phi^*(T))} D_{\phi,\tau^{en}} w_{\widetilde{\theta}}(T, \phi^*(T)) \right] \ . \end{aligned}$$

Since τ^{ex} is progressive, the labor response $(1/l_{\widetilde{\theta}}) D^{ex}_{\tau^{ex}} l_{\widetilde{\theta}}(T)$ is decreasing in $\widetilde{\theta}$ by Lemma 14. Hence, the identical response $(1/l_{\widetilde{\theta}}) D_{\tau^{en}} l_{\widetilde{\theta}}(T)$ decreases in $\widetilde{\theta}$ as well. Then by Lemma 6, the induced technical change effect $(1/w_{\widetilde{\theta}}) D_{\phi,\tau^{en}} w_{\widetilde{\theta}}(T, \phi^*(T))$ must also decrease in $\widetilde{\theta}$. We therefore obtain

$$\begin{aligned} 0 \geq& \frac{d}{d\theta} \left[\frac{1}{w_{\widetilde{\theta}}(T, \phi^*(T))} D_{\phi,\tau^{en}} w_{\widetilde{\theta}}(T, \phi^*(T)) \right] \\ =& D_{\phi,\tau^{en}} \widehat{w}_{\widetilde{\theta}}(T, \phi^*(T)) \ . \end{aligned}$$

By equation (B.51), this implies that the consumption difference $D_{\tau^{en}} c_\theta(T) - D^{ex}_{\tau^{ex}} c_\theta(T)$ is decreasing in θ. Moreover, from equations (B.47) and (B.49), we have

$$\begin{aligned} \int_{\underline{\theta}}^{\overline{\theta}} h_\theta D_{\tau^{en}} c_\theta(T)\, d\theta =& \left. \frac{d}{d\mu} F(l(T + \mu\tau^{en}), \phi^*(l(T + \mu\tau^{en}))) \right|_{\mu=0} \\ =& \left. \frac{d}{d\mu} F(l(T + \mu\tau^{ex}), \phi^*(l(T))) \right|_{\mu=0} \\ =& \int_{\underline{\theta}}^{\overline{\theta}} h_\theta D^{ex}_{\tau^{ex}} c_\theta(T)\, d\theta \ , \end{aligned}$$

where the second equality uses an envelope argument. By implication:

$$\int_{\underline{\theta}}^{\overline{\theta}} h_\theta \left(D_{\tau^{en}} c_\theta(T) - D^{ex}_{\tau^{ex}} c_\theta(T) \right) \, d\theta \; .$$

So, inspecting equation (B.50) reveals that if g_θ were constant in θ, we would have

$$D_{\tau^{en}} \widetilde{W}(T) - D^{ex}_{\tau^{ex}} \widetilde{W}(T) = 0 \; .$$

But since both g_θ and $D_{\tau^{en}} c_\theta(T) - D^{ex}_{\tau^{ex}} c_\theta(T)$ are decreasing in θ, we must have

$$\int_{\underline{\theta}}^{\overline{\theta}} g_\theta h_\theta \left(D_{\tau^{en}} c_\theta(T) - D^{ex}_{\tau^{ex}} c_\theta(T) \right) \, d\theta \geq 0$$

and thereby

$$D_{\tau^{en}} \widetilde{W}(T) - D^{ex}_{\tau^{ex}} \widetilde{W}(T) \geq 0 \; .$$

So,

$$D_{\tau^{en}} \widetilde{W}(T) > 0 \; .$$

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 14 to obtain that τ^{en} must be progressive. We have thereby shown that

$$D_{\tau^{en}} \widetilde{W}(T) > 0$$

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 $\mathcal{T}^{ex} \subseteq \mathcal{T}$. □

Alternative Optimal Tax Formula

Proof of Proposition 15. To derive equation (B.42), we start from equation (B.32) and replace $\widehat{l}_\theta$ by $\epsilon^w_\theta \widehat{w}_\theta$:

$$\frac{T'(y_\theta)}{1 - T'(y_\theta)} = \left(1 + \frac{1}{e_\theta}\right) \frac{1 - H_\theta}{h_\theta} (1 - \widetilde{g}_\theta) \widehat{w}_\theta \left(1 - (\gamma^{CES} + \rho^{CES}) \epsilon^w_\theta \right) + \gamma^{CES} (1 - g_\theta) + \rho^{CES} (1 - g_\theta) \; .$$

Using m_y and M_y to denote density and cumulative distribution function of income, a change-of-variable implies $h_\theta = m_{y_\theta} y' - \theta$. Using this in the previous expression for marginal tax rates, we obtain:

$$\begin{aligned}\frac{T'(y_\theta)}{1-T'(y_\theta)} &= \left(1+\frac{1}{e_\theta}\right)\frac{1-M_{y_\theta}}{m_{y_\theta}y_\theta}\frac{1}{\widehat{w}_\theta+\widehat{l}_\theta}(1-\widetilde{g}_\theta)\widehat{w}_\theta\left(1-(\gamma^{CES}+\rho^{CES})\epsilon^w_\theta\right) \\ &\quad+\gamma^{CES}(1-g_\theta)+\rho^{CES}(1-g_\theta) \\ &= \left(1+\frac{1}{e_\theta}\right)\frac{1-M_{y_\theta}}{m_{y_\theta}y_\theta}\frac{1}{\widehat{w}_\theta(1+\epsilon^w_\theta)}(1-\widetilde{g}_\theta)\widehat{w}_\theta\left(1-(\gamma^{CES}+\rho^{CES})\epsilon^w_\theta\right) \\ &\quad+\gamma^{CES}(1-g_\theta)+\rho^{CES}(1-g_\theta) \\ &= \left(1+\frac{1}{e_\theta}\right)\frac{1-M_{y_\theta}}{m_{y_\theta}y_\theta}\frac{1}{1+\epsilon^w_\theta}(1-\widetilde{g}_\theta)\left(1-(\gamma^{CES}+\rho^{CES})\epsilon^w_\theta\right) \\ &\quad+\gamma^{CES}(1-g_\theta)+\rho^{CES}(1-g_\theta)\ .\end{aligned}$$

From the expression for ϵ^w_θ and ϵ^R_θ in Section 3.3.2, it is straightforward to show that

$$\left(1+\frac{1}{e_\theta}\right)\frac{1}{1+\epsilon^w_\theta}=\frac{1}{\epsilon^R_\theta}\ ,$$

and hence

$$\frac{T'(y_\theta)}{1-T'(y_\theta)}=\frac{1-(\gamma^{CES}+\rho^{CES})\epsilon^w_\theta}{\epsilon^R_\theta}\frac{1-M_{y_\theta}}{m_{y_\theta}y_\theta}(1-\widetilde{g}_\theta)+\gamma^{CES}(1-g_\theta)+\rho^{CES}(1-g_\theta)\ .$$

By definition of the aggregate elasticity $\bar{\epsilon}^R_\theta$ (see Appendix B.3.1), this yields equation (B.42) from Proposition 5:

$$\frac{T'(y_\theta)}{1-T'(y_\theta)}=\frac{1}{\bar{\epsilon}^R_\theta}\frac{1-M_{y_\theta}}{m_{y_\theta}y_\theta}(1-\widetilde{g}_\theta)+\gamma^{CES}(1-g_\theta)+\rho^{CES}(1-g_\theta)\ .$$

□

C Appendix to Chapter 4

C.1 Omitted Proofs

This section collects all proofs omitted from the main text.

C.1.1 Discussion of Assumption (4.6)

To gain insights about the properties of W, v, and c, under which assumption (4.6) holds, we rewrite the expert's quality choice as a direct choice of consumer utility. In particular, using

$$u_b = v(a_b) - p_b \ ,$$

we can write expert utility as (in the bilateral setting of Section 4.4)

$$\overline{u}_e = W(p_b - \tilde{c}(u_b + p_b)) + u_b \ .$$

Here, $\tilde{c}(x) \equiv c(v^{-1}(x))$ measures the cost of providing utility-from-treatment (i.e., utility gross of the price) of x to the consumer.

Given a price offer p_b, expert e now chooses u_b to maximize her utility $\overline{u}_e$. The derivative of $\overline{u}_e$ with respect to u_b is

$$-W'(p_b - \tilde{c}(u_b + p_b))\tilde{c}'(u_b + p_b) + 1 \ .$$

By assumption (4.4), this derivative is weakly negative at $p_b = \tilde{c}(v(0))$ and $u_b = v(0) - \tilde{c}(v(0))$. For concreteness, suppose now that this assumption indeed holds exactly, that is,

$$W'(0)\tilde{c}'(v(0)) = 1 \ .$$

Then, since the second derivative of $\overline{u}_e$ with respect to u_b is strictly negative everywhere, the effect of raising p_b on the optimal choice u_b is qualitatively given by the sign of

$$\frac{\partial^2 \overline{u}_e}{\partial u_b \partial p_b} = -W''(p_b - \tilde{c}(u_b + p_b))(1 - \tilde{c}(u_b + p_b))\tilde{c}'(u_b + p_b) - W'(p_b - \tilde{c}(u_b + p_b))\tilde{c}''(u_b + p_b) \ .$$

At the competitive equilibrium values $p_b = \tilde{c}(v(0))$ and $u_b = v(0) - \tilde{c}(v(0))$, the cross-derivative is positive if and only if

$$-\frac{W''(0)}{W'(0)}(1 - \tilde{c}'(v(0))) > \frac{\tilde{c}''(v(0))}{\tilde{c}(v(0))} .$$

If and only if this is satisfied, assumption (4.6) holds and the collectively optimal price offer of consumers exceeds the competitive price $\tilde{c}(v(0))$. Hence, for this to be true, the cost function $\tilde{c}$ must have sufficiently small curvature at $v(0)$. Put differently, the marginal cost of providing additional utility-from-treatment to consumers must not increase too quickly around the competitive equilibrium.

C.1.2 Proof of Lemma 9

Let

$$A(p_b, y_e, a_b) := W'(y_e)\,[p_b - c(a_b)] + v(a_b) - p_b$$

denote the marginal utility for expert e of adding consumer b to her set of clients B_e if she provides quality a_b to b.

Expert e's actual quality choice for consumer b follows from conditions (4.7) as a function of the expert's income y_e. Denote this quality by $a_b^{IC}(y_e)$. Then, the expert's actual marginal utility from serving consumer b, taking into account her quality choice $a_b^{IC}(y_e)$, becomes

$$A^{IC}(p_b, y_e) := A(p_b, y_e, a_b^{IC}(y_e)) .$$

Expert e will accept an offer p_b if and only if $A^{IC}(p_b, y_e) \geq 0$. Hence, the equality $A^{IC} = 0$ defines the acceptance threshold described by Lemma 9.

Before deriving the claimed properties of the threshold, note that

$$A^{IC}(p_b, y_e) = \max_{a \geq 0} A(p_b, y_e, a) \tag{C.1}$$

by definition of $a_b^{IC}(y_e)$. In words, the expert chooses the service quality for b such as to maximize her utility from serving b.

Case 1: $y_e \leq 0$. By assumption (4.4), we have $a_b^{IC}(y_e) = 0$ for all $y_e \leq 0$. Hence, $A^{IC}(c(0), y_e) = 0$ for all $y_e \leq 0$. That is, if the expert has negative income, she just accepts an offer at $c(0)$. Since A^{IC} is strictly increasing in p_b, we have that for all $y_e \leq 0$, $A^{IC}(p_b, y_e) \geq 0$ if and only if $p_b \geq c(0)$. This proves the first piece of the acceptance threshold in Lemma 9.

Case 2: $y_e > 0$. As in Lemma 9, denote the acceptance threshold for $y_e > 0$ by $\tilde{p}(y_e)$, that is, $A^{IC}(\tilde{p}(y_e), y_e) = 0$.

First note that $A(c(0), y_e, 0) = 0$ for all y_e. Hence, $A^{IC}(c(0), y_e) \geq 0$ for all y_e. Therefore, the acceptance threshold satisfies $\tilde{p}(y_e) \leq c(0)$ for all y_e.

It remains to show that $\tilde{p}(y_e)$ is decreasing in y_e. For that, consider $y_e^{(2)} > y_e^{(1)} > 0$. From the definition of A we see that A is increasing in y_e if $p_b \leq c(0)$. Since $\tilde{p}(y_e^{(1)}) \leq c(0)$, we obtain the following inequalities:

$$\begin{aligned} A^{IC}\left(\tilde{p}\left(y_e^{(2)}\right), y_e^{(2)}\right) &= 0 \\ &= A\left(\tilde{p}\left(y_e^{(1)}\right), y_e^{(1)}, a_b^{IC}\left(y_e^{(1)}\right)\right) \\ &\leq A\left(\tilde{p}\left(y_e^{(1)}\right), y_e^{(2)}, a_b^{IC}\left(y_e^{(1)}\right)\right) \\ &\overset{\text{by (C.1)}}{\leq} A\left(\tilde{p}\left(y_e^{(1)}\right), y_e^{(2)}, a_b^{IC}\left(y_e^{(2)}\right)\right) \\ &= A^{IC}\left(\tilde{p}\left(y_e^{(1)}\right), y_e^{(2)}\right) \ . \end{aligned}$$

Using that A^{IC} is always increasing in p_b, the inequality between the first and the last expression implies $\tilde{p}(y_e^{(2)}) \leq \tilde{p}(y_e^{(1)})$.

C.1.3 Proof of Proposition 6

We prove Proposition 6 via the following lemma.

Lemma 18. *Consider the game described by stages 1 and 2. In any subgame perfect equilibrium all offers are symmetric, $p_b = p_{b'}$ for all $b, b' \in B$, all offers are accepted, and all quality levels are symmetric, $a_b = a_{b'}$ for all $b, b' \in B$.*

Proof. **Step 1.** The thresholds in Lemma 9 imply that an offer $p_b = c(0)$ is always accepted. Since $v(0) - c(0) \geq \underline{v}$ and agents always opt against their outside option in case of indifference, consumers always prefer to make the offer $c(0)$ over any offer that is not accepted. Hence, offers that are not accepted are strictly dominated and cannot be part of a subgame perfect equilibrium.

Step 2. Consider now all consumers $b \in B_e$ for a given expert e. By the Kuhn-Tucker conditions (4.7), these consumers all receive the same quality level. Moreover, they face the same acceptance threshold. Since all consumers take expert e's income as given, they anticipate the quality they receive to be independent of their offers. Hence, they offer exactly the acceptance threshold, which is the same across all consumers.

Step 3. By Step 2, any expert e receives the same offers from all consumers matched to her. Suppose now that these offers are strictly higher for some expert e than for another expert e'. Denote the offer level for e by p and for e' by p'. By Step 1, all offers are accepted. So, experts' revenue equals their offer level,

$$\int_{B_e} p_b \, db = p > \int_{B_{e'}} p_b \, db = p' \, .$$

Using this in the Kuhn-Tucker conditions (4.7), it is easy to show that expert e will also have greater income than expert e', $y_e \geq y_{e'}$. But then, by Lemma 9, the acceptance threshold of expert e is smaller than that of expert e'. Hence, consumers matched to e offer lower payments than consumers matched to e'. This contradicts the initial assumption of $p > p'$.

We have therefore established that all consumers offer the same payments and all offers are accepted in any subgame perfect equilibrium. The Kuhn-Tucker conditions (4.7) then immediately imply that quality levels are the same for all consumers in any subgame perfect equilibrium as well. □

Proposition 6 is now proven as follows. By Lemma 18, there is a common offer level $p = p_b$ for all $b \in B$. By Lemma 9, offers $p_b = c(0)$ are always accepted. Moreover, consumers always offer payments exactly equal to the expert's acceptance threshold. So, the common offer level p can be at most $c(0)$.

Suppose that $p < c(0)$. Then, any expert e has negative income, $y_e \leq 0$. But for $y_e \leq 0$, Lemma 9 says that offers below $c(0)$ are rejected. Hence, we must have $p = c(0)$ in any subgame perfect equilibrium. The Kuhn-Tucker conditions (4.7) then imply $a_b = 0$ for all $b \in B$ in any subgame perfect equilibrium.

C.1.4 Proof of Proposition 8

The only part of the proposition that remains to be shown is that an allocation is fully efficient if and only if $a_b = a^{**}$ for almost all $b \in B$.

(⇒) We first prove the "only if" part of the claim. To show that no allocation other than those described above is fully efficient, take an arbitrary allocation q, $\{p_b^q\}_{b \in B}$, $\{B_e\}_{e \in E}$, $\{a_b^q\}_{b \in B}$, with $a_b^q \neq a^{**}$ for some non-zero measure of consumers. Construct a new allocation r with $a_b^r = a^{**}$ for all $b \in B$, $B_e^r = B_e^q$ for all $e \in E$, and

$$p_b^r = p_b^q + v(a_b^r) - v(a_b^q) \, .$$

Comparing r to q, the utility of consumers is unchanged by construction of r. For an expert e the utility change is $W(y_e^r) - W(y_e^q)$. Its sign depends on the difference in incomes $y_e^r - y_e^q$. Using the construction of payments p_b^r in allocation r, this income difference becomes

$$y_e^r - y_e^q = \int_{B_e^q} \left[v(a_b^r) - c(a_b^r) - v(a_b^q) + c(a_b^q) \right] \, db \, .$$

Since a^{**} uniquely maximizes $v(a) - c(a)$, the income difference is positive, $y_e^r - y_e^q > 0$. Hence, experts strictly prefer allocation r to q. Since consumers are indifferent between the two, allocation r Pareto-dominates q. Allocation q can therefore not be fully efficient.

(⇐) To see that any allocation with $a_b = a^{**}$ for almost all b is fully efficient, suppose such an allocation (call it s) is Pareto-dominated by some other allocation (call it t). If t has $a_b \neq a^{**}$ for a non-zero measure of consumers, part (⇒) above implies that there exists an allocation t' with $a_b^{t'} = a^{**}$ almost everywhere that Pareto-dominates t. By transitivity, t' will then also Pareto-dominate s. Hence, we can focus on allocations t that feature $a_b^t = a^{**}$ for almost all b.

Allocations s and t then only differ in the distribution of payments over experts and consumers. Since this distribution is zero-sum, none of the allocations can Pareto-dominate the other. We have thereby established that any allocation with $a_b = a^{**}$ almost everywhere is fully efficient.

C.1.5 Proof of Lemma 10

Given a non-empty set of active experts E, the subgame described by stages 2' and 3' is very similar to the game with exogenous entry described by stages 1 and 2 in Section 4.5. The main difference is that expert e's marginal cost of serving an additional consumer b is $c(a_b) + k'(|B_e|)$ instead of $c(a_b)$ only. The proof of the acceptance threshold in Lemma 10 therefore proceeds in close analogy to the proof of the acceptance threshold from the exogenous entry setting in Lemma 9.

Let

$$\hat{A}(p_b, B_e, \hat{y}_e, a_b) := W'(\hat{y}_e) \left[p_b - c(a_b) - k'(|B_e|) \right] + v(a_b) - p_b - v(0) + k'(|B_e|)$$

denote expert e's marginal utility from adding consumer b to her set of clients B_e if she provides quality a_b to the consumer.

Expert e's actual quality choice follows from the Kuhn-Tucker conditions (4.7) as a function of $\hat{y}_e$. Denote this quality by $\hat{a}_b^{IC}(\hat{y}_e)$. Then, the expert's actual marginal utility from accepting the offer p_b, taking into account her quality choice $\hat{a}_b^{IC}(\hat{y}_e)$, becomes

$$\hat{A}^{IC}(p_b, B_e, \hat{y}_e) := \hat{A}(p_b, B_e, \hat{y}_e, \hat{a}_b^{IC}(\hat{y}_e)) \, .$$

Expert e will accept p_b if and only if $\hat{A}^{IC}(p_b, B_e, \hat{y}_e) \geq 0$. The equality $\hat{A}^{IC} = 0$ therefore defines the acceptance threshold from Lemma 10.

Note at this point that

$$\hat{A}^{IC}(p_b, B_e, \hat{y}_e) = \max_{a \geq 0} A(p_b, B_e, \hat{y}_e, a) \tag{C.2}$$

by definition of $\hat{a}_b^{IC}(\hat{y}_e)$.

Case 1: $\hat{y}_e \leq 0$. Assumption (4.4) implies $\hat{a}_b^{IC}(\hat{y}_e) = 0$ for all $\hat{y}_e \leq 0$. So, $\hat{A}^{IC}(k'(|B_e|), \hat{y}_e) = 0$ for all $\hat{y}_e \leq 0$. That is, at negative income the expert just accepts an offer at marginal cost $k'(|B_e|)$. Since $\hat{A}^{IC}$ is strictly increasing in p_b, it holds for all $\hat{y}_e \leq 0$ that $\hat{A}^{IC}(p_b, B_e, \hat{y}_e) \geq 0$ if and only if $p_b \geq k'(|B_e|)$. We have thus proven the first piece of the acceptance threshold in Lemma 10.

Case 2: $\hat{y}_e > 0$. Denote the acceptance threshold for $\hat{y}_e > 0$ by $\hat{p}(\hat{y}_e, B_e)$, that is, $\hat{A}^{IC}(\hat{p}(\hat{y}_e, B_e), B_e, \hat{y}_e) = 0$.

Note that $\hat{A}(k'(|B_e|), B_e, \hat{y}_e, 0) = 0$ for all $\hat{y}_e$ and B_e. Thus, $\hat{A}^{IC}(k'(|B_e|), B_e, \hat{y}_e) \geq 0$ for all $\hat{y}_e$ and B_e. Hence, we have $\hat{p}(\hat{y}_e, B_e) \leq k'(|B_e|)$ for all $\hat{y}_e$ and B_e.

It remains to prove that $\hat{p}(\hat{y}_e, B_e)$ is decreasing in $\hat{y}_e$. Take any B_e and any two income levels $\hat{y}_e^{(2)} > \hat{y}_e^{(1)} > 0$. From the definition of $\hat{A}$, it is clear that $\hat{A}$ increases in $\hat{y}_e$ if $p_b \leq k'(|B_e|)$. Since $\hat{p}(\hat{y}_e^{(1)}, B_e) \leq k'(|B_e|)$, the following applies:

$$\begin{aligned}
\hat{A}^{IC}\left(\hat{p}\left(\hat{y}_e^{(2)}, B_e\right), B_e, \hat{y}_e^{(2)}\right) &= 0 \\
&= A\left(\hat{p}\left(\hat{y}_e^{(1)}, B_e\right), B_e, y_e^{(1)}, \hat{a}_b^{IC}\left(\hat{y}_e^{(1)}\right)\right) \\
&\leq \hat{A}\left(\hat{p}\left(\hat{y}_e^{(1)}, B_e\right), B_e, \hat{y}_e^{(2)}, \hat{a}_b^{IC}\left(\hat{y}_e^{(1)}\right)\right) \\
&\overset{\text{by (C.2)}}{\leq} \hat{A}\left(\hat{p}\left(\hat{y}_e^{(1)}, B_e\right), B_e, \hat{y}_e^{(2)}, \hat{a}_b^{IC}\left(\hat{y}_e^{(2)}\right)\right) \\
&= \hat{A}^{IC}\left(\hat{p}\left(\hat{y}_e^{(1)}, B_e\right), B_e, \hat{y}_e^{(2)}\right) .
\end{aligned}$$

Since $\hat{A}^{IC}$ always increases in p_b, the inequality between the first and the last expression implies $\hat{p}(\hat{y}_e^{(2)}, B_e) \leq \hat{p}(\hat{y}_e^{(1)}, B_e)$.

C.1.6 Proof of Lemma 11

To prepare the proofs of Lemma 11 and Proposition 9, we prove the following lemma.

Lemma 19. *Take any non-empty set of active experts E and consider the subgame after E described by stages 2' and 3'. In any subgame perfect equilibrium of this subgame all offers are symmetric, $p_b = p_{b'}$ for all $b, b' \in B$, all offers are accepted, and all quality levels are symmetric, $a_b = a_{b'}$ for all $b, b' \in B$.*

Proof. Take a non-empty set of active experts E and consider the subgame after E described by stages 2' and 3'. This subgame is almost equivalent to the game with exogenous entry described by stages 1 and 2 in Section 4.5. Hence, the proof of Lemma 19 closely follows the proof of Lemma 18.

Step 1. The maximum size of B_e for any expert e is M. Hence, Lemma 10 implies that experts always accept an offer $p_b \geq k'(M)$. Since $v(0) - k'(M) \geq \underline{v}$ and agents always decide against their outside option in case of indifference, any consumer b prefers the offer $p_b = k'(M)$ over any offer that is not accepted. So, consumers only make offers that are accepted in equilibrium.

Step 2. This step is identical to step 2 in the proof of Lemma 18. We repeat it here for convenience. Consider all consumers $b \in B_e$ for a given expert e. By the Kuhn-Tucker conditions (4.7) (using $\hat{y}_e$ instead of y_e in the conditions), these consumers all receive the same quality level. Moreover, they face the same acceptance threshold. Since all consumers take expert e's income as given, they anticipate the quality they receive to be independent of their offers. Hence, they offer exactly the acceptance threshold, which is the same across all consumers.

Step 3. By Step 2, any expert e receives the same offers from all consumers matched to her. To derive a contradiction, suppose that these offers are strictly higher for some expert e than for another expert e'. Denote the offer level for e by p and for e' by p'. By Step 1, all offers are accepted. So, the revenues of e and e' are given by

$$\int_{B_e} p_b \, db = \frac{M}{N} p > \frac{M}{N} p' = \int_{B_{e'}} p_b \, db \, .$$

Using this together with the fact that $|B_e| = |B_{e'}|$, the Kuhn-Tucker conditions (4.7) imply that expert e will have a greater income than e', $\hat{y}_e \geq \hat{y}_{e'}$. Then, again because $|B_e| = |B_{e'}|$, Lemma 10 implies that the acceptance threshold of expert e is smaller than that of e'. So, consumers matched to e make smaller offers than those matched to e', contradicting the initial assumption $p > p'$.

We have therefore established that all consumers offer the same payments and all offers are accepted in any subgame perfect equilibrium. The Kuhn-Tucker conditions (4.7) then immediately imply that quality levels are the same for all consumers in any subgame perfect equilibrium as well. □

We prove now each of the three cases of Lemma 11. Since by Lemma 19 all offers are accepted, we can set $|B_e| = M/N$ for all active experts $e \in E$ throughout the proof.

1. We first show that $\hat{y}_e > 0$ for all $e \in E$. To derive a contradiction, suppose that $\hat{y}_e \leq 0$ for some $e \in E$. Using Lemma 10, this implies that all consumers $b \in B_e$ offer $p_b = k'(M/N)$. Moreover, the Kuhn-Tucker conditions (4.7) imply that $a_b = 0$ for all $b \in B_e$. But then we

obtain for expert e's income:

$$\hat{y}_e = \frac{M}{N} k'\left(\frac{M}{N}\right) - k\left(\frac{M}{N}\right) - F > 0 ,$$

a contradiction.

So, $\hat{y}_e > 0$ for all $e \in E$. From Lemma 10 we then obtain $p_b \leq k'(M/N)$ for all $b \in B$.

For experts' utility, note that $a_b \geq 0$ and $p_b \leq k'(M/N)$ for all b imply

$$v(a_b) - p_b - v(0) + k'\left(\frac{M}{N}\right) > 0 .$$

Hence, using $\hat{y}_e > 0$,

$$W(\hat{y}_E) + \int_{B_e} \left[v(a_b) - p_b - v(0) + k'\left(\frac{M}{N}\right)\right] db > W(0)$$

for all $e \in E$.

2. We show that $\hat{y}_e = 0$ for all $e \in E$. To derive a contradiction, suppose first that $\hat{y}_e > 0$ for some $e \in E$. But then $p_b \leq k'(M/N)$ for all $b \in B_e$ by Lemma 10. Together with $a_b \geq 0$ for all b, this implies

$$\hat{y}_e \leq \frac{M}{N} k'\left(\frac{M}{N}\right) - k\left(\frac{M}{N}\right) - F = 0 ,$$

a contradiction. Suppose now that $\hat{y}_e < 0$ for some $e \in E$. Then, $p_b = k'(M/N)$ for all $b \in B$ by Lemma 10. Moreover, expert e's quality choice yields $a_b = 0$ for all $b \in B_e$ by conditions (4.7). Hence we obtain for expert e's income:

$$\hat{y}_e = \frac{M}{N} k'\left(\frac{M}{N}\right) - k\left(\frac{M}{N}\right) - F = 0 ,$$

a contradiction.

So, $\hat{y}_e = 0$ for all $e \in E$. Using Lemma 10, we obtain $p_b = k'(M/N)$ for all $b \in B$.

Moreover, $\hat{y}_e = 0$ for all $e \in E$ implies $a_b = 0$ for all $b \in B$. So,

$$v(a_b) - p_b - v(0) + k'(M/N) = 0$$

for all $b \in B$. Experts' utility thus becomes

$$W(0) + \int_{B_e} \left[v(0) - k'\left(\frac{M}{N}\right) - v(0) + k'\left(\frac{M}{N}\right) \right] db = W(0)$$

for all $e \in E$.

3. We first show that $\hat{y}_e < 0$ for all $e \in E$. To derive a contradiction, suppose $\hat{y}_e \geq 0$ for some $e \in E$. Then, $p_b \leq k'(M/N)$ for all $b \in B_e$ by Lemma 10. Using $a_b \geq 0$ for all b, we obtain

$$\hat{y}_e \leq \frac{M}{N} k'\left(\frac{M}{N}\right) - k\left(\frac{M}{N}\right) - F < 0 \ ,$$

a contradiction.

So, $\hat{y}_e < 0$ for all $e \in E$. With Lemma 10 we then obtain $p_b = k'(M/N)$ for all $b \in B$.

Moreover, $\hat{y}_e < 0$ for all e implies $a_b = 0$ for all b. Experts' utility hence satisfies

$$W(\hat{y}_e) + \int_{B_e} \left[v(0) - k'\left(\frac{M}{N}\right) - v(0) + k'\left(\frac{M}{N}\right) \right] db < W(0)$$

for all $e \in E$.

C.1.7 Proof of Proposition 9

Since all offers are accepted by Lemma 19, we can again set $|B_e| = M/N$ throughout the proof.

From conditions (4.11) and (4.12), we have $M/N \to m$ as $M \to \infty$. Moreover,

$$\frac{M}{N} k'\left(\frac{M}{N}\right) - k\left(\frac{M}{N}\right) - F \to 0 \ .$$

We first show that $\hat{y}_e \to 0$ for all $e \in E$ as $M \to \infty$. For that, take any unbounded sequence of consumer masses M. To derive a contradiction, suppose first that there exists a subsequence such that $\hat{y}_e$ is positive and bounded away from zero along this subsequence. Since $p_b \leq k'(M/N)$ for all $b \in B$ by Lemma 10 and because $a_b \geq 0$ for all b, we have

$$\hat{y}_e \leq \frac{M}{N} k'\left(\frac{M}{N}\right) - k\left(\frac{M}{N}\right) - F \ .$$

But the right-hand-side of the inequality converges to zero along the subsequence. Hence, $\hat{y}_e$ cannot be positive and bounded away from zero.

Suppose now that there is a subsequence of consumer masses along which $\hat{y}_e$ remains negative and bounded away from zero for some $e \in E$. Then by Lemma 10, $p_b = k'(M/N)$ along the subsequence. Moreover, $a_b = 0$ for all $b \in B_e$ along the subsequence by conditions (4.7). Thus,

$$\hat{y}_e = \frac{M}{N} k' \left(\frac{M}{N} \right) - k \left(\frac{M}{N} \right) - F \to 0 \; ,$$

a contradiction.

We have therefore established that $\hat{y}_e \to 0$ for all $e \in E$ as $M \to \infty$. From conditions (4.7), we then immediately obtain $a_b \to 0$ for all $b \in B$.

Finally by Lemma 19, there is a common payment level p and a common quality level a for all consumers. Income of expert e thus becomes

$$\hat{y}_e = \frac{M}{N} p - \frac{M}{N} c(a) - k \left(\frac{M}{N} \right) - F \; ,$$

and hence

$$p = \frac{N}{M} \hat{y}_e + c(a) + \frac{N}{M} k \left(\frac{M}{N} \right) + \frac{N}{M} F \; .$$

Since $M/N \to m$, $a \to 0$, and $\hat{y}_e \to 0$, we can use the definition of m to show that the right-hand-side of the equation goes to $k'(m)$ as $M \to \infty$. Therefore, $p_b \to k'(m)$ for all $b \in B$.

C.1.8 Proof of Proposition 10

Part 1. Consider first the regulation $(\hat{p}^*, \hat{N})$. In the main text we have already shown that the proposed regulation Pareto-dominates the unregulated (or, competitive) outcome for sufficiently large M if the actual number of active experts $\tilde{N}$ equals the cap $\hat{N}$. To see that we will indeed have $\tilde{N} = \hat{N}$, consider the competitive outcome at a given M. From Proposition 9, it is easy to see that experts' utility in the competitive outcome approaches $W(0)$ as $M \to \infty$. Again from Proposition 9, we know that $p_b \to k'(m)$ for all b as $M \to \infty$. Hence, for sufficiently large M the regulated price $\hat{p}^*$ strictly exceeds the competitive price. Holding the number of active exerts constant at $\hat{N}$, an increase in the level of payments strictly increases experts' utility. So for large M and holding the number of experts at $\hat{N}$, experts' utility from the regulated price $\hat{p}^*$ strictly exceeds $W(0)$. But that means that all $\hat{N}$ experts indeed choose to enter the market under the regulation $(\hat{p}^*, \hat{N})$ for sufficiently large M. Hence, the cap of $\hat{N}$ is binding, $\tilde{N} = \hat{N}$.

Part 2. Consider next the pure price regulation $(\hat{p}^*, \infty)$. Denote the number of active experts under this regulation by $\tilde{N}$ and compare it to the regulated number of experts $\hat{N}$ from Part 1. By Part 1, experts' utility under the joint regulation $(\hat{p}^*, \hat{N})$ converges to a level strictly above $W(0)$. Moreover as $M \to \infty$, the impact of an additional entrant on experts' utility approaches zero. Hence, without entry regulation the expert $\hat{N}+1$ finds it beneficial to enter the market. So, $\tilde{N} > \hat{N}$. Since experts' utility declines in the number of active experts for given prices, experts' utility is strictly smaller under the pure price regulation than under the joint regulation of Part 1.

Moreover, suppose that experts' income $\hat{y}_e$ is greater under the pure price regulation than under the joint regulation. This would imply that service quality is higher under the pure price regulation as well. But with a higher service quality and a larger number of active experts, income must be strictly smaller under the pure price regulation than under joint regulation. Hence, experts' income is indeed strictly smaller under the pure price than under the joint regulation.

Finally, under the joint regulation we have $a_b > 0$ for all consumers. So experts' quality choice problem has an interior solution. In the neighborhood of such an interior solution, quality strictly decreases in income. So, service quality must be strictly smaller under the pure price regulation than under the joint regulation. Since the payments p_b are the same in both cases, we obtain that consumers' utility is strictly smaller under the pure price regulation than under the joint regulation. This establishes that the joint regulation Pareto-dominates the pure price regulation.

C.2 Price Competition

In this section we present an alternative trading mechanism where experts instead of consumers make price offers. The environment is the same as in the main text, that is, the one introduced in Section 3.3. The mechanism works as follows.

Stage 1'' Each expert $e \in E$ makes price offers $\{p_{e,b}\}_{b \in B}$ to all consumers.

Stage 2'' Each consumer $b \in B$ observes his offers $\{p_{e,b}\}_{e \in E}$ but not the offers received by other consumers. Consumer b then accepts or rejects each of his offers. Each consumer can accept at most one offer.

Stage 3'' For each expert e, let $B_e \subset B$ denote the set of consumers who accepted e's offers. Expert e observes consumers' acceptance decisions and chooses the service quality a_b for each consumer $b \in B_e$.[1]

For each consumer $b \in \cup_{e \in E} B_e$, set p_b equal to the offer consumer b accepted, that is, $p_b = p_{e,b}$ for e such that $b \in B_e$. Then, each expert receives utility 4.2. Each consumer $b \in \cup_{e \in E} B_e$ receives utility 4.1, and all other consumers receive the outside option $\underline{v}$.

[1]Whether experts observe only the acceptance decisions on their own offers or on all experts' offers does not matter for our results. For concreteness we assume here that experts observe all acceptance decisions of all consumers.

Note that in contrast to the consumer-proposing mechanism from the main text, consumers receive offers from all experts instead of being matched to only one expert each. Our results are robust to adding a matching stage where consumers are matched to only a few, but at least two, experts whom they receive offers from. The minimum number of two experts per consumer is necessary to initiate price competition.

The second noteworthy assumption is that consumers do not observe the offers received by other consumers. This seems appropriate in the context of service provision, where sellers interact directly, and often privately, with each buyer to deliver the service. The assumption is not relevant for our first result on the existence of an equilibrium that replicates the outcome of the consumer-proposing mechanism from the main text. The structure of other equilibria however may change when making a different informational assumption.

C.2.1 Competitive Outcome

Stages 1” to 3” describe a sequential game of (complete, but) imperfect information. We study its perfect Bayesian equilibria (PBE) in the following. We start by constructing a PBE that replicates the competitive outcome of the consumer-proposing mechanism from Proposition 6.

Proposition 16. *Consider the game described by stages 1” to 3”. There exists a PBE in which all consumers accept offers at marginal cost, $p_b = c(0)$ for all $b \in B$, and receive a service of zero quality, $a_b = 0$ for all $b \in B$.*

Proof. We construct a PBE with the desired properties. The PBE consists of the following elements.

- Expert strategies (for all $e \in E$): for any set B_e, expert e’s quality choices on stage 3” are determined by the Kuhn-Tucker conditions (4.7). Moreover, expert e’s price offers on stage 1 are $p_{e,b} = c(0)$ for all $b \in B$.

- Consumer strategies (for all $b \in B$): for any set of offers $\{p_{e,b}\}_{e \in E}$, consumer b accepts the smallest offer if

$$\min_{e \in E} p_{e,b} \leq v(0) - \underline{v} \; . \tag{C.3}$$

  Otherwise, b rejects all offers. If there are multiple smallest offers satisfying equation (C.3), b chooses one of them randomly (the exact distribution of the randomization does not matter).

- Expert beliefs: experts’ beliefs about the history at any of their information sets is consistent with their observations. Since they observe all events, this uniquely identifies experts’ beliefs.

- Consumer beliefs: at any of his information sets, any consumer $b \in B$ believes that all experts $e \in E$ offered $p_{e,b'} = c(0)$ to all other consumers $b' \in B \setminus \{b\}$.

Note first that the proposed beliefs are consistent with equilibrium strategies.

Second, strategies strategies are sequentially rational. To see this, start with experts' quality choices given B_e. Since experts' problem of choosing quality levels to maximize utility is (strictly differentially) concave, the Kuhn-Tucker conditions (4.7) identify the unique solution to this problem. Moreover, given that consumers always accept the lowest price if it does not exceed the threshold $v(0) - \underline{v}$ and given that all other experts make offers at $c(0)$, there is no profitable deviation from the proposed equilibrium offers. Hence, offers $p_{e,b} = c(0)$ for all $b \in B$ are rational for all experts $e \in E$.

Turning to consumers, note that any consumer b's belief together with other consumers' equilibrium strategies implies $y_e = 0$ for all experts $e \in E$ and at any information set of b. Hence, consumers believe to receive zero quality at all of their information sets. So, choosing any of the lowest offers if they are below $v(0) - \underline{v}$ and rejecting all offers otherwise is rational for consumers given their belief. □

The intuition behind Proposition 16 is standard. Consumers accept the lowest prices and experts undercut each other's prices until they hit marginal cost.

In contrast to standard price competition à la Bertrand, however, equilibria with other outcomes exist. Such equilibria are of two types. In the first type, consumers coordinate to buy only from certain sellers but not from others. Suppose for example that all consumers accept the offer of expert 1 as long as it does not exceed a certain threshold level. Expert 1 will then offer the threshold price and all other experts' offers become irrelevant. Consumers may act rationally in this situation because all experts except for expert 1 have zero income and would therefore provide low quality services.

In the second type of equilibrium, consumers coordinate to buy only from those experts who offer a specific price. As soon as some expert deviates from this offer, consumers believe her profits to be zero, because they believe that no other consumer buys from this expert anymore. So, consumers believe that such a deviating expert provides zero quality and may thus indeed shun her rationally.

Both types of equilibria require a high degree of coordination between consumers. For the first type, consumers must believe all other consumers to accept offers only from a certain, arbitrary set of experts. For the second type, they must believe all other consumers to accept only offers at a certain, arbitrary price. We consider such coordination among consumers implausible as a description of many real-world credence goods markets.

To make this reasoning precise, we propose two criteria for equilibrium selection tailored to our environment.The criteria restrict consumers' ability to coordinate. Both of them leave only those equilibria that lead to the competitive outcome described in Proposition 16.

C.2.2 Equilibrium Selection by Insufficient Reason

Any consumer's decision problem is affected by other consumers' actions exclusively via experts' income levels. Beliefs about experts' incomes are hence crucial for sustaining coordination among consumers. In particular, the types of coordination described above require consumers to entertain different beliefs about different experts' incomes at some of their information sets. To curb such coordination we therefore require consumers' strategies to be optimal even under a belief that treats all experts' incomes identically.

A belief that treats all experts' incomes identically is reminiscent of the Principle of Insufficient Reason. Facing a set of events and no particular reason to believe that one of them is more likely than the others, the Principle of Insufficient Reason advises to assign equal probability to all events. Here, from the perspective of a given consumer, differences in experts' incomes can only stem from other consumers' strategies. Since many such strategies are compatible with PBE, a given consumer has little reason to perceive one set of other consumers' strategies as more likely than another. Hence, according to the Principle of Insufficient Reason, he entertains a belief that does not discriminate between experts.[2]

Definition 9. A PBE is robust to insufficient reason if and only if consumer strategies satisfy the following. Take any set of offers $\{p_{e,b}\}_{e\in E}$ for any consumer b. Let $\infty_{(e,b)}$ be an indicator function equal to one if b accepts $p_{e,b}$ and zero otherwise, and let $a^{IC}(y_e)$ denote the solution to the Kuhn-Tucker conditions (4.7) given y_e. Then, consumer b's acceptance decision following the offers $\{p_{e,b}\}_{e\in E}$ must maximize

$$\int_{\mathbb{R}^N}\left[\sum_{e\in E}\infty_{(e,b)}\left(v(a^{IC}(y_e))-p_{e,b}\right)\right]\pi(y_1,y_2,...,y_N)\,d(y_1,y_2,...,y_N)+\left(1-\sum_{e\in E}\infty_{(e,b)}\right)\underline{v} \tag{C.4}$$

for some probability density function ϕ such that the marginal distributions of the y_e are identical for all e, that is,

$$\tilde{\pi}_e=\tilde{\pi}_{e'} \quad \text{for all } e,e'\in E\ ,$$

[2] The Principle of Insufficient Reason is known to fail as a positive theory of choice under uncertainty when individuals face a decision between a risky (with known probabilities) and an uncertain option (with unknown probabilities). See the Ellsberg Paradox (Ellsberg, 1961). Here, there is no way for consumers to escape the uncertainty about other consumers' choices (and hence experts' incomes). So, the critique based on the Ellsberg Paradox does not apply.

where $\tilde{\pi} : y_e \mapsto \mathbb{R}_+$,

$$\tilde{\pi}_e(y_e) := \int_{\mathbb{R}^{N-1}} \pi(y_1, y_2, ..., y_N)\, d(y_1, ..., y_{e-1}, y_{e+1}, ..., y_N) \ ,$$

is the marginal density for y_e.

Robustness to insufficient reason rules out all PBE with consumer strategies that are optimal only under beliefs that discriminate between experts. Since consumer coordination as described above requires such discriminatory beliefs, the robustness criterion excludes all PBE that rely on consumer coordination.

It turns out that only those PBE survive the selection that lead to the competitive outcome of Proposition 16.

Proposition 17. *Consider the game described by stages 1" to 3". In any PBE that is robust to insufficient reason (see Definition 9), all consumers accept offers at marginal cost, $p_b = c(0)$ for all $b \in B$, and receive services of zero quality, $a_b = 0$ for all $b \in B$.*

Proof. **Step 1.** Robustness to insufficient reason imposes a clear structure on consumer strategies. In particular, since the marginal distributions of experts' incomes are identical under π, maximizing (C.4) is equivalent to choosing the least price offer if

$$\min_{e \in E} p_{e,b} \leq \int_{\mathbb{R}} v(a^{IC}(y_e))\tilde{\pi}(y_e)\, dy_e - \underline{v}$$

and rejecting all offers otherwise. Since $a^{IC} \geq 0$,

$$\int_{\mathbb{R}} v(a^{IC}(y_e))\tilde{\pi}(y_e)\, dy_e \geq v(0) \ .$$

So, if the minimal offer is unique and equal to $c(0)$, it is accepted with certainty.

Step 2. Given the consumer strategies from step 1 the standard logic of Bertrand competition implies that we can never have a situation where consumers accept offers strictly greater than $c(0)$. Moreover, suppose some consumer b accepts no offer. Then, some expert e could offer $p_{e,b} = c(0)$ and consumer b would accept. Both e and b would decide for this deviation, because we assumed that all agents decide against their outside option in case of indifference. So, the only PBE that are robust to insufficient reason have all consumers accept offers at marginal cost $c(0)$.

Step 3. Finally by step 2, we have $y_e = 0$ for all $e \in E$ while all consumers accept some offer. The Kuhn-Tucker conditions (4.7) then imply $a_b = 0$ for all $b \in B$. This must again hold in any PBE that is robust to insufficient reason. □

C.2.3 Equilibrium Selection by Ambiguity Aversion

A critique of robustness to insufficient reason is that consumer strategies must be optimal only under a specific belief π. If consumers cannot coordinate and there are many different equilibrium strategies for consumers, where should such a specific belief come from?

Our second criterion allows consumers to entertain many beliefs and perceive experts' incomes as ambiguous, or uncertain in the Knightian sense. If we additionally assume that consumers are ambiguity averse in the sense of Gilboa and Schmeidler (1989), we obtain the following robustness criterion.

Definition 10. A PBE is robust to strategic ambiguity if and only if consumer strategies satisfy the following. Take any set of offers $\{p_{e,b}\}_{e\in E}$ for any consumer b. Let $\infty_{(e,b)}$ be an indicator function equal to one if b accepts $p_{e,b}$ and zero otherwise, and let $a^{IC}(y_e)$ denote the solution to the Kuhn-Tucker conditions (4.7) given y_e. Then, consumer b's acceptance decision following the offers $\{p_{e,b}\}_{e\in E}$ must maximize

$$\min_{(y_1,y_2,\ldots,y_N)\in\mathbb{R}^N} \sum_{e\in E} \infty_{(e,b)} \left(v(a^{IC}(y_e)) - p_{e,b}\right) + \left(1 - \sum_{e\in E} \infty_{(e,b)}\right) \underline{v} \, . \tag{C.5}$$

In a PBE that is robust to strategic ambiguity, consumer strategies are supported by two considerations. First, as is usual in a PBE, consumers can anticipate other agents' strategies, form beliefs about unobserved events accordingly, and choose their strategies as a best response to the anticipated behavior of others. Second, consumers may perceive the behavior of others as ambiguous and choose the strategies that optimize the worst-case outcome.[3]

The only PBE that are robust to strategic ambiguity are those leading to the competitive outcome of Proposition 16.

Proposition 18. *Consider the game described by stages 1" to 3". In any PBE that is robust to strategic ambiguity (see Definition 10), all consumers accept offers at marginal cost,* $p_b = c(0)$ *for all* $b \in B$*, and receive services of zero quality,* $a_b = 0$ *for all* $b \in B$*.*

Proof. In analogy to the proof of Proposition 17, robustness to strategic ambiguity has clear implications for consumer strategies. In particular, the worst-case outcome for consumers for any acceptance decision they make is when $y_e \leq 0$ for all $e \in E$. So, maximizing (C.5) is equivalent to

[3]Moreover, the combination of the usual PBE requirements with robustness to strategic ambiguity allows consumers to engage in considerations of the following type in equilibrium. Any given consumer anticipates that all other consumers perceive others' behavior as ambiguous and optimize their worst-case outcomes. The given consumer then chooses his strategy as a best response to this anticipated behavior of others.

maximizing

$$\sum_{e\in E} \infty_{(e,b)} \left(v(0) - p_{e,b}\right) + \left(1 - \sum_{e\in E} \infty_{(e,b)}\right) \underline{v}\ .$$

This expression is maximized by accepting the least price offer if

$$\min_{e\in E} p_{e,b} \leq v(0) - \underline{v}$$

and rejecting all offers otherwise. This is essentially the same result as obtained from step 1 in the proof of Proposition 17. The remainder of the proof is then analogous to steps 2 and 3 of the proof of Proposition 17. □

C.2.4 Special Case: Two Experts

As a final remark, for $N = 2$ experts the selection criteria can be relaxed substantially. In particular, with two experts it is sufficient to restrict the off-equilibrium part of consumers' strategies. For expositional reasons we focus on robustness to strategic ambiguity here.

Definition 11. A PBE is weakly robust to strategic ambiguity if and only if any consumer b's actions following any off-equilibrium set of offers $\{p_{e,b}\}_{e\in E}$ satisfy the requirements of robustness to strategic ambiguity described in Definition 10.

The reduction to off-equilibrium actions is substantial. The weakened criterion allows consumers to believe in coordination on any arbitrary set of strategies. Only once they observe an event that is incompatible with the strategies they believed in, consumers revert to ambiguity-averse behavior without committing to any specific new belief about other agents' actions.

For two experts, the weak robustness criterion is sufficient to exclude all outcomes except for the competitive one.

Proposition 19. *Consider the game described by stages 1" to 3" and suppose that $N = 2$. Then in any PBE that is weakly robust to strategic ambiguity (see Definition 10) and has experts play pure strategies, all consumers accept offers at marginal cost, $p_b = c(0)$ for all $b \in B$, and receive services of zero quality, $a_b = 0$ for all $b \in B$.*

Proof. Note first that all consumers under all circumstances prefer to accept an offer smaller or equal to $v(0) - \underline{v}$ to rejecting all offers.

Suppose now that in some PBE as described in the proposition, some consumer b accepts no offer. Then in such a PBE, all offers for consumer b must be strictly above $v(0) - \underline{v}$. But then, expert 1 could deviate to offer $p_{1,b} = v(0) - \underline{v}$. This deviation makes consumer b optimize his worst-case outcome according to weak robustness to strategic ambiguity. Thus, b accepts the least price offer if it does not strictly exceed $v(0) - \underline{v}$. Hence, b accepts $p_{1,b}$. But since $p_{1,b} = v(0) - \underline{v} \geq c(0)$, expert 1 is better off through her initial deviation. So there cannot be a PBE as described in the proposition where some consumer rejects all offers.

Next suppose that in some PBE as described in the proposition, some consumer b accepts an offer $p_{2,b} > c(0)$. Then, expert 1 can deviate to some offer $p_{1,b}$ such that $p_{1,b} < p_{2,b}$ and $p_{1,b} \in [c(0), v(0) - \underline{v}]$. The deviation again makes b optimize his worst-case outcome, so b accepts $p_{1,b}$. This makes expert 1 better off, so the deviation is profitable for expert 1. Thus, there cannot be a PBE as described in the proposition where some consumer accepts an offer above marginal cost.

Hence we have shown that in any PBE as described in the proposition, all consumers accept offers at marginal cost $c(0)$. This immediately implies $y_e = 0$ for all experts and, by conditions (4.7), $a_b = 0$ for all consumers. □

Bibliography

Abeliansky, Ana and Klaus Prettner (2017) "Automation and Demographic Change," May, Working Paper.

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.

——— (2003) "Patterns of Skill Premia," *Review of Economic Studies*, 70 (2), 199–230, 10.1111/1467-937X.00242.

——— (2007) "Equilibrium Bias of Technology," *Econometrica*, 75 (5), 1371–1409.

——— (2009) *Introduction to modern economic growth*: MIT Press.

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

——— (2018b) "Demographics and Automation," March, Working Paper.

——— (2018c) "Modeling Automation," *AEA Papers and Proceedings*, 108, 48–53.

——— (2019) "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, forthcoming.

Aghion, Philippe and Peter Howitt (1992) "A Model of Growth Through Creative Destruction," *Econometrica*, 60 (2), 323–351.

Aghion, Phillippe, Benjamin F. Jones, and Charles I. Jones (2017) "Artificial Intelligence and Economic Growth," October, Working Paper.

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

Alvaredo, Facundo, Lucas Chancel, Thomas Piketty, Emmanuel Saez, and Gabriel Zucman eds. (2018) *The World Inequality Report: 2018*: Harvard University Press.

Andreoni, James, Nikos Nikiforakis, and Jan Stoop (2017) "Are the Rich More Selfish than the Poor, or Do They Just Have More Money? A Natural Field Experiment," March, NBER Working Paper No. 23229.

Arnott, Richard, Bruce Greenwald, and Joseph E. Stiglitz (1994) "Information and Economic Efficiency," *Information Economics and Policy*, 6 (1), 77–82.

Arntz, Melanie, Terry Gregory, and Ulrich Zierahn (2016) "The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis," Social, Employment and Migration Working Paper 189, OECD.

Atkeson, Andrew, Christian Hellwig, and Guillermo Ordonez (2015) "Optimal Regulation in the Presence of Reputation Concerns," *Quarterly Journal of Economics*, 130 (1), 415–464.

Atkinson, Anthony B. (1970) "On the Measurement of Inequality," *Journal of Economic Theory*, 2, 244–263.

Autor, David H. and David Dorn (2013) "The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market," *American Economic Review*, 103 (5), 1553–1597.

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney (2006) "The Polarization of the U.S. Labor Market," *American Economic Review*, 96 (2), 189–94.

Backer, Koen De, Carlo Menon, Isabelle Desnoyers-James, and Laurent Moussiegt (2016) "Reshoring: Myth or Reality?" Science, Technology and Industry Policy Paper 27, OECD.

Bartling, Björn, Vanessa Valero, and Roberto A. Weber (2019) "The Causal Effect of Income Growth on Consumer Social Responsibility," Working Paper.

Bester, Helmut and Matthias Dahm (2018) "Credence Goods, Costly Diagnosis and Subjective Evaluation," *Economic Journal*, 128 (611), 1367–1394.

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 (2018) "The UK Wage Premium Puzzle: How Did a Large Increase in University Graduates Leave the Education Premium Unchanged?", Working Paper.

Bound, John and George Johnson (1992) "Changes in the Structure of Wages in the 1980's: An Evaluation of Alternative Explanations," *American Economic Review*, 82 (3), 371–392.

Brewer, Mike, Emmanuel Saez, and Andrew Shepard (2010) "Means-testing and Tax Rates on Earnings," in Adam, Stuart, 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*, 90–201: Oxford University Press.

Caballero, Ricardo J., Emmanuel Farhi, and Pierre-Olivier Gourinchas (2017) "Rents, Technical Change, and Risk Premia – Accounting for Secular Trens in Interest Rates, Returns on Capital, Earning Yields, and Factor Shares," *American Economic Review: Papers & Proceedings*, 107 (5), 614–620.

Carneiro, Pedro, Kai Liu, and Kjell G. Salvanes (2019) "The Supply of Skill and Endogenous Technical Change: Evidence From a College Expansion Reform," June, Working Paper.

Carpenter, Jeffrey, Eric Verhoogen, and Stephen Burks (2005) "The Effect of Stakes in Distribution Experiments," *Economics Letters*, 86 (3), 393–398.

Chetty, Raj (2012) "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply," *Econometrica*, 80 (3), 96–1018.

Chowdhury, Subhasish M and Joo Young Jeon (2014) "Impure altruism or inequality aversion?: An experimental investigation based on income effects," *Journal of Public Economics*, 118, 143–150.

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, Dave Donaldson, Jonathan Vogel, and Iván Werning (2015) "Comparative Advantage and Optimal Trade Policy," *Quarterly Journal of Economics*, 130 (1), 77–128.

Costinot, Arnaud and Jonathan Vogel (2010) "Matching and Inequality in the World Economy," *Journal of Political Economy*, 118 (4), 747–786.

——— (2015) "Beyond Ricardo: Assignment Models in International Trade," *Annual Review of Economics*, 7, 31–62.

Costinot, Arnaud and Ivan Werning (2018) "Robots, Trade, and Luddism: A Sufficient Statistic Approach to Optimal Technology Regulation," September, Working Paper.

Darby, Michael R. and Edi Karni (1973) "Free Competition and the Optimal Amount of Fraud," *Journal of Law and Economics*, 16 (1), 67–88.

Dauth, Wolfgang, Sebastian Findeisen, Jens Suedekum, and Nicole Woessner (2017) "German Robots – The Impact of Industrial Robots on Workers," Discussion Paper 12306, CEPR.

Dechezlepretre, Antoine, David Hemous, Morten Olsen, and Carlo Zanella (2019) "Automating Labor: Evidence from Firm-level Patent Data," October, 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.

Dixit, Avinash, Gene Grossman, and Elhanan Helpman (1997) "Common Agency and Coordination: General Theory and Application to Government Policy Making," *Journal of Political Economy*, 105 (4), 752–769.

Djankov, Simeon, Rafael La Porta, Florencio Lopez de Silanes, and Andrei Shleifer (2002) "The Regulation of Entry," *Quarterly Journal of Economics*, 117 (1), 1–37.

Dulleck, Uwe and Rudolf Kerschbamer (2006) "On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods," *Journal of Economic Literature*, 44 (1), 5–42.

Dulleck, Uwe, Rudolf Kerschbamer, and Matthias Sutter (2011) "The Economics of Credence Goods: An Experiment on the Role of Liability, Verifiability, Reputation, and Competition," *American Economic Review*, 101 (2), 526–55.

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.

Ellsberg, Daniel (1961) "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics*, 75 (4), 643–669.

Emons, Winand (1997) "Credence goods and fraudulent experts," *RAND Journal of Economics*, 107–119.

Engel, Christoph (2011) "Dictator games: A meta study," *Experimental Economics*, 14 (4), 583–610.

Englmaier, Florian and Achim Wambach (2010) "Optimal incentive contracts under inequity aversion," *Games and Economic Behavior*, 69 (2), 312–328.

European Economic and Social Committee (2014) "The State of Liberal Professions Concerning their Functions and Relevance to European Civil Society," European Union.

Fadinger, Harald and Karin Mayr (2014) "Skill-Biased Technological Change, Unemployment and Brain Drain," *Journal of the European Economic Association*, 12 (2), 397–431.

Feldstein, Martin S. (1969) "The Effects of Taxation on Risk Taking," *Journal of Political Economy*, 77 (5), 755–764.

Feng, Andy and Georg Graetz (2018) "Training Requirements, Automation, and Job Polarization," July, Working Paper.

Frey, Carl B. and Michael A. Osborne (2017) "The future of employment: How susceptible are jobs to computerisation?" *Technological Forecasting and Social Change*, 114 (C), 254–280.

Goldin, Claudia and Lawrence F. Katz (2008) "The Race between Education and Technology: The Evolution of U.S. Educational Wage Differentials, 1890 to 2005," Working Paper 12984, NBER.

Graetz, Georg and Guy Michaels (2018) "Robots at Work," *Review of Economics and Statistics*, 100 (5), 753–768.

Gruber, Jon and Emmanuel Saez (2002) "The Elasticity of Taxable Income: Evidence and Implications," *Journal of Public Economics*, 84 (1), 1–32.

Guerreiro, Joao, Sergio Rebelo, and Pedro Teles (2018) "Should Robots be Taxed?" Working Paper 23806, NBER.

Heathcote, Jonathan, Kjetil Storesletten, and Gianluca Violante (2017) "Optimal Tax Progressivity: An Analytical Framework," *Quarterly Journal of Economics*, 132 (4), 1693–1754.

Hémous, David and Morten Olsen (2018) "The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality," March, Working Paper.

Jacobs, Bas and Uwe Thuemmel (2018a) "Optimal Linear Income Taxation and Education Subsidies under Skill-Biased Technical Change," September, Working Paper.

——— (2018b) "Optimal Taxation of Income and Human Capital and Skill-Biased Technical Change," September, Working Paper.

Jagadeesan, Ravi (2019) "Optimal Taxation with an Endogenous Growth Rate," February, Working Paper.

Jones, Charles I. (2019) "Taxing Top Incomes in a World of Ideas," July, Working Paper.

Karabarbounis, Loukas and Brent Neiman (2014) "The Global Deline of the Labor Share," *Quarterly Journal of Economics*, 129 (1), 61–103.

Katz, Lawrence F. and Kevin M. Murphy (1992) "Changes in Relative Wages, 1963-87: Supply and Demand Factors," *Quarterly Journal of Economics*, 107 (1), 35–78.

Kerschbamer, Rudolf, Matthias Sutter, and Uwe Dulleck (2017) "How Social Preferences Shape Incentives in (Experimental) Markets for Credence Goods," *Economic Journal*, 127 (600), 393–416.

Kiley, Michael T. (1999) "The Supply of Skilled Labour and Skill-Biased Technological Progress," *Economic Journal*, 109 (458), 708–724.

Kleiner, Morris M. and Alan B. Krueger (2013) "Analyzing the Extent and Influence of Occupational Licensing on the Labor Market," *Journal of Labor Economics*, 31 (2), 173–202.

Korenok, Oleg, Edward L Millner, and Laura Razzolini (2012) "Are Dictators Averse to Inequality?" *Journal of Economic Behavior & Organization*, 82 (2-3), 543–547.

Kőszegi, Botond (2014) "Behavioral Contract Theory," *Journal of Economic Literature*, 52 (4), 1075–1118.

Krenz, Astrid, Klaus Prettner, and Holger Strulik (2018) "Robots, Reshoring, and the Lot of Low-Skill Workers," Discussion Paper 351, CEGE.

Lewis, Ethan (2011) "Immigration, Skill Mix, and Capital Skill Complementarity," *Quarterly Journal of Economics*, 126 (2), 1029–1069.

List, John A. (2011) "The Market for Charitable Giving," *Journal of Economic Perspectives*, 25 (2), 157–180.

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," November, Working Paper.

——— (2018) "An Elementary Theory of Directed Technical Change and Wage Inequality," August, Working Paper.

Lordan, Grace and David Neumark (2018) "People Versus Machines: The Impact of Minimum Wages on Automatable Jobs," Working Paper 23667, NBER.

Lucas, Robert E. (1988) "On the mechanics of economic development," *Journal of Monetary Economics*, 22 (1), 3–42.

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 Adam, Stuart, 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*, 202–274: Oxford University Press.

Milgrom, Paul and John Roberts (1996) "The LeChatelier Principle," *American Economic Review*, 86 (1), 173–179.

Milgrom, Paul and Ilya Segal (2002) "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70 (2), 583–601.

Milgrom, Paul and Chris Shannon (1994) "Monotone Comparative Statics," *Econometrica*, 62 (1), 157–180.

Mimra, Wanda, Alexander Rasch, and Christian Waibel (2016) "Price competition and reputation in credence goods markets: Experimental evidence," *Games and Economic Behavior*, 100, 337–352.

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," Working Paper 24078, NBER.

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.

Nedelkoska, Ljubica and Glenda Quintini (2018) "Automation, skills use and training," Social, Employment and Migration Working Paper 202, OECD.

OECD (2011) "Taxation and Employment," Tax Policy Studies 21, OECD.

Paris, Valérie, Marion Devaux, and Lihan Wei (2010) "Health Systems Institutional Characteristics: A Survey of 29 OECD Countries," April, OECD Health Working Papers No. 50.

Pesendorfer, Wolfgang and Asher Wolinsky (2003) "Second Opinions and Price Competition: Inefficiency in the Market for Expert Advice," *Review of Economic Studies*, 70 (2), 417–437.

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.

Pitchik, Carolyn and Andrew Schotter (1987) "Honesty in a Model of Strategic Information Transmission," *American Economic Review*, 77 (5), 1032–1036.

Prescott, Edward C. and Robert M. Townsend (1984) "Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard," *Econometrica*, 52 (1), 21–46.

Rakotoarisoa, Manitra A., Massimo Iafrate, and Marianna Paschali (2012) *Why has Africa become a net food importer? Explaining Africa agricultural and food trade deficits*: Trade and Markets Division, FAO.

Rasch, Alexander and Christian Waibel (2018) "What Drives Fraud in a Credence Goods Market?– Evidence from a Field Study," *Oxford Bulletin of Economics and Statistics*, 80 (3), 605–624.

Romer, Paul M. (1986) "Increasing Returns and Long-Run Growth," *Journal of Political Economy*, 94 (5), 1002–1037.

——— (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) "Redisributive 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," November, Working Paper.

Sevinc, Orhun (2018) "Skill-Biased Technical Change and Labor Market Polarization: The Role of Skill Heterogeneity Within Occupations," March, Working Paper.

Slavik, Ctirad and Hakki Yazici (2014) "Machines, Buildings, and Optimal Dynamic Taxes," *Journal of Monetary Economics*, 66, 47–61.

Smeets, Paul, Rob Bauer, and Uri Gneezy (2015) "Giving behavior of millionaires," *Proceedings of the National Academy of Sciences*, 112 (34), 10641–10644.

Smith, Vincent H, Michael R Kehoe, and Mary E Cremer (1995) "The private provision of public goods: Altruism and voluntary giving," *Journal of Public Economics*, 58 (1), 107–126.

Stiglitz, Joseph E. (1982) "Self-Selection and Pareto Efficient Taxation," *Journal of Public Economics*, 17 (2), 213–240.

Teulings, Coen N. (1995) "The Wage Distribution in a Model of the Assignment of Skills to Jobs," *Journal of Political Economy*, 103 (2), 280–315.

Thuemmel, Uwe (2018) "Optimal Taxation of Robots," August, Working Paper.

de Vries, Gaaitzen, Marcel Timmer, and Klaas de Vries (2015) "Structural Transformation in Africa: Static Gains, Dynamic Losses," *Journal of Development Studies*, 51 (6), 674–688.

Wiepking, Pamala and René Bekkers (2012) "Who gives? A literature review of predictors of charitable giving. Part Two: Gender, family composition and income," *Voluntary Sector Review*, 3 (2), 217–245.

Wolinsky, Asher (1993) "Competition in a market for informed experts' services," *RAND Journal of Economics*, 380–398.

World Bank (2016) "Digital Dividends," World Development Report 2016, World Bank.