An Elementary Theory of Directed Technical Change and Wage Inequality*

Jonas Loebbing†

August 2018

This paper generalizes central results from the theory of (endogenously) directed technical change to settings where technology does not take a labor-augmenting form and with arbitrarily many levels of skill. Building on simple notions of complementarity, the results remain intuitive despite their generality. The developed theory allows to study the endogenous determination of labor-replacing, that is, automation technology through the lens of directed technical change theory. In an assignment model with a continuum of differentially skilled workers and capital, where capital takes the form of machines that perfectly substitute for labor in the production of tasks, any increase in the relative supply of skilled workers stimulates both automation and investment into improving the productivity of machines, potentially leading skill premia to increase in relative skill supply. Relatedly, trade with a skill-scarce country discourages automation and machine improvements, potentially reversing the standard Heckscher-Ohlin effects.

JEL: J24, J31, O33, **Keywords**: Directed Technical Change, Endogenous Technical Change, Wage Inequality, Automation, Assignment Model, Monotone Comparative Statics.

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

*I thank Peter Funk for helpful comments, discussion, and guidance. Moreover, I thank Daron Acemoglu, Raphael Flore, Tobias Foell, Philipp Giesa, Emanuel Hansen, Anna Kerkhof, Ilse Lindenlaub, Xianjia Ye, and participants at various presentations for helpful comments and discussion. Previous versions circulated under the titles "A LeChatelier Principle for Relative Demand and Implications for Directed Technical Change" and "Relative Equilibrium Bias of Technology".

†Center for Macroeconomic Research, University of Cologne. Website: https://cmr.uni-koeln.de/en/team/phd-students/research-teaching-assistant/loebbing/. Email: loebbing[at]wiso.uni-koeln.de. Telephone: +49 221 470 8650.

[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 (2018a), and Dauth, Findeisen, Suedekum and Woessner (2017) on the effects of automation technology in particular.

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), 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, 2018b; 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.

[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 5.4.

[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.

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, 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,

[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.

which in turn are combined to produce a single final good. In line with recent forecasts on the future automation potential 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, 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 com-

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

petition 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 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 3 presents these results for the case with only two different levels of skill. Section 4 generalizes them to skill supply of arbitrary dimension. Section 5 applies the results to endogenous automation technology in assignment models, and Section 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, 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

[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).

[8] See Acemoglu and Restrepo (2018d) 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.

with differentially labor-augmenting technology. Most closely related to the second part of the paper are recent papers by Acemoglu and Restrepo (2018b), Hémous and Olsen (2018), Feng and Graetz (2018), Acemoglu and Restrepo (2018c), and Krenz, Prettner and Strulik (2018). Acemoglu and Restrepo (2018b) 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 (2018c) 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 5.4, their model does not feature endogenous investment into the productivity of automation technology and hence (by the results of Section 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. 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 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 1 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'

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

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

$$w(L, \theta) = \nabla_L F(L, \theta), \tag{1}$$

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

$$\theta^*(L) \in \underset{\theta \in \Theta}{\operatorname{argmax}} F(L, \theta). \tag{2}$$

Moreover, wages are given by $\nabla_L F(L, \theta^*(L))$, so equation (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 B.1 shows that equations (1) and (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 (1) and (2) in symmetric exogenous or, respectively, endogenous technology equilibria.[10]

Appendix B.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

[10] Acemoglu (2007) presents three more models (his economies C, M, and O) of endogenous technical change that satisfy equations (1) and (2). His models are related to but distinct from those presented in Appendix B.1.

becomes relevant in Section 3.2 below.[11]

Since equations (1) and (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 1. The goal is to answer the following questions.

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 (1) and (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 (1). Second, the equilibrium technology $\theta^*(L)$ responds according to equation (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

[11] While all specific models presented in this paper are static, the models in Appendix B.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).

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). 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 $\text{argmax}_\theta F(L,\theta)$, 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 B.1 or in the applied Section 5), weak conditions guarantee that $\text{argmax}_\theta F(L,\theta)$ is a singleton, so the selection of a unique $\theta^*(L)$ does not seem very restrictive.

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.

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' 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

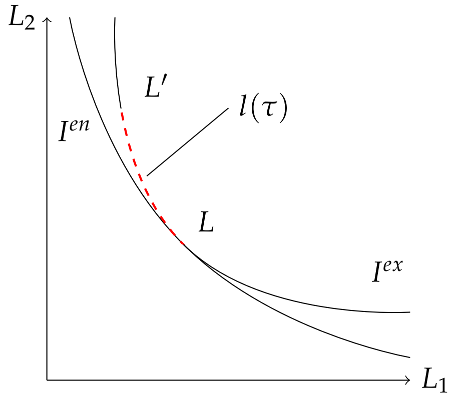

Figure 1. I^{ex} and I^{en} are the exogenous and the endogenous technology isoquants through L. The move from L to L' is parameterized by $l(\tau)$ such that $l(0) = L$ and $l(1) = L'$. Moving along $l(\tau)$ leaves output constant when $\theta^*(L)$ is fixed, but must raise output when technology is 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.

$\tau \in [0,1]$. 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 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

$$F(l(0), \theta^*(L')) \leq F(l(0), \theta^*(L)) = F(l(1), \theta^*(L)) \leq F(l(1), \theta^*(L')). \tag{3}$$

In Figure 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 (3) must be strict, because θ^* is selected as the supremum of the maximizer set in equation (2). (If both inequalities were equalities, we would either select $\theta^*(L)$ 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

$$\begin{aligned} 0 &< \nabla_L F(l(\tau'), \theta^*(L')) \frac{\mathrm{d}\, l(\tau')}{\mathrm{d}\, \tau} \\ &= w_1(l(\tau'), \theta^*(L')) \frac{\mathrm{d}\, l_1(\tau')}{\mathrm{d}\, \tau} + w_2(l(\tau'), \theta^*(L')) \frac{\mathrm{d}\, l_2(\tau')}{\mathrm{d}\, \tau}. \end{aligned} \tag{4}$$

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

$$w_1(l(\tau'), \theta^*(L)) \frac{\mathrm{d}\, l_1(\tau')}{\mathrm{d}\, \tau} + w_2(l(\tau'), \theta^*(L)) \frac{\mathrm{d}\, l_2(\tau')}{\mathrm{d}\, \tau} = 0. \tag{5}$$

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

Finally, rearranging and combining equations (4) and (5) yields

$$\frac{w_2(l(\tau'),\theta^*(L'))}{w_1(l(\tau'),\theta^*(L'))} > \frac{w_2(l(\tau'),\theta^*(L))}{w_1(l(\tau'),\theta^*(L))}. \tag{6}$$

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 B.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 $\text{argmax}_\theta\, F(L,\theta)$ 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 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 5.

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. 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 shows two alternative shapes of the endogenous technology isoquant through L. In one case

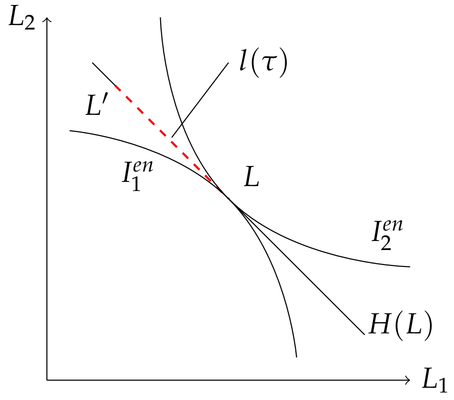

Figure 2. The figure shows two alternative endogenous technology isoquants, I_1^{en} for the case of a quasiconvex endogenous technology production function $\overline{F}$ and I_2^{en} for the quasiconcave case. In the quasiconcave case, moving from L to L' along the line $H(L)$ reduces $\overline{F}$ 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.

(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, 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 B.1 and B.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 B.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 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.

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.

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. □

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 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. □

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 5.3 demonstrates its usefulness in a case where Theorem 1 would not be applicable.
Finally, Appendix B.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).

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 $w(L', \theta^*(L')) \not\preceq^p w(L, \theta^*(L))$.

(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. □

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 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 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 3.2 for detailed discussion).

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

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

it holds that

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

Proof. See Appendix A.1. □

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, 2018a; 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

[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).

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 as international trade. Section 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 3.1 to 4.2 apply.

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, 2018b), 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, applying the results developed in the previous part of the paper.

[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, 2018).

[15]See Acemoglu and Restrepo (2018d) 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 (2018a); Dauth et al. (2017); Abeliansky and Prettner (2017); Acemoglu and Restrepo (2018c), 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).

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

$$Y = \beta \left(\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \, \mathrm{d}\, x \right)^{\frac{\epsilon}{\epsilon-1}} \tag{7}$$

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.

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 $x < x'$.

[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.

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, 2018a).[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 $\tilde{x}$ such that capital performs all tasks below $\tilde{x}$ while labor performs all tasks above. Moreover, more 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 $\tilde{x} \in X$ and a strictly increasing and continuous matching function $m : S \to [\tilde{x}, \overline{x}]$ such that*

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

Proof. See Appendix A.1. □

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

[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 (2018a) 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).

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 (7);	(final good firms)
(E3)	$m(s) \in \operatorname{argmax}_{x \in X} \gamma(s, x) p_x$ for all s;	(task producers)
(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)	$\frac{w_{\underline{s}}}{\gamma(\underline{s}, \widetilde{x})} = \frac{r}{\alpha(\widetilde{x})}$.	

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]

$$\frac{\mathrm{d} \log w_s}{\mathrm{d}\, s} = \frac{\partial \log \gamma(s, m(s))}{\partial s}. \tag{8}$$

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.

[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.

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. 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.

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. □

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,*

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

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 $\tilde{x}$ corresponds to skill-biased technical change in the current model, $\tilde{x} \leq \tilde{x}' \Leftrightarrow \tilde{x} \preceq^b \tilde{x}'$ (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 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 $\tilde{x}$) 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, $\overline{F}(L) := F(L, \tilde{x}^*(L))$.
Here, firms make their automation decisions individually and independently of each other, as in the baseline model of endogenous technology choices in Section 2. 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, $\overline{F}(L) := F(L, \tilde{x}^*(L))$ (where F is as defined in Lemma 2), is quasiconcave.*

Proof. See Appendix A.1. □

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

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

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

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

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 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.

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]

[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.

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 B.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

$$Y = \int_0^1 \beta_i q_{\beta,i}^{\kappa} \,\mathrm{d}\, i \left(\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \,\mathrm{d}\, x \right)^{\frac{(1-\kappa)\epsilon}{\epsilon-1}}, \tag{9}$$

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

$$p_{\alpha,i} = \kappa \alpha_i q_{\alpha,i}^{\kappa-1} \int_0^{\widetilde{x}} p_x K_x^{1-\kappa} \,\mathrm{d}\, x, \tag{10}$$

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

$$p_{\beta,i} = \kappa \beta_i q_{\beta,i}^{\kappa-1} \left(\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \,\mathrm{d}\, x \right)^{\frac{(1-\kappa)\epsilon}{\epsilon-1}}. \tag{11}$$

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

[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.

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

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

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:

$$\rho \kappa q_{\alpha,i}^\kappa \int_{\underline{x}}^{\tilde{x}} p_x K_x^{1-\kappa} \, \mathrm{d}\, x = p_D \alpha_i^{\frac{1-\rho}{\rho}}. \tag{12}$$

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

$$\rho \kappa q_{\beta,i}^\kappa \left(\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \, \mathrm{d}\, x \right)^{\frac{(1-\kappa)\epsilon}{\epsilon-1}} = p_D \beta_i^{\frac{1-\rho}{\rho}}. \tag{13}$$

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

- an automation threshold $\tilde{x}$, a matching function $m : S \to [\tilde{x}, \overline{x}]$, 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 $x < \tilde{x}$ and $Y_x = \gamma(m^{-1}(x), x) L_{m^{-1}(x)} \frac{\mathrm{d}\, m^{-1}(x)}{\mathrm{d}\, x}$ if $x \geq \tilde{x}$;	(market clearing)
(E2)′	$p_x = \frac{\partial Y}{\partial Y_x}$ for all x, where Y is given by (9);	(final good firms)
(E3)′	q_β satisfies equation (11);	
(E4)′	$m(s) \in \operatorname{argmax}_{x \in X} \gamma(s, x) p_x$ for all s;	(task producers)
(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 $x < \tilde{x}$ and $p_c = r$;	
(E7)′	$\frac{w_{\underline{s}}}{\gamma(\underline{s}, \tilde{x})} = \left(\frac{p_\alpha}{\kappa \alpha}\right)^\kappa \left(\frac{r}{(1-\kappa)\alpha}\right)^{1-\kappa}$;	
(E8)′	q_α satisfies equation (10);	
(E9)′	$p_\alpha = \frac{\eta_\alpha}{\kappa}$ and $p_\beta = \frac{\eta_\beta}{\kappa}$;	(technology firms)
(E10)′	α, β, and p_D satisfy equations (12), (13), and $\alpha^{\frac{1}{\rho}} + \beta^{\frac{1}{\rho}} = D$.	

[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 (2018b).

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 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

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

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. 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. □

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 B.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 $[\underline{s}, \tilde{s}]$ who have no discernible comparative advantage over machines. Formally, for $s \in [\underline{s}, \tilde{s}]$, $\gamma(s, x)$ 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 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

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

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. □

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)′,

$$w_{\underline{s}} = \left(\frac{p_\alpha}{\kappa\alpha}\right)^{\kappa} \left(\frac{r}{(1-\kappa)\alpha}\right)^{1-\kappa} \gamma(\underline{s}, \widetilde{x}), \tag{14}$$

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 $s \le \widetilde{s}$ by noting that the ratios $w_s / w_{\underline{s}}$ are fixed for all $s \le \widetilde{s}$, 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 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 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.

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.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

[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.

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.2. □

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, 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 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.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 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]

[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

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 (2018a).[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, 2018a, 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 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

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.

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

[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 shortage of alternative employment possibilities emphasized by the theory.

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.

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.

References

Abeliansky, Ana and Klaus Prettner (2017) "Automation and Demographic Change," 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.

——— (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) "Robots and Jobs: Evidence from US Labor Markets," Working Paper.

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

——— (2018c) "Demographics and Automation," Working Paper.

——— (2018d) "Modeling Automation," *AEA Papers and Proceedings*, **108**, 48–53.

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," Working Paper.

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.

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.

Blum, Bernardo S. (2010) "Endowments, Output, and the Bias of Directed Innovation," *Review of Economic Studies*, **77** (2), 534–559.

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.

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 (2018) "The Supply of Skill and Endogenous Technical Change: Evidence From a College Expansion Reform," Working Paper.

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.

Dauth, Wolfgang, Sebastian Findeisen, Jens Suedekum, and Nicole Woessner (2017) "German Robots – The Impact of Industrial Robots on Workers," Discussion Paper 12306, CEPR.

Feng, Andy and Georg Graetz (2018) "Training Requirements, Automation, and Job Polarization," 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*, forthcoming.

Hémous, David and Morten Olsen (2018) "The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality," 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.

Kiley, Michael T. (1999) "The Supply of Skilled Labour and Skill-Biased Technological Progress," *Economic Journal*, **109** (458), 708–724.

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.

Loebbing, Jonas (2016) "A LeChatelier Principle for Relative Demand and Implications for Directed Technical Change," 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.

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.

Morrow, Peter M. and Daniel Trefler (2017) "Endowments, Skill-Biased Technology, and Factor Prices: A Unified Approach to Trade," Working Paper 24078, NBER.

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.

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.

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.

Sevinc, Orhun (2018) "Skill-Biased Technical Change and Labor Market Polarization: The Role of Skill Heterogeneity Within Occupations," Working Paper.

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.

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.

World Bank (2016) "Digital Dividends," World Development Report 2016, World Bank.

A. Omitted Proofs and Derivations

A.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 $\mathrm{d}\,L$:

$$\nabla_L\left[\nabla_\theta F(L,\theta^*(L))\,\mathrm{d}\,\theta^*\right]\mathrm{d}\,L \geq 0. \tag{15}$$

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 (15) above. $\square$

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:

$$\overline{F}(l(\overline{\tau})) - \overline{F}(l(0)) = \int_0^{\overline{\tau}} \nabla_L F(l(\tau), \theta^*(l(\tau))) \frac{\mathrm{d}\, l(\tau)}{\mathrm{d}\,\tau}\, \mathrm{d}\,\tau > 0$$
$$\overline{F}(l(1)) - \overline{F}(l(\overline{\tau})) = \int_{\overline{\tau}}^1 \nabla_L F(l(\tau), \theta^*(l(\tau))) \frac{\mathrm{d}\, l(\tau)}{\mathrm{d}\,\tau}\, \mathrm{d}\,\tau < 0.$$

It follows that there exist $\tau_1 < \tau_2$ such that

$$\nabla_L F(l(\tau_2), \theta^*(l(\tau_2))) \frac{\mathrm{d}\, l(\tau_2)}{\mathrm{d}\,\tau} < 0 < \nabla_L F(l(\tau_1), \theta^*(l(\tau_1))) \frac{\mathrm{d}\, l(\tau_1)}{\mathrm{d}\,\tau}.$$

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 $\theta^*(L)\not\sim^b\theta^*(L')$, because otherwise the statement of the theorem is trivially satisfied. In these cases, we have

$$F(l(0),\theta^*(L'))\leq F(l(0),\theta^*(L))=F(l(1),\theta^*(L))\leq F(l(1),\theta^*(L')),\tag{16}$$

with at least one of the inequalities being strict because we select the supremum of the maximizer set in equation (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 (16) implies that there is a $\tau'\in(0,1)$ such that

$$\nabla_L F(l(\tau'),\theta^*(L'))\frac{\mathrm{d}\,l(\tau')}{\mathrm{d}\,\tau}>0.\tag{17}$$

Let $\tilde{s}$ denote a skill level such that $\mathrm{d}\,l_s(\tau')/\,\mathrm{d}\,\tau\geq 0$ for all $s>\tilde{s}$ and $\mathrm{d}\,l_s(\tau')/\,\mathrm{d}\,\tau\leq 0$ for all $s\leq\tilde{s}$. 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 (17) to

$$\left[\nabla_L F(l(\tau'),\theta^*(L'))-\frac{w_{\tilde{s}}(l(\tau'),\theta^*(L'))}{w_{\tilde{s}}(l(\tau'),\theta^*(L))}\nabla_L F(l(\tau'),\theta^*(L))\right]\frac{\mathrm{d}\,l(\tau')}{\mathrm{d}\,\tau}>0.\tag{18}$$

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 $\tilde{s}$. 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 $\tilde{s}$, since

$$\frac{w_s(l(\tau'),\theta^*(L')}{w_s(l(\tau'),\theta^*(L))} \gtrless \frac{w_{\tilde{s}}(l(\tau'),\theta^*(L')}{w_{\tilde{s}}(l(\tau'),\theta^*(L))} \quad \text{if} \quad s \lesseqgtr \tilde{s}. \tag{19}$$

But this implies that the product of the two vectors is weakly negative, in contradiction to inequality (18). 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 $\tilde{\tau}$ such that

$$F(l(\tilde{\tau}),\theta^*(L')) = F(l(\tilde{\tau}),\theta^*(L)) = F(l(\tau),\theta^*(L)) < F(l(\tau),\theta^*(L'))$$

for all $\tau \in (\tilde{\tau},1]$.
The mean value theorem then implies existence of a $\tau' \in (\tilde{\tau},1)$ such that

$$\nabla_L F(l(\tau'),\theta^*(L'))\frac{\mathrm{d}\,l(\tau')}{\mathrm{d}\,\tau} > 0, \tag{20}$$

analogous to inequality (17) above. From here on, the proof follows exactly the proof of Theorem 1, starting at inequality (17). Note that inequality (19) holds here, because the choice of $\tilde{\tau}$ 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 (19) 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 $\tilde{s} \in (0,1)$ be the skill such that $\mathrm{d}\,l_s(1)/\,\mathrm{d}\,\tau > 0$ for all $s > \tilde{s}$ and $\mathrm{d}\,l(1)/\,\mathrm{d}\,\tau \leq 0$ for $s \leq \tilde{s}$. Such an $\tilde{s}$ 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 $\tilde{s}' \in (0,1)$

such that $\mathrm{d}\, l_s(1)/\,\mathrm{d}\,\tau < 0$ if $s < \tilde{s}'$. Moreover, by hypothesis, we have

$$\frac{w_s(l(1),\theta^*(l(1)))}{w_s(l(0),\theta^*(l(0)))} \gtreqless \frac{w_{\tilde{s}}(l(1),\theta^*(l(1)))}{w_{\tilde{s}}(l(0),\theta^*(l(0)))} \quad \text{if} \quad s \gtreqless \tilde{s},$$

with strict inequality for a strictly positive measure of skills.[30]

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

$$\left[w(l(1),\theta^*(l(1))) - \frac{w_{\tilde{s}}(l(1),\theta^*(l(1)))}{w_{\tilde{s}}(l(0),\theta^*(l(0)))} w(l(0),\theta^*(l(0)))\right] \frac{\mathrm{d}\, l(1)}{\mathrm{d}\,\tau} > 0,$$

because the left-hand side is the inner product of two vectors with positive (negative) entries for s above (below) $\tilde{s}$, and these vectors are simultaneously different from zero at a subset of skills of strictly positive measure. It follows that

$$w(l(1),\theta^*(l(1))) \frac{\mathrm{d}\, l(1)}{\mathrm{d}\,\tau} > \frac{w_{\tilde{s}}(l(1),\theta^*(l(1)))}{w_{\tilde{s}}(l(0),\theta^*(l(0)))} w(l(0),\theta^*(l(0))) \frac{\mathrm{d}\, l(1)}{\mathrm{d}\,\tau} = 0, \tag{21}$$

where the equality follows from the construction of $l(\tau)$. Using that wages are identical to the L-derivative of F, inequality (21) 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:

$$\overline{F}(l(\overline{\tau})) - \overline{F}(l(0)) = \int_0^{\overline{\tau}} \nabla_L F(l(\tau),\theta^*(l(\tau))) \frac{\mathrm{d}\, l(\tau)}{\mathrm{d}\,\tau}\, \mathrm{d}\,\tau < 0$$
$$\overline{F}(l(1)) - \overline{F}(l(\overline{\tau})) = \int_{\overline{\tau}}^1 \nabla_L F(l(\tau),\theta^*(l(\tau))) \frac{\mathrm{d}\, l(\tau)}{\mathrm{d}\,\tau}\, \mathrm{d}\,\tau > 0.$$

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.

[30]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.

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 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

$$\left[w(l(\tau_1), \theta^*(l(\tau_1))) - \frac{w_{\widetilde{s}}(l(\tau_1), \theta^*(l(\tau_1)))}{w_{\widetilde{s}}(l(\tau_2), \theta^*(l(\tau_2)))} w(l(\tau_2), \theta^*(l(\tau_2))) \right] \frac{\mathrm{d}\, l(\tau_1)}{\mathrm{d}\, \tau} < 0. \tag{22}$$

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 (22) 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 (22). □

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

$$\left[w(l(\tau_1),\theta^*(l(\tau_1))) - \frac{w_{\widetilde{s}}(l(\tau_1),\theta^*(l(\tau_1)))}{w_{\widetilde{s}}(l(\tau_2),\theta^*(l(\tau_2)))} w(l(\tau_2),\theta^*(l(\tau_2)))\right] \frac{\mathrm{d}\, l(\tau_1)}{\mathrm{d}\,\tau} > 0. \tag{23}$$

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 (23) 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 (23). 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 30).

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

$$w(l(1),\theta^*(l(1)))\frac{\mathrm{d}\, l(1)}{\mathrm{d}\,\tau} < \frac{w_{\widetilde{s}}(l(1),\theta^*(l(1)))}{w_{\widetilde{s}}(l(0),\theta^*(l(0)))} w(l(0),\theta^*(l(0)))\frac{\mathrm{d}\, l(1)}{\mathrm{d}\,\tau} = 0, \tag{24}$$

where the equality follows from the construction of $l(\tau)$. Using that wages are identical to the

L-derivative of F, inequality (24) 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

$$F(l(0), \theta^*(L')) \leq F(l(0), \theta^*(L)) = F(l(1), \theta^*(L)) < F(l(1), \theta^*(L')), \tag{25}$$

where the last inequality is strict because we select θ^* as the supremum of the maximizer set in equation (2) (here this means that we select the most polarizing technology from the maximizer set).
Using the mean value theorem, equation (25) 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

$$\left[w(l(\tau'), \theta^*(L')) - c w(l(\tau'), \theta^*(L)) \right] \frac{\mathrm{d}\, l(\tau')}{\mathrm{d}\, \tau} > 0. \tag{26}$$

But now set c equal to $w_2(l(\tau'), \theta^*(L')) / w_2(l(\tau'), \theta^*(L))$. Inequality (26) then reduces to

$$\begin{aligned} &\left[w_1(l(\tau'), \theta^*(L')) - \frac{w_2(l(\tau'), \theta^*(L'))}{w_2(l(\tau'), \theta^*(L))} w_1(l(\tau'), \theta^*(L)) \right] \frac{\mathrm{d}\, l_1(\tau')}{\mathrm{d}\, \tau} \\ &\qquad + \left[w_3(l(\tau'), \theta^*(L')) - \frac{w_2(l(\tau'), \theta^*(L'))}{w_2(l(\tau'), \theta^*(L))} w_3(l(\tau'), \theta^*(L)) \right] \frac{\mathrm{d}\, l_3(\tau')}{\mathrm{d}\, \tau} > 0. \end{aligned}$$

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 $\tilde{x} \in X$ such that (i) $K_x > 0$ only if $x \leq \tilde{x}$ and (ii) $L_{s,x} > 0$ only if $x \geq \tilde{x}$ for all s. Moreover, it is obvious that $K_x > 0$ for all $x < \tilde{x}$, 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 $\tilde{x}$, the assignment of labor to tasks in $[\tilde{x}, \overline{x}]$ is the same as in a model without capital and with a task set of $[\tilde{x}, \overline{x}]$. 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 $\tilde{x}$ 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 $\tilde{x}$ 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 $K_{\tilde{x}} = 0$ and $L_{s,\tilde{x}} > 0$. □

Proof of Lemma 2. Consider first existence and uniqueness of the exogenous technology equilibrium. For any $\tilde{x} \in (\underline{x}, \overline{x})$, the assignment of labor to the task set $[\tilde{x}, \overline{x}]$ 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 $[\underline{x}, \tilde{x})$ 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 (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 $\tilde{x}$ 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 $\tilde{x}^*$ satisfies

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

in any endogenous technology equilibrium (otherwise task producers could choose a different $\tilde{x}$ and earn positive profits thereby). Moreover, any such $\tilde{x}$ 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 $F(L, \tilde{x})$ is continuous in $\tilde{x}$ and X is compact.[31]

For uniqueness, suppose that there are two equilibrium technologies $\tilde{x}_1^* < \tilde{x}_2^*$ at some labor supply L. Then, using the assumption about comparative advantage between capital and labor (Assumption 1), (E6) implies that the least-skilled worker earns less under $\tilde{x}_1^*$ than under $\tilde{x}_2^*$, that is,

$$w_{\underline{s}}(L, \tilde{x}_1^*) < w_{\underline{s}}(L, \tilde{x}_2^*). \tag{27}$$

On the other hand, we know from point 3 of Lemma 2 that skill premia must also be smaller under $\tilde{x}_1^*$, that is, $w(L, \tilde{x}_1^*) \preceq^p w(L, \tilde{x}_2^*)$. Finally, since $F(L, \tilde{x}_1^*) = F(L, \tilde{x}_2^*)$ and F is linear homogeneous in L, we have

$$\int_S w_s(L, \tilde{x}_1^*)\, \mathrm{d}s = \int_S w_s(L, \tilde{x}_2^*)\, \mathrm{d}s.$$

In combination with the fact that skill premia are greater under $\tilde{x}_2^*$, this requires that the least-skilled worker earns less under $\tilde{x}_2^*$, a contradiction to inequality (27). □

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

$$\int_X Y_x^{**} p_x^*\, \mathrm{d}x > \int_X Y_x^* p_x^*\, \mathrm{d}x, \tag{28}$$

[31] To see that $F(L, \tilde{x})$ is continuous in $\tilde{x}$, note that (E1) and the final good production function (equation (7)) imply existence of an aggregate production function $F'(L, m, \{K_x\}_x, \tilde{x})$ that is continuous in $\tilde{x}$. Since in the exogenous technology equilibrium $\{K_x\}_x$ and m are such that they maximize aggregate production, the reduced form production function $F(L, \tilde{x})$ is the upper envelope of $\{F'(L, m, \{K_x\}_x, \tilde{x})\}_{m, K_x}$. As the upper envelope of a family of continuous (in $\tilde{x}$) functions, F is itself continuous.

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 (28)). 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 equilibrium 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 (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 (7) in the setting without technology-embodying intermediate goods. Analogously, solving equation (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
$$\frac{w_{\underline{s}}(L, \alpha, \beta)}{w_{\underline{s}}(L, \alpha', \beta')} = \frac{\alpha'}{\alpha} \frac{\gamma(\underline{s}, \widetilde{x}(L, \alpha, \beta))}{\gamma(\underline{s}, \widetilde{x}(L, \alpha', \beta'))} > 1. \tag{29}$$
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 (29).

Consider now the endogenous technology equilibrium, where α and β are determined via (E10)′. Take any
$$(\alpha^*(L), \beta^*(L)) \in \operatorname*{argmax}_{(\alpha,\beta)\in\mathcal{D}} F(L, \alpha, \beta),$$
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 (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 (12) and (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.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 5.4. The Northern economy has the same structure as the closed economy of Section 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

$$Y_x^N = \int_0^1 \left(\alpha_i^{N\frac{1}{\kappa}} q_{\alpha,i}^{NN} + \alpha_i^{S\frac{1}{\kappa}} q_{\alpha,i}^{SN} \right)^{\kappa} \mathrm{d}\, i\, K_x^{N1-\kappa} + \gamma^N(m^{N-1}(x), x) L^N_{m^{N-1}(x)} \frac{\mathrm{d}\, m^{N-1}(x)}{\mathrm{d}\, x}.$$

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),

$$p_\alpha^N = \kappa\alpha^N q_\alpha^{NN\,\kappa-1} \int_{\underline{x}}^{\tilde{x}^T} p_x K_x^{N\,1-\kappa}\,\mathrm{d}\,x,$$

and in the South,

$$p_\alpha^N = \kappa\alpha^N q_\alpha^{NS\,\kappa-1} \int_{\underline{x}}^{\tilde{x}^T} p_x K_x^{S\,1-\kappa}\,\mathrm{d}\,x.$$

Now let $q_\alpha := q_\alpha^N + q_\alpha^S = q_\alpha^N$ denote world production, and let $s_\alpha := q_\alpha^{NN}/q_\alpha$ 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 $x < \tilde{x}^T$ 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 $x < \tilde{x}^T$, where $K_x := K_x^N + K_x^S$ 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

$$\begin{aligned} p_\alpha^N &= \kappa\alpha^N s_\alpha^{\kappa-1} q_\alpha^{\kappa-1} \int_{\underline{x}}^{\tilde{x}^T} p_x K_x^{N\,1-\kappa}\,\mathrm{d}\,x \\ &= \kappa\alpha^N q_\alpha^{\kappa-1} \int_{\underline{x}}^{\tilde{x}^T} p_x \left(\frac{K_x^N}{s_\alpha}\right)^{1-\kappa} \mathrm{d}\,x \\ &= \kappa\alpha^N q_\alpha^{\kappa-1} \int_{\underline{x}}^{\tilde{x}^T} p_x K_x^{1-\kappa}\,\mathrm{d}\,x, \end{aligned} \tag{30}$$

which has the same form as the corresponding inverse demand in the closed economy (see equation (10)). Profits of α-monopolists at a given α are then given by

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

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):

$$\rho\kappa q_\alpha^\kappa \int_{\underline{x}}^{\tilde{x}^T} p_x K_x^{1-\kappa}\,\mathrm{d}\,x = p_D \alpha^{N\,\frac{1-\rho}{\rho}}. \tag{31}$$

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

$$Y_x = \gamma^N(m^{T-1}(x), x) L^N_{m^{T-1}(x)} \frac{\mathrm{d}\, m^{T-1}(x)}{\mathrm{d}\, x} + \gamma^S(m^{T-1}(x), x) L^S_{m^{T-1}(x)} \frac{\mathrm{d}\, m^{T-1}(x)}{\mathrm{d}\, x}$$
$$= \gamma^N(m^{T-1}(x), x) \left(L^N_{m^{T-1}(x)} + \Delta L^S_{m^{T-1}(x)} \right) \frac{\mathrm{d}\, m^{T-1}(x)}{\mathrm{d}\, x}$$

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:

$$p^N_{\beta,i} = \kappa \beta^N_i q^{NN\,\kappa-1}_{\beta,i} \left[\int_{\underline{x}}^{\overline{x}} \left(Y^{NN}_x + Y^{SN}_x \right)^{\frac{\epsilon-1}{\epsilon}} \mathrm{d}\, x \right]^{\frac{\epsilon(1-\kappa)}{\epsilon}}$$
$$= \kappa \beta^N_i q^{NN\,\kappa-1}_{\beta,i} s^{1-\kappa}_\beta \left[\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \mathrm{d}\, x \right]^{\frac{\epsilon(1-\kappa)}{\epsilon}},$$

and inverse demand in the South:

$$p^N_{\beta,i} = \kappa \beta^N_i q^{NS\,\kappa-1}_{\beta,i} \left[\int_{\underline{x}}^{\overline{x}} \left(Y^{SS}_x + Y^{NS}_x \right)^{\frac{\epsilon-1}{\epsilon}} \mathrm{d}\, x \right]^{\frac{\epsilon(1-\kappa)}{\epsilon}}$$
$$= \kappa \beta^N_i q^{NS\,\kappa-1}_{\beta,i} (1 - s_\beta)^{1-\kappa} \left(\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \mathrm{d}\, x \right)^{\frac{\epsilon(1-\kappa)}{\epsilon}}.$$

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

$$p_\beta = \kappa \beta^N q^{\kappa-1}_\beta \left(\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \mathrm{d}\, x \right)^{\frac{\epsilon(1-\kappa)}{\epsilon}}, \tag{32}$$

which is the desired inverse demand equation in terms of world quantities. Profits of β-monopolists are then given by

$$\pi_{\beta,i}(\beta^N_i) = \max_q \left\{ \kappa \beta^N_i q^\kappa \left(\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \mathrm{d}\, x \right)^{\frac{\epsilon(1-\kappa)}{\epsilon}} - \eta_\beta q \right\} - p_D \beta^{N\,\frac{1}{\rho}}_i,$$

such that the first order condition for a profit maximum in β^N_i becomes (using symmetry to drop the i):

$$\rho\kappa q_\beta^\kappa \left(\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \, \mathrm{d}\, x \right)^{\frac{(1-\kappa)\epsilon}{\epsilon-1}} = p_D {\beta^N}^{\frac{1-\rho}{\rho}}. \tag{33}$$

Finally, let $Y = Y^N + Y^S$ be world (gross) production of the final good. Using the share s_β, we can write

$$\begin{aligned} Y &= \beta^N (s_\beta q_\beta)^\kappa \left(\int_{\underline{x}}^{\overline{x}} (s_\beta Y_x)^{\frac{\epsilon-1}{\epsilon}} \, \mathrm{d}\, x \right)^{\frac{(1-\kappa)\epsilon}{\epsilon-1}} + \beta^N ((1-s_\beta) q_\beta)^\kappa \left(\int_{\underline{x}}^{\overline{x}} ((1-s_\beta) Y_x)^{\frac{\epsilon-1}{\epsilon}} \, \mathrm{d}\, x \right)^{\frac{(1-\kappa)\epsilon}{\epsilon-1}} \\ &= \beta^N q_\beta^\kappa \left(\int_{\underline{x}}^{\overline{x}} Y_x^{\frac{\epsilon-1}{\epsilon}} \, \mathrm{d}\, x \right)^{\frac{(1-\kappa)\epsilon}{\epsilon-1}}. \end{aligned} \tag{34}$$

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 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 (34);
- (E3)″ q_β satisfies equation (32);
- (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 (30);
- (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 (31), (33), 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 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 $L^N + \Delta L^S$.

B. Supplementary Material

B.1. Further Models of Endogenous Technical Change

This section derives equations (1) and (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 $\operatorname{argmax}_{\theta \in \Theta} F(L, \theta)$ 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 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 (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

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

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 3. *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 (3) 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) 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* (1). *Moreover, suppose Assumption 3 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) *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 (1) and (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 (1) and (2), define the following "modified production function" (see also Lemma 4):

$$F(L_i, \theta) := \max_{(q_k)_{k=1,2,\ldots,K}} \left\{ F'(L_i, (\theta_k q_k^\kappa)_{k=1,2,\ldots,K}) - \frac{\eta_k}{\kappa} q_k \right\} - \frac{1}{\kappa} \sum_{k=1}^{K} C_k(\theta_k).$$

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.[32] 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 (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 (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 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

[32]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.

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) and wages by (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 (1).*
Moreover, if $F(L, \theta)$ is concave in L, there exists a symmetric endogenous technology equilibrium with equilibrium technology satisfying equation (2) and wages given by (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).

B.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 5. 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 6. 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 $\overline{\theta} \not\preceq^b \theta^*(L')$, 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). From here on, the proof proceeds analogously to the proof of Theorem 1 starting from equation (16): $\bar{\theta}$ 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 (17) and (18) are reversed, and the contradiction at the end is obtained by observing that $\theta^*(L') \preceq^b \bar{\theta}$ implies that the left-hand-side of inequality (18) 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.[33] 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.

[33]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 $p \preceq^b p'$.