DISCUSSION PAPER SERIES

No. 10244

THE RISE OF THE MACHINES: AUTOMATION, HORIZONTAL INNOVATION AND INCOME INEQUALITY

David Hemous and Morten Olsen

INTERNATIONAL MACROECONOMICS

Centre for Economic Policy Research

ISSN 0265-8003

THE RISE OF THE MACHINES: AUTOMATION, HORIZONTAL INNOVATION AND INCOME INEQUALITY

David Hemous and Morten Olsen

Discussion Paper No. 10244
November 2014
Submitted 05 November 2014

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801
www.cepr.org

THE RISE OF THE MACHINES: AUTOMATION, HORIZONTAL INNOVATION AND INCOME INEQUALITY[†]

Abstract

We construct an endogenous growth model of directed technical change with automation (the introduction of machines which replace low-skill labor and complement high-skill labor) and horizontal innovation (the introduction of new products, which increases demand for both types of labor). Such an economy endogenously follows three phases. First, low-skill wages are low, which induces little automation, such that income inequality and labor's share of GDP are constant. Second, as low-skill wages increase, investment in automation is stimulated, which depresses the future growth rate of low-skill wages (potentially to negative), and reduces the total labor share. Finally, the share of automated products stabilizes and the economy moves toward an asymptotic steady state, where low-skill wages grow but at a lower rate than high-skill wages. This model therefore delivers persistently increasing wage inequality and stagnating real wages for low skill workers for an extended period of time, features of modern labor markets which have been difficult to reconcile with the theoretical literature on economic growth. We further include middle-skill workers, which allows the model to generate a phase of wage polarization after one where labor income inequality increases uniformly. Finally, we show that an endogenous labor supply response in this framework can quantitatively account for the evolution of the skill premium, the skill ratio and the labor share in the US since the 1960s.

JEL Classification: E23, E25, O31, O33 and O41
Keywords: automation, capital-skill complementarity, directed technical change, factor share, horizontal innovation, income inequality and wage polarization

David Hemous david.hemous@insead.edu
INSEAD and CEPR

Morten Olsen molsen@iese.edu
IESE

[†] Morten Olsen gratefully acknowledges the financial support of the European Commission under the Marie Curie Research Fellowship programme (Grant Agreement PCIG11-GA-2012-321693) and the Spanish Ministry of Economy and Competitiveness (Project ref: ECO2012-38134). We thank Daron Acemoglu, Philippe Aghion, Pol Antràs, Tobias Broer, Steve Cicala, Per Krusell, Brent Neiman, Jennifer Page, Andrei Shleifer and Che-Lin Su amongst others for helpful comments and suggestions. We also thank seminar and conference participants at IIES, University of Copenhagen, Warwick, UCSD, UCLA Anderson, USC Marshall, Barcelona GSE Summer Forum, the 6th Joint Macro Workshop at Banque de France, the 2014 SED meeting, the NBER Summer Institute and the 2014 EEA meeting. We thank Ria Ivandic and Marton Varga for excellent research assistance.

1 Introduction

How does the automation of production processes affect the distribution of income? By allowing for the use of machines, automation reduces the demand for some type of labor, particularly low-skill labor. This mechanism has received mounting empirical support (e.g. Autor, Levy and Murnane, 2003),[1] and it is often considered a major cause of the sharp rise in income inequality in developed countries since 1970. The seemingly ever increasing capabilities of machines only makes this concern more relevant today. Yet, economists often argue that technological development also creates new products, which boosts the demand for labor, and certainly, many of today's jobs did not exist just a few decades ago.

This paper develops an analytical framework to explore these dynamics, in which the creation of new products and the automation of the production process interact to drive changes in inequality and growth. Such a framework must center on the economic incentives faced by innovators. Consequently, our framework is one of directed technical change between automation and horizontal innovation (the introduction of new products), yet it departs from the previous literature in that it does not focus on factor-augmenting technical change (such as Acemoglu, 1998). Contrary to that literature, our framework is able to account for the following salient features of the evolution of the income distribution in the last 40 years: a continuous increase in labor income inequality, a decline in the labor share, and stagnating (possibly declining) real wages for low-skill workers. In addition, the framework is malleable, which allows us to develop several extensions. One such extension can account for a phase of wage polarization following a phase of uniform increase in income inequality—consistent with the US experience. The model can further be used quantitatively to jointly account for the evolutions of the college premium and the labor share since the 1960s.

Our model is an expanding variety growth model with low-skill and high-skill workers. Horizontal innovation, modeled as in Romer (1990), increases the demand for both low- and high-skill workers. Automation allows for the replacement of low-skill workers with machines in production. It takes the form of a secondary innovation in existing product lines, similar to secondary innovations in Aghion and Howitt (1996) (though their focus is the interplay between applied and fundamental research, and not automation). Within

[1] Autor, Katz and Krueger (1998) and Autor, Levy and Murnane (2003) use cross-sectional data to demonstrate that computerization is associated with relative shifts in demand favoring college educated workers. Such evidence also exists at the firm level (Bartel, Ichniowski and Shaw, 2007).

a firm, automation increases the demand for high-skill workers but reduces the demand for low-skill workers. We refer to firms that can use machines as 'automated' firms. 'Non-automated' firms can only produce using low-skill and high-skill labor. Once invented, a specific machine is produced with the same technology as the consumption good.

We study the transitional dynamics of this economy and our results highlight the role played by low-skill wages. The cost advantage of an automated over a non-automated firm increases with the real wage (since the price of a machine is the price of the consumption good). As a result, an economy with an initially low level of technology first goes through a phase where growth is mostly generated by horizontal innovation and the skill premium and the labor share are constant. Only when low-skill wages are sufficiently high will firms invest in automation (firms are forward-looking so naturally they consider the entire future path of low-skill wages). During this second phase, the share of automated firms increases, low-skill workers lose relatively to high-skill workers and, depending on parameters, the real low-skill wage may temporarily decrease. The total labor share decreases progressively, in line with recent evidence (Karabarbounis and Neiman, 2013). Finally, the share of automated products stabilizes as the entry of new, non-automated products compensates for the automation of existing ones that had yet to be automated. In this third phase, low-skill wages grow asymptotically; intuitively, the presence of non-automated products ensures that low-skill workers and machines are only imperfect substitutes at the aggregate level. With an increasing quantity of machines, the relative cost of a low-skill worker and a machine, which here is also the real wage, must grow. Yet, low-skill wages must grow at a lower rate than high-skill wages, since an increase in low-skill wages further increases the cost advantage of automated firms and thereby the skill premium (thus the economy does not feature a balanced growth path). The total labor share stabilizes and growth results mostly from automation.

We show that the automation technology has an ambiguous impact on low-skill wages. An economy with a more productive automation technology enters the second phase sooner and automation is more intensive resulting in temporarily lower low-skill wages. Yet, a more productive automation technology also boosts horizontal innovation in the long-run, so that, asymptotically, low-skill wages grow faster.

We extend the model to include a supply response in the skill distribution, and calibrate it to match the evolution of the skill premium, the skill ratio, the labor share and GDP growth since the 60's. This exercise demonstrates that our model is able to replicate the trends in the data quantitatively.

Recent empirical work has increasingly found that workers in the middle of the income distribution are most adversely affected by technological progress. To address this, we extend the baseline model to include middle-skill workers as a separate skill-group. Firms either rely on low-skill workers or middle-skill workers (but not both) and the two skill-groups are symmetric except that automating to replace middle-skill workers is more costly (or machines are less productive in middle-skill firms). This means that the automation of low-skill workers' tasks happens first, with a delayed automation process for the tasks of middle-skill workers. We show that this difference can reproduce important trends in the United States income distribution. In a first period, there is a uniform dispersion of the income distribution, as low-skill workers' products are rapidly automated but middle-skill ones are not; while in the second period there is wage polarization: low-skill workers' share of automated products is stabilized, and middle-skill products are more rapidly automated.

In another extension, we relax the assumption that, once invented, machines are produced with the same technology as the consumption good. In this setting, the real price of machines declines over time, which may lead to decreasing real low-skill wages in the long-run.

There is a small theoretical literature on labor-replacing technology. In Zeira (1998), firms have access to two technologies which differ in their capital intensity. Adoption of the capital intensive technology is analogous to automation in our model, and both are encouraged by higher low-skill wages. In his model, higher low-skill wages are generated exogenously by an increase in TFP, while in ours, low-skill wages endogenously increase through horizontal innovation.[2] Acemoglu (2010) shows that labor scarcity induces innovation (the Habbakuk hypothesis), if and only if innovation is labor-saving, that is, if it reduces the marginal product of labor. Neither paper analyzes labor-replacing innovation in a fully dynamic model nor focuses on income inequality, as we do. Peretto and Seater (2013) build a dynamic model of factor-eliminating technical change where firms learn how to replace labor with capital, a process which bears strong similarities with automation in our framework. They do not, however, focus on income inequality either and since their model only features one source of growth, wages are constant so that the incentive to automate does not change over time.

A large literature has used skill-biased technical change (SBTC) as a possible explanation for the increase in the skill premium in developed countries since the 1970's,

[2] More recently, Zeira (2013) uses a model with a similar mechanism to show that globalization can generate diverging growth rates across countries.

despite a large increase in the relative supply of skilled workers (see Hornstein, Krusell and Violante, 2005, for a more complete literature review). One can categorize theoretical papers into one of three strands. The first strand emphasizes the hypothesis of Nelson and Phelps (1966) that more skilled workers are better able to adapt to technological change, in which case a technological revolution (like the IT revolution) increases the relative demand for skilled workers and increases income inequality. Several papers have formalized this idea (including Aghion and Howitt, 1997; Lloyd-Ellis, 1999; Caselli, 1999; Galor and Moav, 2000, and Aghion, Howitt and Violante, 2002). However, such theories mostly explain transitory increases in inequality whereas inequality has been increasing for decades. Our model, on the contrary, introduces a mechanism that creates permanent and widening inequality.

A second strand sees the complementarity between capital and skill as the source for the increase in the skill premium. Krusell, Ohanian, Ríos-Rull and Violante (2000) formalize this idea by developing a framework where capital equipment and high-skill labor are complements. To this, they add the empirically observed increase in the stock of capital equipment, and show that their model can account for most of the variation in the skill premium. Our model shares features with their framework: machines play an analogous role to capital equipment in their model, since they are more complementary with high-skill labor than with low-skill labor. The focus of our paper is different though since we seek to explain why innovation has been directed towards automation in the first place.

Finally, a third branch of the literature, originally presented by Katz and Murphy (1992), considers technology to be either high-skill or low-skill labor augmenting. This approach has been used empirically within a relative supply and demand framework of these two skill groups—typically college and non-college graduates—to infer the extent of skill-biased technical change from changes in the relative labor supply and the skill premium. For instance, Goldin and Katz (2008) find that in the US, technical change has been skill-biased throughout the 20^{th} century. On the theory side, the directed technical change literature (most notably Acemoglu, 1998, 2002 and 2007) also uses factor-augmenting technical change models to endogenize the bias of technical change. In particular, it shows that an increase in the supply of high-skill workers may foster skill biased technical change and increase the skill premium. Such models deliver important insights about inequality and technical change, but they have no role for labor-replacing technology (a point emphasized in Acemoglu and Autor, 2011). In addition, even though income inequality varies, neither high-skill nor low-skill wages can decrease in absolute terms, and their asymptotic growth rates must be the same. The present model is also a directed technical change framework as economic incentives determine whether technical change takes the form of horizontal innovation or automation, but it deviates from the assumption of factor-augmenting technologies and explicitly allows for labor-replacing automation, generating the possibility for (temporary) absolute losses for low-skill workers, and permanently increasing income inequality.

More recently, Autor, Katz, and Kearney (2006, 2008) and Autor and Dorn (2013), amongst others, show that whereas income inequality has continued to increase above the median, there has been a reversal below the median. They argue that the (routine) tasks performed by many middle-skill workers—storing, processing and retrieving information—are more easily done by computers than those performed by low-skill workers, now predominantly working in service occupations. This 'wage polarization' has been accompanied by a 'job polarization' as employment has followed the same pattern of decreasing employment in middle-skill occupations.[3] Acemoglu and Autor (2011) argue that a task-based model where technological progress explicitly allows the replacement of one input, e.g. labor, by another, e.g. capital, in the production of some tasks provides a better explanation for wage and job polarization than 'factor augmenting technical change' models (and in addition, allows for a decrease in the absolute level of wages).[4] In the present paper, automation similarly replaces labor with machines in the production of some goods (and one could interpret the different products as different tasks). However, whereas the 'tasks' literature has considered static models, our framework is dynamic and endogenizes the arrival of automation. In addition, it provides a unified explanation for the relative decline of middle-skill wages since the mid-1980s and the relative decline of low-skill wages in the period before.

Section 2 introduces the baseline model and defines the equilibrium. Section 3 describes the evolution of the economy through three phases and derives the asymptotic steady-state. Section 4 extends the model to analyze wage polarization. Section 5 considers several extensions. Section 6 calibrates the model to the US economy since the

[3]This phenomenon has also been observed and associated with the automation of routine tasks in Europe (Spitz-Oener, 2006; Goos and Manning, 2007 and Goos, Manning and Salomons, 2009). Another explanation for polarization stems from the consumption side and relates the high growth rate of wages for the least-skilled workers with an increase in the demand for services from the most-skilled—and richest—workers, see Mazzolari and Ragusa (2013) and Bárány and Siegel (2014).

[4]A related literature analyzes this non-monotonic pattern in inequality changes through the lens of assignment models where workers of different skill levels are matched to tasks of different skill productivities (e.g. Costinot and Vogel, 2010and Burstein, Morales and Vogel, 2014).

1960s. Section 7 concludes.

2 The Baseline Model

2.1 Preferences and production

We consider a continuous time infinite-horizon economy populated by H high-skill and L low-skill workers. Both types of workers supply labor inelastically and have identical preferences over a single final good of:

$$U_{k,t} = \int_t^\infty e^{-\rho(\tau-t)} \frac{C_{k,\tau}^{1-\theta}}{1-\theta} d\tau,$$

where ρ is the discount rate, $\theta \geq 1$ is the inverse elasticity of intertemporal substitution and $C_{k,t}$ is consumption of the final good at time t by group $k \in \{H, L\}$. H and L are kept constant in our baseline model, but we consider the case where workers choose occupations based on relative wages and heterogeneous skill-endowments in Section 5.1.

The final good is produced by a competitive industry combining an endogenous set of intermediate inputs, $i \in [0, N_t]$ using a CES aggregator:

$$Y_t = \left(\int_0^{N_t} y_t(i)^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}},$$

where $\sigma > 1$ is the elasticity of substitution between these inputs and $y_t(i)$ is the use of intermediate input i at time t. As in Romer (1990), an increase in N_t represents a source of technological progress and increases the production possibility set. We normalize the price of Y_t to 1 at all points in time and drop time subscripts when there is no ambiguity. The demand for each variety is:

$$y(i) = p(i)^{-\sigma} Y, \tag{1}$$

where $p(i)$ is the price of intermediate input i and the normalization implies that the ideal price index, $[\int_0^{N_t} p(i)^{1-\sigma} di]^{1/(1-\sigma)}$ equals 1.

Each intermediate input is produced by a monopolist who owns the perpetual rights of production. She can produce the intermediate input by combining low-skill labor, $l(i)$, high-skill labor, $h(i)$, and, possibly, type-i machines, $x(i)$, using the production function:

$$y(i) = \left[l(i)^{\frac{\epsilon-1}{\epsilon}} + \alpha(i) \left(\tilde{\varphi} x(i) \right)^{\frac{\epsilon-1}{\epsilon}} \right]^{\frac{\epsilon\beta}{\epsilon-1}} h(i)^{1-\beta},$$

where $\alpha(i) \in \{0, 1\}$ is an indicator function for whether or not the firm has access to an automation technology which allows for the use of machines. If the firm is not automated ($\alpha(i) = 0$), production takes place using a standard Cobb-Douglas production function with only low-skill and high-skill labor with a low-skill factor share of β. If the firm is automated ($\alpha(i) = 1$) it can also use machines in the production process. We assume that the elasticity of substitution between machines and low-skill workers, ϵ, is strictly greater than 1 and allow for perfect substitutability, in which case $\epsilon = \infty$ and the production function is $y(i) = [l(i) + \alpha(i)\tilde{\varphi}x(i)]^{\beta} h(i)^{1-\beta}$. The parameter $\tilde{\varphi}$ is the relative productivity advantage of machines over low-skill workers. Firm i becomes automated when $\alpha(i)$ switches from 0 to 1. We will denote by G_t the share of automated firms. It is because $\alpha(i)$ is not fixed, but changes over time, that our model captures the notion that machines can replace low-skill labor in new tasks. A model in which $\alpha(i)$ were fixed for each firm would only allow for machines to be used more intensively in production, but always for the same tasks. Throughout the paper we will refer to x as 'machines', though our interpretation also includes any form of computer inputs, algorithms, the services of cloud-providers, etc. In our baseline model, machines are an intermediate input in production and are wholly consumed in the production process. In Section 5.3, we examine the case where machines do not depreciate immediately and are part of a capital stock. In addition, we assume here that, once invented, machines of type i are produced competitively one for one with the final good, so that the price of an existing machine for an automated firm is always equal to 1. Though a natural starting point, this is an important assumption and Section 5.2 presents a version of the model which relaxes it. Importantly, this does not imply that our model cannot capture the notion of a decline in the real cost of equipment: indeed, automation for firm i can equivalently be interpreted as a decline of the price of machines i from infinity to 1.

2.2 Innovation

There are two sources of technological progress in this model: automation and horizontal innovation. We assume that an automated firm remains so forever and that becoming automated requires an investment. More specifically, a non-automated firm which hires $h_t^A(i)$ high-skill workers in automation research, becomes automated according to a Poisson process with rate $\eta G_t^{\widetilde{\kappa}} N_t^{\kappa} h_t^A(i)^{\kappa}$. $\eta > 0$ denotes the productivity of the automation technology, $\kappa \in (0, 1)$ measures the concavity of the automation technology, $G_t^{\widetilde{\kappa}}$, $\widetilde{\kappa} \in [0, \kappa]$, represents possible knowledge spillovers from the share of automated products, and N_t^{κ} represents knowledge spillovers from the total number of intermediate inputs (these spillovers are necessary to ensure that both automation and horizontal innovation can take place in the long-run).[5] We define the total mass of high-skill workers working in automation: $H_t^A \equiv \int_0^{N_t} h_t^A(i)di$. Our set-up can be interpreted in two ways. From one standpoint, machines are intermediate input-specific and each producer needs to invent his own machine, which, once invented, is produced with the same technology as the consumption good.[6] From a second standpoint, machines are produced by the final good sector, and each intermediate input producer must spend resources in adapting the machine to his product line.

Horizontal innovation occurs in a standard manner. New intermediary inputs are developed by high-skill workers according to a linear technology with productivity γN_t (where $\gamma > 0$ measures the productivity of the horizontal innovation technology). With H_t^D high-skill workers pursuing horizontal innovation, the mass of intermediate inputs evolves according to:

$$\dot{N}_t = \gamma N_t H_t^D.$$

We assume that firms do not exist before their product is created. Coupled with our assumption that automation follows a continuous Poisson process, new products must then be born non-automated. This feature of the model is motivated by the idea that when a task is new and unfamiliar, the flexibility and outside experiences of workers allow them to solve unforeseen problems. As the task becomes routine and potentially codifiable a machine (or an algorithm) can perform it (as argued by Autor, 2013). In reality, some new tasks may be sufficiently close to older ones that no additional investment would be required to automate them immediately. As explained in section 3.7, our

[5] An alternative way of interpreting this functional form is the following: let there be a mass 1 of firms with N_t products (instead of assuming that each individual i is a distinct firm), then this functional form means that when a firm hires a mass $N_t h_t^A$ of high-skill workers in automation each of its products gets independently automated with a Poisson rate of $\eta G_t^{\widetilde{\kappa}} \left(N_t h_t^A\right)^{\kappa}$.

[6] Alternatively, machine-i may be invented by an outside firm and then sold to the intermediate input producer. With such market structure the rents from automation would be divided between the intermediate input producer and the machine producer. Up to a constant representing the bargaining power of each party, it would not affect any of our results. Yet another alternative would be to have entrants undertaking automation and potentially displacing the original firm. This would not qualitatively affect the equilibrium as long as the incumbent has a positive probability of becoming automated.

results still carry through if we assume that only a share of the new products are born non-automated or even if automation is only undertaken at the entry stage.

Define $H_t^P \equiv \int_0^{N_t} h_t(i)di$ as the total mass of high-skill workers involved in production. Factor markets clearing implies that

$$\int_0^{N_t} l_t(i)di = L,\ H_t^A + H_t^D + H_t^P = H. \tag{2}$$

2.3 Equilibrium wages

In this subsection, we take the technological levels N, G and the mass of high-skill workers in production H^P as given and show how low-skill wages (denoted w) and high-skill wages (denoted v) are determined in equilibrium. First, note that all automated firms are symmetric and therefore behave in the same way. Similarly all non-automated firms are symmetric. The unit cost of intermediate input i is given by:

$$c(w, v, \alpha(i)) = \beta^{-\beta}(1-\beta)^{-(1-\beta)}\left(w^{1-\epsilon} + \varphi\alpha(i)\right)^{\frac{\beta}{1-\epsilon}} v^{1-\beta}, \tag{3}$$

where $\varphi \equiv \tilde{\varphi}^\epsilon$, $c(\cdot)$ is strictly increasing in both w and v and $c(w, v, 1) < c(w, v, 0)$ for all $w, v > 0$ (automation reduces costs). The monopolist charges a constant markup over costs such that price is $p(i) = \sigma/(\sigma - 1) \cdot c(w, v, \alpha(i))$.

Using Shepard's lemma and equations (1) and (3) delivers the demand for low-skill labor of a single firm.

$$l(w, v, \alpha(i)) = \beta\frac{w^{-\epsilon}}{w^{1-\epsilon} + \varphi\alpha(i)}\left(\frac{\sigma-1}{\sigma}\right)^\sigma c(w, v, \alpha(i))^{1-\sigma}Y, \tag{4}$$

which is decreasing in w and v. The effect on demand for low-skill labor in a given firm following automation is generally ambiguous. This is due to the combination of a negative *substitution* effect (the ability of the firm to substitute machines for low-skill workers) and a positive *scale* effect (the ability of the firm to employ machines decreases overall costs, lowers prices and increases production). Here we focus on labor-substituting innovation and impose throughout that $\mu \equiv \beta(\sigma - 1)/(\epsilon - 1) < 1$ (that is the elasticity of substitution between machines and low-skill labor is large enough), which is necessary and sufficient for the substitution effect to dominate and ensure $l(w, v, 1) < l(w, v, 0)$ for all $w, v > 0$. It is straightforward to show that the relative use of skilled labor is increasing in automation, consistent with a large literature that finds that technological improvements — in particular computerization — are associated with increased relative skill-use.

Let $x\,(w, v)$ denote the use of a machines by an automated firm. The relative use of machines and low-skill labor for such a firm is then:

$$\frac{x(w,v)}{l(w,v,1)} = \varphi w^{\epsilon}, \tag{5}$$

which is decreasing in w as the real wage is also the price of low-skill labor relative to machines.

The iso-elastic demand (1), coupled with constant mark-up $\sigma/(\sigma-1)$, implies that revenues are given by $R(i) = \left((\sigma-1)/\sigma\right)^{\sigma-1} c(i)^{1-\sigma} Y$ and that a share $1/\sigma$ of revenues accrues to the monopolists as profits: $\pi(i) = R(i)/\sigma$. Using (3), the relative revenues (and relative profits) of non-automated and automated firms are then given by:

$$\frac{R(w,v,0)}{R(w,v,1)} = \frac{\pi(w,v,0)}{\pi(w,v,1)} = \left(1+\varphi w^{\epsilon-1}\right)^{-\mu}, \tag{6}$$

which is a decreasing function of w. Since non-automated firms rely more heavily on low-skill labor, their relative market share drops with higher low-skill wages.

Since monopolists capture the same share of the revenues as profits whether a firm is automated or not, aggregate profits are a constant share $1/\sigma$ of output Y (though not of GDP, see below). Similarly, the factor share of high-skill workers in production is the same constant whether a firm is automated or not, so that high-skill workers in production receive a share $\nu_h = (1-\beta)(\sigma-1)/\sigma$ of output. Aggregating over all high-skill workers in *production*, we get that

$$vH^P = (1-\beta)\frac{\sigma-1}{\sigma}Y. \tag{7}$$

Using factor demand functions, the share of revenues accruing to low-skill labor is given by $\nu_l(w,v,\alpha(i)) = \frac{\sigma-1}{\sigma}\beta\left(1+\varphi w^{\epsilon-1}\alpha(i)\right)^{-1}$, and is lower for automated than non-automated firms. The aggregate revenues of low-skill workers can be obtained by summing up over all intermediate inputs:

$$wL = N\left[GR(w,v,1)\nu_l(w,v,1) + (1-G)R(w,v,0)\nu_l(w,v,0)\right],$$

with an analogous expression for high-skill workers in production. Taking the ratio of these two expressions and using the expression for relative revenues (6) gives the following lemma.

Lemma 1. *For $\epsilon < \infty$, the high-skill wage premium is given by*[7]

$$\frac{v}{w} = \frac{1-\beta}{\beta}\frac{L}{H^P}\frac{G + (1-G)(1+\varphi w^{\epsilon-1})^{-\mu}}{G\left(1+\varphi w^{\epsilon-1}\right)^{-1} + (1-G)(1+\varphi w^{\epsilon-1})^{-\mu}}. \tag{8}$$

For given L/H^P and $G > 0$, the skill premium is increasing in the absolute level of low-skill wages, which means that if G is bounded above 0, low-skill wages cannot grow at the same rate as high-skill wages in the long-run. This is the case both because automated firms substitute machines for low-skill workers (the term $(1+\varphi w^{\epsilon-1})^{-1}$) and because higher wages improve automated firms cost-advantage and thereby their market share (as reflected by the term $(1+\varphi w^{\epsilon-1})^{-\mu}$ in equation 8). Only in the special case of no automated firms, $G = 0$, will the aggregate factor shares inherit that of a non-automated firm and $v/w = (1-\beta)L/(\beta H^P)$.

With constant mark-ups, the cost equation (3) and the price normalization gives:

$$\left(G\left(\varphi + w^{1-\epsilon}\right)^{\mu} + (1-G)w^{\beta(1-\sigma)}\right)^{\frac{1}{1-\sigma}} v^{1-\beta} = \frac{\sigma-1}{\sigma}\beta^{\beta}\left(1-\beta\right)^{1-\beta} N^{\frac{1}{\sigma-1}}. \tag{9}$$

We label this a *productivity* condition, as it shows the positive relationship between real wages and the level of technology given by N, the number of intermediate inputs, and G the share of automated firms. Together (8) and (9) determine real wages as a function of technology N, G and the share of high-skill workers engaged in production H^P.

Though the production function implies that, at the firm level, the elasticity of substitution between high-skill labor and machines is equal to that between high-skill and low-skill labor, this does not imply that the same holds at the aggregate level. Therefore our paper is not in contradiction to Krusell et al. (2000), who argue that the aggregate elasticity of substitution between high-skill and low-skill labor is greater than the one between high-skill labor and machines. In fact, one can show that the Morishima elasticity of substitution between high-skill labor and low-skill labor is close to 1 when low-skill wages are low and close to $1 + \beta(\sigma - 1)$ when they are high; while the one between high-skill labor and machines is close to ϵ in the former case and to 1 in the latter, so that the ordering is reversed as low-skill wages grow.[8]

[7] When machines and low-skill workers are perfect substitutes, $\epsilon = \infty$, the skill premium is given by $\frac{v}{w} = \frac{1-\beta}{\beta}\frac{L}{H^P}$ if $w < \widetilde{\varphi}^{-1}$ such that no firm uses machines, and $\frac{v}{w} = \frac{1-\beta}{\beta}\frac{L}{H^P}\frac{G+(1-G)(\widetilde{\varphi}w)^{-1}}{(1-G)(\widetilde{\varphi}w)^{-1}}$ if $w > \widetilde{\varphi}^{-1}$.

[8] There is no universal agreement on the proper extension of the elasticity of substitution when

Given the amount of resources devoted to production (L, H^P), the static equilibrium is closed by the final good market clearing condition:

$$Y = C + X \tag{10}$$

where $C = C_L + C_H$ is total consumption and $X = \int_0^N x(i)di$ is total use of machines. Y differs from GDP for two reasons: it includes intermediate inputs and it does not include R&D investments, which are done by high-skill labor. Hence, we have:

$$GDP = Y - X + v\left(H^D + H^A\right). \tag{11}$$

2.4 Innovation allocation

We now study how innovation is determined in equilibrium. To do this, we first define V_t^A the value of an automated firm and r_t the economy wide interest rate. The asset pricing equation for an automated firm is then given by

$$r_t V_t^A = \pi_t^A + \dot{V}_t^A, \tag{12}$$

where we ease notation by defining $\pi_t^A \equiv \pi(w_t, v_t, 1)$ as profits at time t by an automated firm. The equation states that the required return on holding an automated firm, V_t^A, must equal the instantaneous profits plus appreciation. An automated firm only maximizes instantaneous profits and has no intertemporal investment decisions to make.

A non-automated firm, on the other hand, has to decide how much to invest in automation. Denoting by V_t^N the value of a non-automated firm, we get the corresponding asset pricing equation:

$$r_t V_t^N = \pi_t^N + \eta G_t^{\widetilde{\kappa}}\left(N_t h_t^A\right)^{\kappa}\left(V_t^A - V_t^N\right) - v_t h_t^A + \dot{V}_t^N, \tag{13}$$

where $\pi_t^N \equiv \pi(w_t, v_t, 1)$ and h_t^A is the mass of high-skill workers hired in automation research by a single non-automated firm (so that $H_t^A = (1 - G_t)N_t h_t^A$). This equation has an analogous interpretation to equation (12), except that profits are augmented by the instantaneous expected gain from innovation $\eta G_t^{\widetilde{\kappa}}\left(N_t h_t^A\right)^{\kappa}\left(V_t^A - V_t^N\right)$ net of expenditure on automation research, $v_t h_t^A$. The first order condition for automation

production, as here, uses more than 2 factors. Blackorby and R. Russell (1981) discuss often used definitions and argue that the Morishima elasticity of substitution is the most appropriate.

innovation follows as:

$$\kappa G_t^{\widetilde{\kappa}} N_t^{\kappa} \left(h_t^A\right)^{\kappa-1} \left(V_t^A - V_t^N\right) = v_t, \tag{14}$$

which must hold at all points in time for all non-automated firms. The mass of high-skill workers hired in automation increases with the difference in value between automated and non-automated firms, and as such is increasing in current and future low-skill wages—all else equal.

Since non-automated firms get automated at Poisson rate $\eta G_t^{\widetilde{\kappa}} \left(N_t h_t^A\right)^{\kappa}$, and since new firms are born non-automated, the share of automated firms obeys:

$$\dot{G}_t = \eta G_t^{\widetilde{\kappa}} \left(N_t h_t^A\right)^{\kappa} (1 - G_t) - G_t g_t^N, \tag{15}$$

where g_t^N denotes the growth rate of N_t, the number of products.

Free-entry in horizontal innovation guarantees that the value of creating a new firm cannot be greater than its opportunity cost:

$$\gamma N_t V_t^N \leq v_t, \tag{16}$$

with equality whenever there is strictly positive horizontal innovation ($\dot{N}_t > 0$).

Finally, the low-skill and high-skill representative households' problems are standard and lead to Euler equations which in combination give

$$\frac{\dot{C}_t}{C_t} = \frac{1}{\theta} \left(r_t - \rho\right), \tag{17}$$

with a transversality condition requiring that the present value of all time-t assets in the economy—the aggregate value of all firms at time t—is asymptotically zero:

$$lim_{t\to\infty} \left(exp\left(-\int_0^t r_s ds\right) N_t \left((1-G_t)V_t^N + G_t V_t^A\right)\right) = 0.$$

2.5 Equilibrium Characterization

We define a feasible allocation and an equilibrium as follows:

Definition 1. A feasible allocation is defined by time paths of stock of varieties and share of those that are automated, $[N_t, G_t]_{t=0}^{\infty}$, time paths of use of low-skill labor, high-skill labor, and machines in the production of intermediary inputs $[l_t(i), h_t(i), x_t(i)]_{i\in[0,N_t],t=0}^{\infty}$, time path of intermediary inputs production $[y_t(i)]_{i\in[0,N_t],t=0}^{\infty}$, time paths of high-skill workers engaged in automation $[h_t^A(i)]_{i\in[0,N_t],t=0}^{\infty}$, and in horizontal innovation $[H_t^D]_{t=0}^{\infty}$, time paths of final good production levels and consumption levels $[Y_t, C_t]_{t=0}^{\infty}$ such that factor markets clear ((2) holds) and good market clears ((10) holds).

An equilibrium is a feasible allocation, a time path of intermediary input prices $[p_t(i)]_{i\in[0,N_t],t=0}^{\infty}$, a time path for low-skill wages, high-skill wages, interest rate and the value of non-automated and automated firms $\left[w_t, v_t, r_t, V_t^N, V_t^A\right]_{t=0}^{\infty}$ such that $[y_t(i)]_{i\in[0,N_t],t=0}^{\infty}$ maximizes final good producer profits, $[p_t(i), l_t(i), h_t(i), x_t(i)]_{i\in[0,N_t],t=0}^{\infty}$ maximize intermediary inputs producers' profits, $[h_t^A(i)]_{i\in[0,N_t],t=0}^{\infty}$ maximizes the value of non-automated firms, $[H_t^D]_{t=0}^{\infty}$ is determined by free entry, $[C_t]_{t=0}^{\infty}$ is consistent with consumer optimization and the transversality condition is satisfied.

In order to work with a system with an asymptotic steady-state; we introduce $n_t \equiv N_t^{-\beta/[(1-\beta)(1+\beta(\sigma-1))]}$ and $\omega_t \equiv w_t^{\beta(1-\sigma)}$ which both tend towards 0 as N_t and w_t tend towards infinity respectively, we define the normalized mass of high-skill workers in automation ($\hat{h}_t^A \equiv N_t h_t^A$), normalized high-skill wages ($\hat{v}_t = v_t N_t^{-\psi}$), where $\psi \equiv ((1-\beta)(\sigma-1))^{-1}$ (ψ can equivalently be shown to be equal to the long-run elasticity of GDP with respect to the number of products), and finally define the variable $\chi_t \equiv \hat{c}_t^{\,\theta}/\hat{v}_t$. When there is positive entry in the creation of new products at all points in time, this allows us to define a system of differential equations with two state variables n_t, G_t, two control variables, $\hat{h}_t^A, \chi_t$ and an auxiliary equation defining ω_t (see Appendix 8.1.1 for the derivation, in particular the system is given by equations (25), (26), (28) and (29)). We can then derive:

Proposition 1. *Assume that*

$$\rho\left(\frac{1}{\eta\kappa^{\kappa}(1-\kappa)^{1-\kappa}}\left(\frac{\rho}{\gamma}\right)^{1-\kappa}+\frac{1}{\gamma}\right)<\psi H, \tag{18}$$

then the system of differential equations admits a steady-state $(n^*, G^*, \hat{h}^{A*}, \chi^*)$ *with* $n^* = 0$, $G^* > 0$ *and positive growth* $\left(g^N\right)^* > 0$. *In such steady-state,* $G^* < 1$.

Proof. See Appendix 8.1.2. □

We will refer to the steady-state $(n^*, G^*, \hat{h}^{A*}, \chi^*)$ as as an asymptotic steady-state for our original system of differential equations. In addition, the assumption that $\theta \geq 1$ ensures that the transversality condition always holds. To see the intuition behind equation (18), consider the case in which the efficiency of the automation technology η

is arbitrarily large, such that the model is arbitrarily close to a Romer model where all firms are automated. Then equation (18) becomes $\rho/\gamma < \psi H$, which mirrors the classical condition for positive growth in a Romer model with linear innovation technology. With a small η the present value of a new product is reduced such that the corresponding condition is more stringent. For the rest of the paper we restrict attention to parameters such that there exists a unique saddle-path stable steady-state $(n^*, G^*, \hat{h}^{A*}, \chi^*)$ with $n^* = 0$, $G^* > 0$. Then, for an initial pair $(N_0, G_0) \in (0, \infty) \times [0, 1]$ sufficiently close to the asymptotic steady-state, the model features a unique equilibrium converging towards the asymptotic steady-state.[9] The proposition further stipulates that $G^* < 1$, which is derived from (26) with $\dot{G}_t = 0$ and $g_t^N > 0$.

3 The Three Phases of the Transition

This section analyzes the transitional dynamics from an initial starting point of (N_0, G_0) to the asymptotic steady state. We initially briefly discuss our baseline choice of parameters. Thereafter, we show that the transitional dynamics are best considered as consisting of three phases which we characterize through a combination of analytic and simulation methods.[10] Then, we analytically characterize the steady-state and finally, we consider other parameter choices.

Table 1: Baseline Parameter Specification

σ	ϵ	β	H	L	θ	η	κ	$\tilde{\varphi}$	ρ	$\widetilde{\kappa}$	γ
3	4	2/3	1/3	2/3	2	0.2	0.5	0.25	0.02	0	0.3

Table 1 presents our baseline parameters. Section 6 employs Bayesian techniques to estimate the parameters, but the focus of this section is theoretical and we simply choose 'reasonable' parameters. As our goal is to characterize the evolution of an economy which transitions from automation playing a small to a central role, we choose an initially low level of automation ($G_0 = 0.001$) and an initial mass of intermediate inputs small enough to ensure that the real wage is initially low relative to the productivity of machines. The

[9] Multiple asymptotic steady-states are technically possible but are not likely for reasonable parameter values (see Appendix 8.1.3). In addition, with two state variables (n_t and G_t) saddle path stability requires exactly two eigenvalues with positive real parts, in our numerical investigation, for all parameter combinations which satisfy the previous restrictions, this condition was always met.

[10] We employ the so-called "relaxation" algorithm for solving systems of discretized differential equations (Trimborn, Koch and Steger, 2008). See Appendix 8.2 for details.

characterization of the equilibrium in 3 phases is robust to considering other parameter sets with low G_0 as long as N_0 is sufficiently low to imply little initial incentive to automate. More generally, in the following, we will carefully specify which features of the equilibrium are specific to the set of parameters and which ones are more general. The time unit is 1 year and the other parameters of the model are given in Table 1. Total stock of labor is 1 and we set $L = 2/3$ and $\beta = 2/3$ such that absent automation and if all high-skill workers were in production the skill premium would be 1. The initial mass of products is $N_0 = 1$ and the productivity parameter for machines is $\tilde{\varphi} = 0.25$, which ensures that at $t = 0$, the cost advantage of automated firms is very small (their profits are 0.004% higher). We set $\sigma = 3$ to capture an initial labor share close to 2/3. The elasticity of substitution between machines and low-skill workers in automated firms is $\epsilon = 4$. The innovation parameters (γ, η, κ) are chosen such that GDP growth is close to 2% both initially and asymptotically, and we first consider the case where there is no externality from the share of automated products in the automation technology, $\widetilde{\kappa} = 0$—hereafter, we will refer to this externality as the externality in automation technology (although there is also an externality from the total mass of products). The parameters ρ and θ are chosen such that the interest rate is around 6% (at the beginning and at the end of the transition). For any variable a_t we let $g_t^a \equiv \dot{a}_t/a_t$ denote its growth rate and for future reference we let $g_\infty^a \equiv lim_{t\to\infty}\dot{a}_t/a_t$ denote its asymptotic limit (if such exists).

3.1 Phase 1: Almost Balanced Growth

Figure 1 plots the evolution the economy. We initially focus on the first 100 years of the transition which we denote 'Phase 1'. As we start with a low initial level of N_t, low-skill wages are low, and as shown in Panel C, the profits of an automated firm are only slightly higher than the profits of a non-automated firm (equation (6)). Non-automated firms invest very little in automation and G_t remains low (Panel C). The economy behaves essentially as if the aggregate production function were Cobb-Douglas: wages of both high- and low-skill workers grow at the rate of GDP (Panel A), and the labor share is constant (Panel D). Economic growth is (almost) entirely driven by the introduction of new products.[11]

[11]With a higher G_0 but still a low N_0, firms would still have had a low incentive to automate. As a result G_t would initially decline with the entry of new, non-automated products so that the transitional dynamics would quickly look similar to the present case.

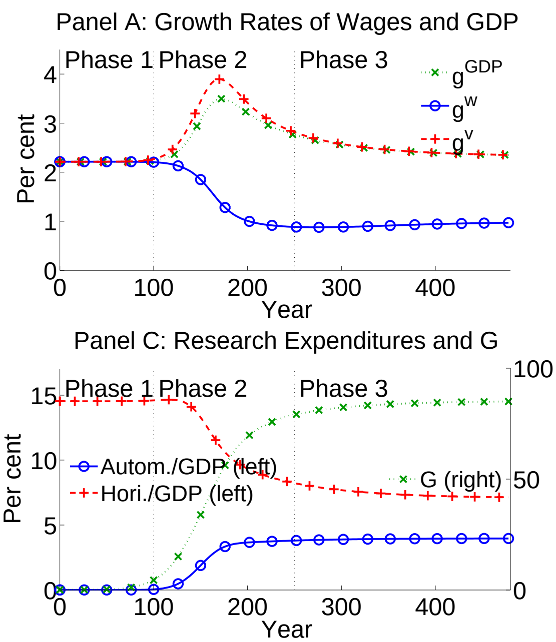

Figure 1: Transitional Dynamics for baseline parameters. Panel A shows yearly growth rates for GDP, low-skill wages and high-skill wages, Panel B the profit ratio of automated and non-automated firms and the relative pay of high-skill and low-skill wages, Panel C the total spending on horizontal innovation and automation as well as the share of products that are automated (G), and Panel D the wage share of GDP for total wages and low-skill wages.

To give further intuition, Figure 2a plots the skill-premium (8) and productivity (9) conditions in (w, v) space. This figure shows how wages depend on the technology levels, parametrized by the number of products N_t and the share of automated products G_t, and the mass of high-skill workers employed in production H_t^P. For G_t close to 0, equation (8) places the skill premium just above the straight line with slope $(1-\beta)L/(\beta H_t^P)$ (represented with a dotted line). During this phase, H_t^P remains nearly constant (as the ratio of research expenditures to GDP remains nearly constant). With nearly constant G_t and H_t^P, the skill premium condition barely moves. The increase in the number of products N_t pushes the productivity condition out, which increases low-skill and high-skill wages proportionally.

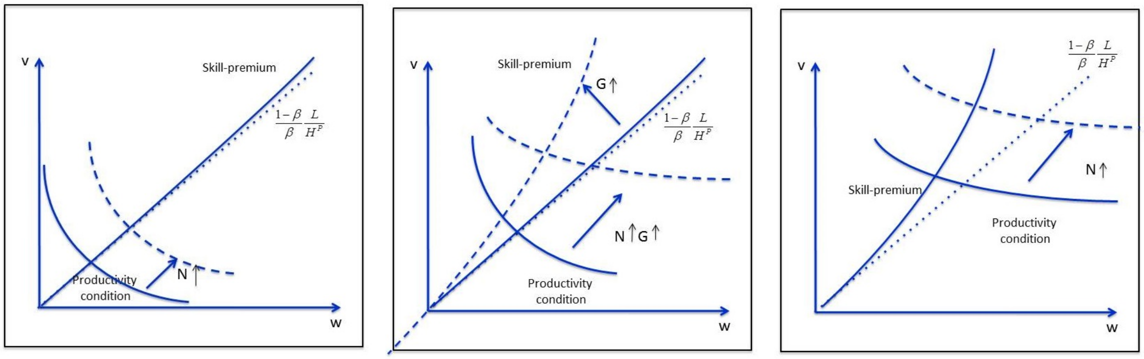

(a) Phase 1 (b) Phase 2 (c) Phase 3

Figure 2: Evolution of high-skill (v_t) and low-skill (w_t) wages . Wages are determined at the intersection of the skill-premium (equation 8) and productivity conditions (9) which move with the total stock of products, N_t, the share of automated firms, G_t and the mass of high-skill workers dedicated to production, H_t^P. Phase 1: The share of automated products G_t is close to zero and nearly constant. The stock of products, N_t, is increasing and pushes out the productivity condition causing low - and high-skill wages to grow at nearly identical rates. Phase 2: Automation increases the share of automated products and pivots the skill-premium condition counter-clockwise increasing the economy's ability to substitute low-skill workers with machines. Increases in both N_t and G_t push out the productivity-condition by improving the economy's productive capability. In general low-skill wages can decrease or increase. Phase 3: as the share of automated products approaches the asymptotic steady state, only the productivity condition is pushed out. Low-skill wages, w_t, grow at positive rate, though lower than that of high-skill wages, v_t.

3.2 Phase 2: Acceleration in Automation

As low-skill wages grow, the relative profitability of automated firms rise (Panel B in Figure 1) and the second phase of the transition is initiated around year 100. To facilitate exposition we describe the evolution of key variables sequentially.

Innovation. The immediate effect of higher relative profitability for automated firms, is an increase in spending on automation from an initial negligible level to around 4 per cent of GDP (Panel C). More precisely, since innovators are forward looking, it is the increase in the relative profitability of automated firms in the future which affects their incentive to automate.[12] In addition the share of spending on horizontal

[12]Note that the cost of automation, namely high-skill wages divided by the number of products, N_t, is also increasing over time. Yet, high-skill wages and aggregate profits grow at the same rate (since they are both proportional to output), therefore when the share of automated products, G_t, is low, high-skill wages divided by N_t and the profits of a non-automated firm grow at the same rate. This is why the increase in the cost of automation is dominated by the increase in its benefits, and therefore

innovation declines, particularly because new (non-automated) products will compete with increasingly productive automated firms and therefore get a smaller initial market share—the increase in automation spending at some point in Phase 2 is a general feature of the model but the decrease in horizontal innovation is not. The change in innovation spending directly increases the fraction of automated products, G_t (Panel C).

Labor income inequality. The increase in the share of automated products, G_t, changes the relative growth rates of low- and high-skill wages. As shown in Panel A in Figure 1, the growth rate of high-skill wages approaches 4%, while the growth rate of low-skill wages goes down to around 1% (since there are no financial constraints, the two types share a common consumption growth rate throughout, see Appendix 8.3.1).

Figure 2b shows how wages depend on the technology levels N_t and G_t for given mass of high-skill workers in production, H_t^P. Following Lemma 1, the skill premium curve bends up when G_t is different from 0. During Phase 2, N_t continues to increase, which pushes out the productivity condition. This increases both high-skill and low-skill wages, though the skill premium rises as higher low-skill wages increase the market share of automated firms, which rely relatively less on low-skill workers.

The increase in G_t has a positive effect on high-skill wages, but an ambiguous effect on low-skill wages. This is so, because an increase in the share of automated products has two opposing effects: i) an *aggregate productivity* effect as higher automation increases the productive capability of the economy and pushes out the productivity condition and ii) an *aggregate substitution* effect as it allows the economy to more easily substitute away from low-skill labor which pivots the skill-premium condition counter-clockwise. In the vocabulary of Acemoglu (2010), automation is low-skill labor saving whenever the aggregate substitution effect dominates the aggregate productivity effect. Which effect dominates here is generally ambiguous but when $\beta/(1-\beta) < \epsilon - 1$, that is when the elasticity of substitution between machines and low-skill workers is sufficiently high (the case for the parameters chosen here), w is decreasing in G for low N and 'inverse u'-shaped in G for a large N. Intuitively, when N and therefore w is low, the productive capabilities of the economy are not much improved by automation and the wage of low-skill workers is always decreasing in N. With higher N, the automation of the first products has a large productivity effect for the economy, while the substitution effect is relatively small since most firms are still non-automated; with the reverse being true for the automation of the last products. When $\beta/(1-\beta) < \epsilon - 1$, it further holds that a

firms start investing more in automation at the beginning of Phase 2.

fully automated economy will give low-skill workers lower wages than a completely non-automated one: $w|_{G=0} > w|_{G=1}$.[13] It is precisely this movement of the skill-premium curve that an alternative model with constant G (i.e. one where the fraction of tasks that can be performed with machines is constant) would not be able to reproduce, and consequently such a model would not feature labor-saving innovation. For this simulation, the increase in G_t always has a negative impact on low-skill wages, but it is sufficiently slow relative to the increase in the number of products that low-skill wages grow at a positive rate throughout. Importantly, this is not a general result. As shown below, there are parameters for which the growth rate of low-skill wages can be negative during Phase 2.

In addition to the effects of changing G_t and N_t, changes in the mass of high-skill workers in production, H_t^P, affect the skill premium. As high-skill labor is the only factor used in innovation, an increase in the mass of high-skill workers used in innovation increases the skill premium.[14] For our present simulation, H_t^P decreases slightly (when automation starts) and then increases later (when horizontal innovation declines), though this is not a general result. These effects on the skill premium are quantitatively dominated by the changes in use of machines in production.

Capital and labor shares. The second important feature is the progressive drop in the labor share of GDP (a feature robust to most parameter values). Profits are a constant share of output (because of the constant mark-up $1/\sigma$), but the increased use of intermediate inputs—which do not count towards GDP—implies a decreasing GDP/Y. Since in this model, capital income corresponds to profits, it is a growing share of GDP. Note that this happens even though machines are not part of a capital stock in this baseline version of the model (see section 5.3 for the alternative case). For the same reason, though the low-skill labor share drops rapidly, the high-skill labor share increases, such that the total labor share drops only slowly over the entire period. This is

[13] More formally, combining (8) and (9) gives the real wages w_t as a function of N_t, G_t and H_t^P. w_t is increasing in N_t and H_t^P. When $\beta/(1-\beta) > \epsilon - 1$: w is either increasing in G (for a low N) or 'inverse u'-shaped in G (for a large N) with $w|_{G=0} < w|_{G=1}$. When $\beta/(1-\beta) = \epsilon - 1$: w is 'inverse u'-shaped in G with $w|_{G=0} = w|_{G=1}$. When $\beta/(1-\beta) < \epsilon - 1$: w depends on G as described in the text. The relative strength of the aggregate productivity effect depends on the ratio $\beta/\left((1-\beta)(\epsilon-1)\right)$. A larger factor share β for machines/low-skill workers makes the aggregate scale effect stronger, whereas a larger elasticity of substitution, ϵ, makes the aggregate substitution effect stronger. Hence, a complete transition from all firms being non-automated to all firms being automated will benefit low-skill workers if and only if $\beta/(1-\beta) > \epsilon - 1$. See Appendix 8.1.4.

[14] Here a parallel can be drawn with the General Purpose Technology (GPT) literature discussed in the introduction, which argues that the arrival of a GPT temporarily increases the demand for high-skill workers.

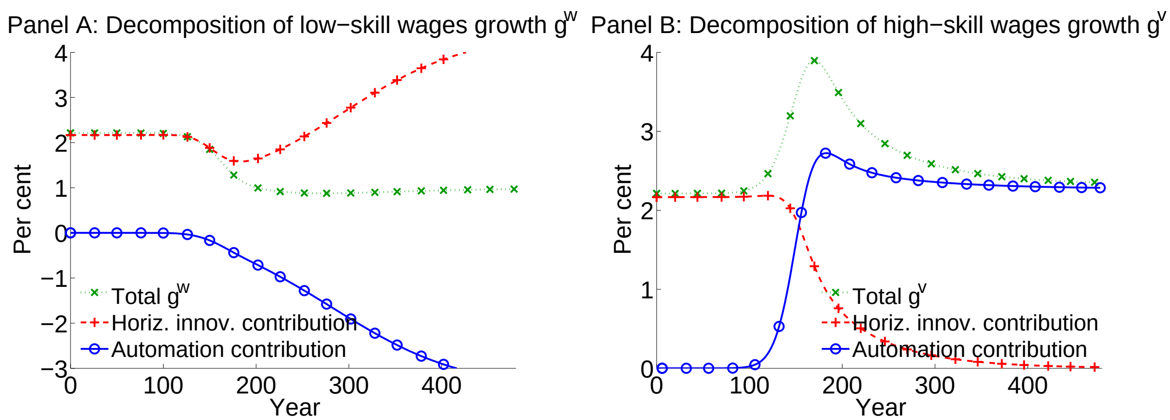

Figure 3: Growth decomposition. Panel A: The growth rate of low-skill wages and the instantaneous contribution from horizontal innovation and automation, respectively. Panel B is analogous for high-skill wages. See text for details.

consistent with recent evidence that has seen a drop in the labor share: Karabarbounis and Neiman (2013) find a global reduction of 5 percentage points in labor's share of corporate gross value added over the past 35 years. Rodriguez and Jayadev (2010) and Elsby, Hobijn and Sahin (2013), for the United States, find similar results. Consistent with recent trends (Piketty and Zucman, 2014, Piketty, 2014), the ratio of wealth to GDP increases since profits are an increasing share of GDP (see Appendix 8.3.1). This effect dominates a temporary increase in the interest rate.

As with the skill premium the total labor share is positively affected by increases in innovation as only high-skill workers work in innovation.

Growth decomposition. Figure 3 performs a growth decomposition exercise for low-skill and high-skill wages by computing separately the instantaneous contribution of each type of innovation. We do so by performing the following thought experiment: at a given instance t, for given allocation of factors, suppose that all innovation of a given type fails. By how much would the growth rates of w and v change this instance? This exercise is complementary to the one performed in Figure 2 which focuses on the impact of technological stocks instead of innovation.[15] In Phase 1, there is little automation, so the only source of growth for both skill-groups is horizontal innovation. In Phase

[15] More specifically, as mentioned in footnote 13, we can write $w_t = f(N_t, G_t, H_t^P)$, using equations

2, automation sets in and an increasing number of firms automate and substitute their low-skill labor with machines. From this point onwards, low-skill labor is continuously reallocated from existing products which get automated, to new, not yet automated, products. Consequently, the immediate impact of automation on low-skill wages is negative, while horizontal innovation has a positive impact, as it both increases the range of available products and decreases the share of automated products (since new products start non-automated). The figure also shows that automation plays a larger and larger role in explaining the growth rate of high-skill wages, while the contribution of horizontal innovation declines. This is because new products capture a smaller and smaller share of the market and therefore do not contribute much to the demand for high-skill labor. Consequently, automation is skill-biased while horizontal innovation is unskilled-biased. We stress that this growth decomposition is for changes in the rate of automation and horizontal innovation at a given point in time. This should not be interpreted as "automation being harmful" to low-skill workers in general. In fact, as we demonstrate in Section 3.5, an increase in the effectiveness of the automation technology, η, though it might have temporary negative impact on low-skill wages, will have positive long-term consequences.

Finally, a decomposition of g_t^{GDP} would look similar to the decomposition of g_t^v, such that as the economy grows, automation becomes an increasingly important source of growth. The increase in growth in Phase 2 is a result of us choosing parameters which imply an asymptotic growth rate around the initial growth rate and is not general. Had we chosen parameters for which asymptotic growth is slower than initial growth, the growth rate of Phase 2 would not necessarily have been much higher than that of Phase 1 (see Appendix 8.3.2 for such a case).[16]

(8) and (9). Differentiating with respect to time and using equation (26) gives:

$$g_t^w = \left(\frac{N_t}{w_t} \frac{\partial f}{\partial N} - \frac{G_t}{w_t} \frac{\partial f}{\partial G} \right) \gamma H_t^D + \frac{1}{w_t} \frac{\partial f}{\partial G} \eta G_t^{\tilde{\kappa}} (1 - G_t) (\hat{h}_t^A)^{\kappa} + \frac{1}{w_t} \frac{\partial f}{\partial H^P}.$$

Figure 3 plots the first two terms as the growth impact of expenses in horizontal innovation and automation, respectively. The third term ends up being negligible for our parameter choices. We perform a similar decomposition for v_t.

[16]There is an ongoing debate about the potential level of long-run growth. Jones (2002) argues that most of recent U.S. growth can be attributed to temporary factors such as a rise in educational attainment. The present model cannot quantitatively speak to potential long-run growth, but just shows that a phase of increased automation can act as an additional temporary factor spurring higher growth.

3.3 Phase 3: Towards the Asymptotic Steady-State

Finally, we discuss the period after year 250, during which the economy approaches its asymptotic steady state. Although the resources devoted to automation continue to increase, eventually the growth rate in G_t slows down and G_t asymptotes a constant, $G_\infty(= G^*$ the steady-state value), strictly below 1. The evolution of G_t results from the difference between two terms: the automation of not already automated products and the introduction of new non-automated products. As long as the automation intensity is bounded there will always be a share of products that are non-automated (see Lemma 2 below).[17]

The growth rates of GDP_t and high-skill wages, v_t, approach the same constant, and the labor share stabilizes at a lower level than that of Phases 1 and 2. Both high-skill workers and capital earn a higher share of GDP than in Phases 1 and 2, while the share going to low-skill workers asymptotes zero (Panel D in Figure 1). The wealth/GDP ratio, not drawn here, also stabilizes at a higher level than in Phases 1 and 2. We represent the evolution of the economy in (w, v) space in Figure 2c. With G_t (and H_t^P) roughly constant, the skill premium condition does not move, while horizontal innovation continues to push out the productivity condition. In a sharp contrast to Phase 2, low-skill wages cannot decrease, and instead grow at a positive nearly constant rate lower than that of high-skill wages (see Panel A). The skill premium grows unboundedly, though at a lower pace than in Phase 2 (Panel B). Further, note that there is no simple one-to-one link between automation spending and rising inequality. Here, automation spending is higher in Phase 3 than in Phase 2 (Panel C), yet the skill premium is increasing at a slower rate.

The properties of the asymptotic steady-state can be derived analytically and for a broader class of models than the baseline model studied here. We do so in the following subsection, which helps clarify the key assumptions behind our results.

3.4 Asymptotics for General Technological Processes

For this subsection, we consider any model where the equilibrium high-skill and low-skill wages satisfy equations (8) and (9). That is, our analysis depends on the "static" part of

[17]Formally, profits of an automated firm are asymptotically proportional to output Y_t divided by the mass of firms N_t. At the same time, wages of high-skill workers are asymptotically proportional to Y_t, so that the first-order condition for automation (14) implies that $N_t h_t^A$ asymptotes a constant. It then follows from (15) that $G_\infty < 1$.

the model (how v_t and w_t depend on the technological levels N_t, G_t and on H_t^P), but it does not rely on our particular specification for the evolution of N_t and G_t (for instance, it holds if the R&D input is the final good instead of high-skill workers, if some inputs are born automated or if some products become obsolete). The following proposition gives the asymptotic growth rates of w_t, GDP_t and v_t.

Proposition 2. *Consider three processes* $[N_t]_{t=0}^{\infty}$, $[G_t]_{t=0}^{\infty}$ *and* $[H_t^P]_{t=0}^{\infty}$ *where* $(N_t, G_t, H_t^P) \in (0,\infty) \times [0,1] \times (0,H]$ *for all* t. *Assume that* G_t, g_t^N *and* H_t^P *all admit strictly positive limits. Then, the growth rates of high-skill wages and output admit limits with:*

$$g_{\infty}^{v} = g_{\infty}^{GDP} = \frac{1}{(1-\beta)(\sigma-1)} g_{\infty}^{N}. \tag{19}$$

Part A) Furthermore, if $0 < G_{\infty} < 1$ *then the asymptotic growth rate of* w_t *exists and is positive at*

$$g_{\infty}^{w} = \frac{1}{1+\beta(\sigma-1)} g_{\infty}^{GDP}. \tag{20}$$

Part B). If $G_{\infty} = 1$ *and* G_t *converges sufficiently fast (more specifically if* $lim_{t\to\infty} (1-G_t) N_t^{\psi(1-\mu)\frac{\epsilon-1}{\epsilon}}$ *exists and is finite) then :*

- If $\epsilon < \infty$ *the asymptotic growth rate of* w_t *is positive at :*

$$g_{\infty}^{w} = \frac{1}{\epsilon} g_{\infty}^{GDP}, \tag{21}$$

where $1+\beta(\sigma-1) < \epsilon$ *by the assumption that* $\mu < 1$.[18]

- If low-skill workers and machines are perfect substitutes then $lim_{t\to\infty} w_t$ *is finite and weakly greater than* $\tilde{\varphi}^{-1}$ *(equal to* $\tilde{\varphi}^{-1}$ *when* $lim_{t\to\infty} (1-G_t) N_t^{\psi} = 0$*)*

Proof. see Appendix 8.1.5. □

This proposition does two things. First, it relates the growth rate of GDP (and high-skill wages) to the growth rate of the number of products. Without automation GDP_t would be proportional to $N_t^{1/(\sigma-1)}$, as in a standard expanding-variety model: the higher the degree of substitutability between inputs the lower the gain in productivity from an increase in N_t. Here, the fact that machines, produced one-for-one with the final good, are a potential input creates an acceleration effect as the higher productivity also increases the supply of machines. Asymptotically, this effect is increasing in the

[18] If $lim_{t\to\infty} (1-G_t) N_t^{\psi(1-\mu)\frac{\epsilon-1}{\epsilon}} = \infty$ then $\frac{1}{\epsilon} g_{\infty}^{Y} \leq g_{\infty}^{w} \leq \frac{1}{1+\beta(\sigma-1)} g_{\infty}^{Y}$.

factor share of low-skill workers/machines, β, under the conditions of Proposition 2.[19] Note also that for a given growth rate of the number of products, the asymptotic growth rate of output is independent of the share of automated firms, as long as it is strictly positive.

Second, it shows that in the case of asymptotically positive growth in N_t, mild assumptions are sufficient to guarantee an asymptotic positive growth rate of w_t. To see why, first consider the case in which $G_\infty < 1$, which includes the baseline model studied prior to this subsection as stipulated in Proposition 1. Since automated and non-automated products are imperfect substitutes, then so are machines and low-skill workers at the aggregate level. With a growing stock of machines and a fixed supply of low-skill labor, the relative price of a worker (w_t) to a machine (p_t^x) must grow at a positive rate. Since machines are produced with the same technology as the consumption good, $p_t^x = p_t^C$, where p_t^C is the price of the consumption good (1 with our normalization), and the real wage $w_t = w_t/p_t^C = (w_t/p_t^x)(p_t^x/p_t^C)$ must also grow at a positive rate—a generalized version of Proposition 2 is presented in section 5 which allows for asymptotic (negative) growth in p_t^x/p_t^C and thereby potentially decreasing real wage for low-skill workers.

The relative market share of automated firms and their reliance on machines also increase, both of which ensure that low-skill wages grow at a lower rate than the economy (see Lemma 1). This contrasts our paper with most of the literature which features a balanced growth path and therefore does not have permanently increasing inequality. For instance, in Acemoglu (1998), low-skill and high-skill workers are imperfect substitutes in production. Yet, since the low-skill augmenting technology and the high-skill augmenting technology grow at the same rate asymptotically, the relative stocks of effective units of low-skill and high-skill labor is constant, leading to a constant relative wage.

For growing low-skill wages, a higher importance of low-skill workers (a higher β) or a higher substitutability between automated and non-automated products (a higher σ) imply a faster loss of competitiveness of the non-automated firms and a lower relative growth rate of low-skill wages. The asymptotic growth rate of w_t is independent of the elasticity of substitution between machines and low-skill workers, ϵ, as the income received by low-skill workers from automated firms becomes negligible relative to the income earned from non-automated firms (this results from our assumption that $\mu < 1$

[19] If all labor could be replaced at some point by machines, then β would effectively be 1. The economy would then reach a "world of plenty" in finite time. In reality, one may think that natural resources would then become the binding factor.

such that automation reduces labor demand in a given firm).

Now, consider the case of $G_\infty = 1$ such that asymptotically all products are automated and the convergence happens fast enough (relative to the growth in N_t) to satisfy the condition in Part B of Proposition 2. If there is less than perfect substitutability between workers and machines within automated firms ($\epsilon < \infty$), then an analogous argument-by-contradiction shows that low-skill wages must increase asymptotically, though the growth rate relative to that of the economy must now be lower as the automated firms more readily substitute workers for machines than the economy substitutes from non-automated to automated products. The relative growth of low-skill wages now depends on ϵ as the fraction of non-automated firms becomes zero. Only in the special case in which machines and low-skill workers are perfect substitutes in the production by automated firms and the share of automated firms is asymptotically 1 will there be economy-wide perfect elasticity of substitution between low-skill workers and machines. In this case, w_t cannot grow asymptotically, but will still be bounded below by $\tilde{\varphi}^{-1}$, since a lower wage would imply that no firm would use machines.

In general, the processes of N_t, G_t and H_t^P will depend on the rate at which new products are introduced, the extent to which they are initially automated, and the rate at which non-automated firms are automated. The following lemma derives condition under which $G_\infty < 1$, as in the baseline model, so that Part A of Proposition 2 applies and in the long-run the economy looks like Phase 3 in our baseline model.

Lemma 2. *Consider processes $[N_t]_{t=0}^{\infty}$, $[G_t]_{t=0}^{\infty}$ and $[H_t^P]_{t=0}^{\infty}$, such that g_t^N and H_t^P admit strictly positive limits. If i) the probability that a new product starts out non-automated is bounded below away from zero and ii) the intensity at which non-automated firms are automated is bounded above and below away from zero, then any limit of G_t must have $0 < G_\infty < 1$.*

Proof. See Appendix 8.1.6. □

Under the conditions of Lemma 2, the mass of new non-automated products is positive, so that the reallocation of low-skill workers to these products ensures that their real wage grows in the long-run. It is only in the special case of all new products starting out automated (or equivalently the intensity with which they are automated increases without bounds) that G_∞ may be 1. In all other cases, Part A of Proposition 2 governs the asymptotic properties of g_t^w.

Interestingly, the intuition given by the combination of Lemma 2 and Part A of Proposition 2 does not rely on our assumption that new products are born identical to older products. In a model where new products are born more productive, the growth rate of high-skill wages and low-skill wages will obey equations (19) and (20), as long as the intensity at which non-automated firms get automated is bounded and the economy grows at a positive but finite rate.

In return, since the crucial element in Phase 2 of the baseline model was the increase in G_t from a low level to a level close to the steady-state value, a model which obeys equations (8) and (9), and satisfy the conditions of Lemma 2, will also feature a period akin to Phase 2 as long as G_0 is initially low relative to the asymptotic value G_∞.

3.5 Sensitivity Analysis

We now revert back to our specific baseline model, and study different scenarios. This section is complemented by Appendix 8.3.4, which carries a systematic comparative statics exercise.

Declining low-skill wages. We now show that our model can accommodate declining low-skill wages in Phase 2. An easy way to generate this pattern is to introduce the externality in automation. Figure 4 shows the evolution of the economy when $\widetilde{\kappa} = 0.49$.[20] The economy continues to be characterized by three phases. Even though the profit ratio becomes significantly different from 1 starting at year 100, expenditures in automation research remain low until around year 180, implying that Phase 2 starts much later (not until year 200). With $\widetilde{\kappa} > 0$, the automation technology is unproductive when G_t is small. Although the benefits from becoming automated are large, the chance of becoming automated remains low, which reduces the amount of investment in automation. As a result, the skill premium remains constant for longer (see Panel C). Yet, when Phase 2 starts, it is much more intense, partly due to the sharp increase in the productivity of the automation technology (following the increase in G_t) and partly because low-skill wages are higher when this happens. The skill premium sharply increases and the labor share drops. Importantly, for this case, low-skill wages decrease for part of Phase 2. This occurs both because automation is more intense and horizontal innovation less. First, the increase in G_t is accelerated so that in Figure 2b, the skill-premium condition pivots counter-clockwise faster (the aggregate substitution effect). The same accelerated increase in G_t pushes out the productivity-condition (the aggregate scale effect), which

[20]We choose this value for $\widetilde{\kappa}$ instead of 0.5, because in that case there is no horizontal innovation for some type periods (that is (16) holds with a strict inequality). This is not an issue in principle but simulating this case would require a specific numerical approach.

explains the high growth rates for v_t and GDP_t. Following our discussion in section 3.2, the substitution effect dominates the scale effect once G_t is large enough resulting in a drop in w_t—accordingly, the drop in g_t^w is delayed compared to the increase in automation. Second, horizontal innovation drops considerably, both because new firms are less competitive than their automated counterparts, and because the high demand for high-skill workers for automation increases the cost of inventing a new product. As a result for a period of time, most of the outward movement of the productivity condition is due to increases in G_t. Yet, the decline in w_t lowers the profit ratio, which in return tends to lower automation. This reflects a general point of the paper: just as increases in w_t tend to encourage automation; so do reductions in w_t discourage the same automation and reduce pressure on low-skill wages. To summarize, a larger automation externality tends to postpone Phase 2 but also to make it more intense, which makes negative growth rates for low-skill wages more likely.

Importantly, $g_t^w < 0$ is also possible (though not common) without the externality ($\widetilde{\kappa} = 0$) for other parameter choices—see Appendix 8.3.3 for an example. For this result, it is crucial that automation expenses are an upfront investment instead of a cost to be paid every period. A model where automation expenses take the latter form will have a hard time generating decreasing low-skill wages without an automation externality, because lower low-skill wages would reduce the incentive to pay this cost and thereby prevent low-skill wages from dropping. With innovation as an upfront investment, however, the level of low-skill wages depends on the stock of knowledge, and affects the incentives to innovate, that is the flow of knowledge, making it easier to generate a decreasing path for low-skill wages.

Innovation parameters. Second, we study the impact of the automation technology parameter on wages. For comparison, we also look at the impact of the horizontal innovation parameter. Figure 5 considers the separate cases of a higher productivity for horizontal innovation $\gamma = 0.32$ (instead of 0.3) and of higher productivity for the automation technology η=0.4 (instead of 0.2). For each case, Panels A-C show the value of selected outcomes relative to the baseline case. A higher productivity for horizontal innovation implies that GDP grows faster than in the baseline (Panel A), and with it low-skill wages (Panel B) as well as the skill premium. In addition, faster growth in low-skill wages implies that Phase 2 starts sooner (Panel D), which explains why the skill premium jumps relative to the baseline case before increasing smoothly. G_t is lower than in the baseline case in the long-run, as there are more new (non-automated)

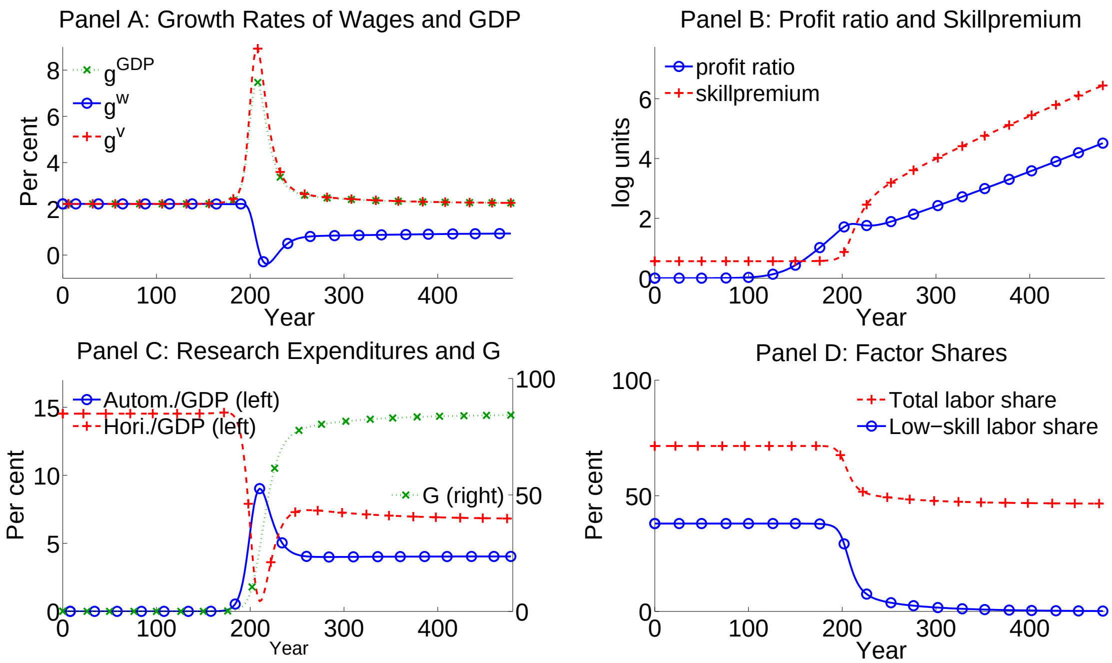

Figure 4: Transitional Dynamics. Note: Same as for figure 1 but with an automation externality of $\tilde{\kappa} = 0.49$

products at every instant. In fact, as shown in Appendix 8.1.7, it is possible to derive analytically that the asymptotic growth rate of the economy is increasing in γ and that the asymptotic share of automated products is decreasing in γ (when the steady-state is unique).

A higher productivity for the automation technology η initially has no impact during Phase 1, but it moves Phase 2 forward (see Panel D) as investing in automation technology starts being profitable for a lower level of low-skill wages. Since automation occurs sooner, the absolute level of low-skill wages drop relative to the baseline case (Panel B), which leads to a fast increase in the skill premium. A higher η also leads to a higher growth-rate asymptotically: it increases the value of a new firm (for a given innovation rate) since new firms are more likely to automate, which, in turn, leads to a faster rate of horizontal innovation (we prove this result analytically in Appendix 8.1.7 when the steady-state is unique). As in the case of a higher γ, a faster rate of horizontal innovation does imply that the skill premium keeps increasing relative to the baseline, but also that low-skill wages are eventually larger than in the baseline case. Therefore, a more productive automation technology only hurts low-skill workers temporarily, while they benefit in the long-run.

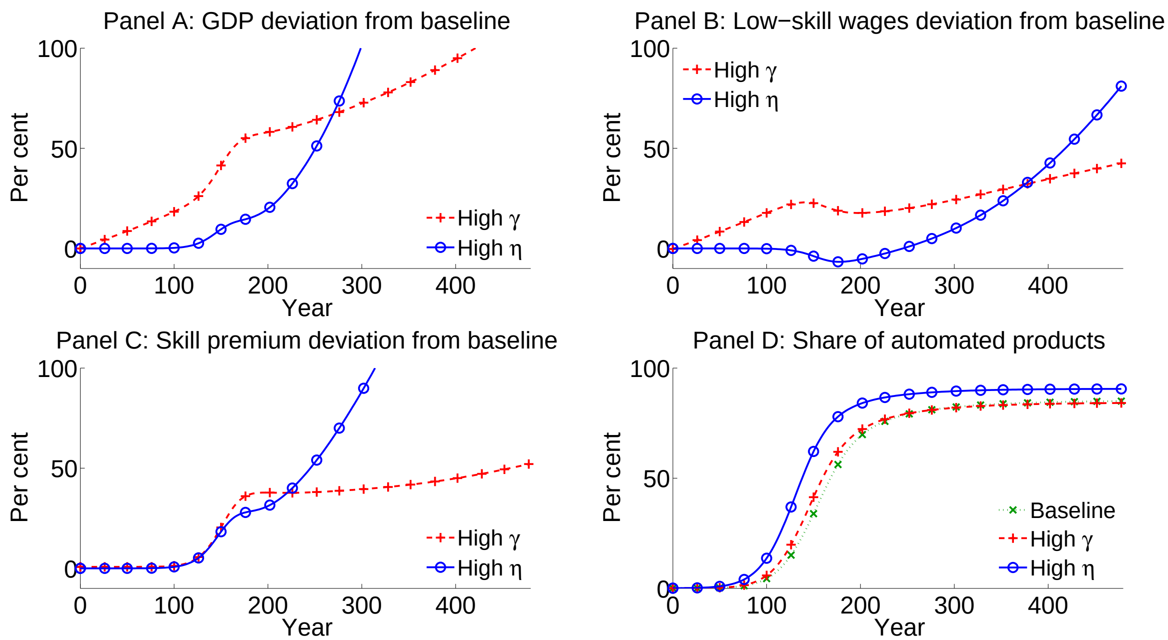

Figure 5: Transitional dynamics for more productive horizontal innovation technology (γ) and more productive automation technology (η).

A delayed decline in the labor share. Empirically, the drop in the labor share is a more recent phenomenon than the increase in the skill premium. In Figure 6, we choose parameters such that this happens. The automation technology is more productive $\eta = 0.4$; the automation technology is less concave $\kappa = 0.9$; and all other parameters are identical to the baseline case. In this case, more high-skill workers get allocated to automation during Phase 2: as shown in Panel C automation expenditures represent a much larger share of GDP during Phase 2 than they do in the baseline case. The mass of high-skill workers engaged in production declines during Phase 2. This results first in a sharper increase in the skill premium (the skill premium condition moves further to the left). In addition, the drop in the labor share is delayed since innovation spending are part of GDP (see equation (11)) while capital income is a constant share of output Y. The growth rates of low-skill and high-skill wages start diverging significantly from around year 135 and by year 150, the high-skill wage growth rate is 2pp higher than the low-skill wage growth rate, while the total labor share only start declining from around year 150 and in fact increases slightly before.

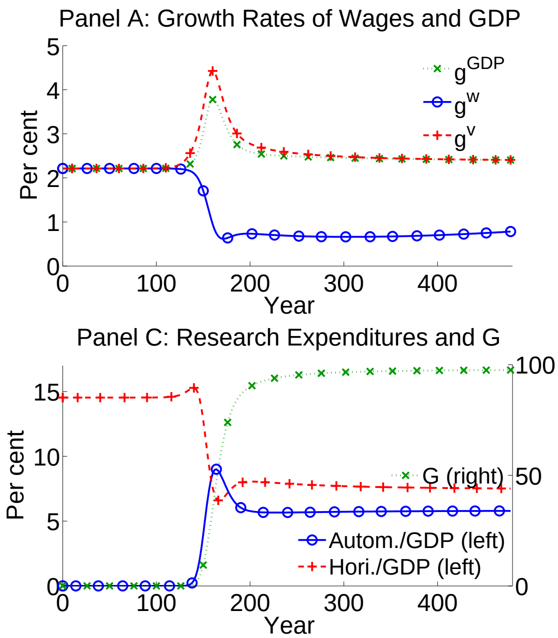

Figure 6: Transitional dynamics with a delayed drop in capital share

3.6 Social planner problem

Here, we briefly discuss the social planner problem associated with our model (Appendix 8.4 gives more details). The main lesson is that the social planner's solution looks qualitatively similar to the equilibrium we described, so that our results are not driven by the market structure we imposed. In particular, when the economy is initially endowed with low levels of technology N_0, G_0, the transitional dynamics still features the three phases described for the equilibrium case. There are four market imperfections that the social planner corrects: a monopoly distortion, a positive externality in horizontal innovation from the total number of products, a positive externality in the automation technology from the total number of products (the term N_t^{κ}) and a positive externality in the automation technology from the share of automated products when $\tilde{\kappa} > 0$ (which we referred so far as the "automation externality"). The optimal allocation can be decentralized using a subsidy to the use of intermediates inputs $1/\sigma$ (to correct for the monopoly distortion), a positive subsidy to horizontal innovation (to correct for the two externalities arising from the number of products), and, if $\tilde{\kappa} > 0$, a positive subsidy to automation (when $\tilde{\kappa} = 0$ there is no externality arising from automation and therefore the subsidy is null); all subsidies are financed with lump-sum taxes.

3.7 Discussion

Automation of new products. The evolution of the economy through the three phases does not depend on our assumption that new products are born non-automated. To emphasize this further, we present in Appendix 8.5 an alternative model where automation can only occur at the entry stage: when a firm is born, its owner can make it automated with probability $\min(\eta(N_t h_t^A)^\kappa, 1)$ by hiring h_t^A high-skill workers in automation. The economy goes through three phases as in the original model. As low-skill wages increase, the benefit from automation increases as well, and while initially most firms are born non-automated, overtime more and more are born automated, until this share stabilizes towards its asymptotic steady-state value. In this alternative model, new firms are more likely to be automated than older ones. Since automation does not take time, for some parameter values, all firms are eventually born automated, so that $G_\infty = 1$; in all cases, the asymptotic dynamics are still governed by Proposition 2.

Episodes of decreasing skill premium. Our baseline model implies an increasing skill-premium, which has not always been the case historically. An immediate explanation is a changing relative supply of skills. Goldin and Katz (2008) show that incorporating changes in the supply of skills into a model of skill-biased technical change captures well the evolution of the skill premium throughout the 20^{th} century. In some cases though, automation itself did not aim at replacing the most unskilled workers, as exemplified by the mechanization of the 19^{th} century, which replaced skilled artisans (including the Luddites), or the computerization of the last 30 years. The next section, which introduces a group of middle-skill workers, helps us account for such events.

Episodes of decreasing capital share and capital income ratio. Similarly, empirically, the capital share of income and the capital income ratio seem to have followed a U-curve in the 20^{th} century (Piketty and Zucman, 2014 and Piketty, 2014). A short decrease in the capital share is perfectly consistent with the transitional dynamics of our model but we cannot account for such large movements. Yet, the earlier decrease in the capital share and the capital income ratio was partly due to the two World Wars and to changes in the tax system. In addition, the transition away from the agricultural sector (and therefore the reduced importance of land) played a crucial role, which we have not modeled.

Structural shifts. The present model imposes that all products are equally substitutable with an elasticity greater than 1, implying that a firm that automates captures a larger market share. Historically, different sectors of the economy have experienced automation at different points in time. This could be captured with a nested structure with an elasticity of substitution between broad sectors of less than 1 . If these sectors differ in how easy it is to automate their intermediate inputs, then the phases of intense automation will happen sequentially. In such an economy as one broad sector experiences intense automation, spending shares in non-automated sectors would increase (as in Acemoglu and Guerrieri, 2008) securing a higher growth rate for low-skill wages. In addition, such an economy could replicate the broad features of an economy switching from agriculture, to manufacturing, and then services, and could generate interesting dynamics for the capital share.

4 Middle-Skill Workers and Wage Polarization

As mentioned in the introduction, a recent literature (including Autor, Katz and Kearney, 2006 and Autor and Dorn, 2013) argue that since the 1990s, wage polarization has taken place. That is, inequality has kept rising in the top half of the distribution, whereas it has narrowed for the lower half of the distribution. They conjecture that these "middle-skill"-workers are performing cognitive routine tasks which are the most easily automated. Our model suggests a related, but distinct explanation: Automating the tasks performed by middle-skill workers is not easier, but more difficult and therefore happens later. Hence, before 1990 and in fact for most of the 20^{th} century low-skill workers were in the process of being replaced by machines as semi-automated factories, mechanical farming, household appliances etc were increasingly used. Now computers are powerful enough to replace middle-skill workers and in the period since 1990, it has mostly been the tasks performed by middle-skill workers that have been replaced by computers, algorithms etc. In fact, Figure 3 in Autor and Dorn (2013) shows that low-skill workers left non-service occupations from the 70's, which is consistent with the view that their tasks in non-service occupations were automated before the middle-skill workers' tasks.

To make this precise, we introduce a mass M of middle-skill worker into the model. We think of these workers as being sequentially 'ranked' such that high-skill workers can perform all tasks, middle-skill workers can perform middle-skill tasks and low-skill tasks, and low-skill workers can perform only low-skill tasks. All newly introduced intermediate products continue to be non-automated, but there is an exogenous probability δ that they require low-skill and high-skill workers as described before, and a probability $1-\delta$ that they require both middle-skill and high-skill workers in an analogous. We refer to the former type of products as "low-skill products" and the latter type as "middle-skill products". This gives the following production functions (for $i \in [0, N_t]$):

$$y_L(i) = \left[l(i)^{\frac{\epsilon-1}{\epsilon}} + \alpha(i) \left(\tilde{\varphi}_L x(i) \right)^{\frac{\epsilon-1}{\epsilon}} \right]^{\frac{\epsilon\beta}{\epsilon-1}} h(i)^{1-\beta},$$

$$y_M(i) = \left[m(i)^{\frac{\epsilon-1}{\epsilon}} + \alpha(i) \left(\tilde{\varphi}_M x(i) \right)^{\frac{\epsilon-1}{\epsilon}} \right]^{\frac{\epsilon\beta}{\epsilon-1}} h(i)^{1-\beta},$$

where $y_L(i)$ and $y_M(i)$ are the production of low-skill and middle-skill products, respectively, and $m(i)$ is the use of middle-skill workers by a firm of the latter type. $\tilde{\varphi}_L$ and $\tilde{\varphi}_M$ are the productivity of machines that replace low-skill and middle-skill workers, respectively. The mass of low-skill products is δN, the mass of middle-skill products is $(1-\delta)N$ (an alternative model would allow all products to be produced by all factors; this would make the analysis substantially more complicated without altering the underlying argument). The final good continues to be produced competitively by a CES aggregator of all intermediate inputs and machines for either type of intermediate input producer are produced one-for-one with the final good keeping a constant price of 1. The shares of automated products, G_L and G_M will in general differ.

Both types of producers have access to an automation technology as before, but we allow the productivity to differ, such that automation happens with intensity $\eta_L G_L^{\widetilde{\kappa}} (N h_L^A)^{\kappa}$ for low-skill products and $\eta_M G_M^{\widetilde{\kappa}} (N h_M^A)^{\kappa}$ for middle-skill products. The equilibrium can be defined in a similar way as in section 2.5 and a proposition analogous to Proposition 1 exists.

To describe the equilibrium, we combine simulation methods and analytical results as in section 3. We want to analyze a situation where low-skill and middle-skill workers are symmetric except that middle-skill workers' tasks are more difficult to automate. To do this, we choose $\delta = 1/2$ and set $L = M = 0.25$ and keep parameters as before except that we choose $\widetilde{\varphi}_M = 0.15$ and $\tilde{\varphi}_L = 0.3$, so that machines are less productive in middle-skill products than in low-skill ones. The situation would be similar had we chosen $\widetilde{\varphi}_M = \tilde{\varphi}_L$, but $\eta_M < \eta_L$ such that the automation technology for middle-skill firms is less productive. Figure 7 describes the equilibrium for two different cases: without an externality in the automation technology ($\widetilde{\kappa} = 0$), and with a large externality ($\widetilde{\kappa} = 0.5$).

The overall picture is similar to that of Figure 1, but with distinct paths for low-skill

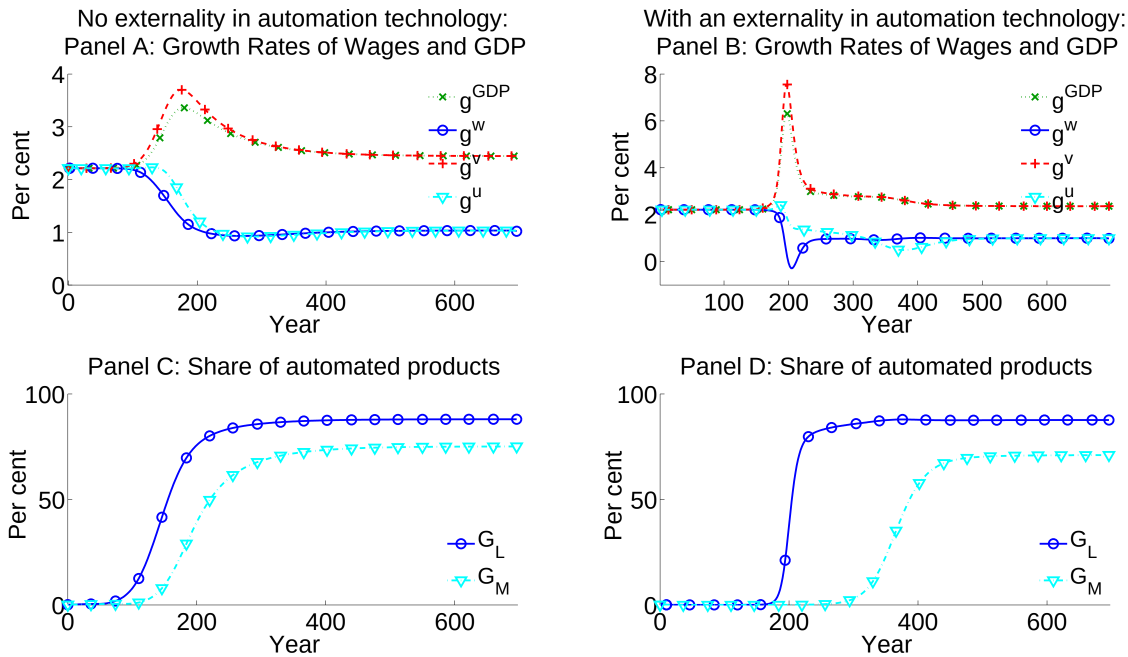

Figure 7: Transitional Dynamics with Middle-Skill Workers. Panel A and C show transitional dynamics for an economy with middle-skill workers but with no automation externality. Panel B and D include an automation externality of $\tilde{\kappa} = 0.5$

and middle-skill wages denoted w and u. The crucial difference is that since machines are less productive for middle-skill than for low-skill products, the intense automation phase occurs later for middle-skill products. One can now distinguish 4 phases. Phase 1 is analogous to Phase 1 in the previous case, and all wages grow at roughly the same rate. From around year 100 when $\widetilde{\kappa} = 0$, low-skill wages become sufficiently high, that low-skill product firms start investing in automation and G_L starts growing. Yet, since machines are less productive in middle-skill workers' tasks, G_M stays low until around year 150. During this second phase, inequality increases uniformly, high-skill wages grow faster than middle-skill wages which again grow faster than low-skill wages. Middle-skill wages do not grow as fast as GDP because automation in low-skill products increases their market share at the expense of the middle-skill products. From around year 150, the economy enters a third phase, where automation in middle-skill products is now intense, so that the growth rate of middle-skill wages drops further. Here it stays higher than the growth rate of low-skill wages until year 250, and is thereafter very slightly below, so that polarization occurs but barely.[21] The polarization phase is more salient with the automation externality, $\widetilde{\kappa} = 0.5$. Since the intense automation phase for middle-skill products occurs later, G_M increases faster which has a negative impact on middle-skill wages u_t relative to the previous case. In addition, G_L is closer to its steady-state level, which has a positive effect on g_t^w. Note, however, that for both cases, and at all points in time $v_t \geq u_t \geq w_t$, so no group has an incentive to be employed below its skill level.

Finally, in a fourth Phase (from around year 350 when there is no externality), G_L and G_M are close to their asymptotic steady-state levels and the economy approaches the asymptotic steady-state, with low-skill and high-skill wages growing positively but at a rate lower than that of the economy. Proposition 2 can be extended to this case. High-skill wages v_t and GDP_t all grow at the same rate which depends on the growth rate of the number of products:

$$g_\infty^v = g_\infty^{GDP} = \psi g_\infty^N.$$

Both low-skill and middle-skill wages grow asymptotically but at a lower rate given by:

$$g_\infty^w = g_\infty^u = \frac{1}{1 + \beta(\sigma - 1)} g_\infty^{GDP}.$$

Our assumption that automation is intrinsically easier for firms hiring low-skill workers than for those hiring middle-skill workers ($\widetilde{\varphi}_M < \widetilde{\varphi}_L$ or $\eta_M < \eta_L$) may seem at odds with empirical papers which argue that automation now predominantly hurts middle-skill workers. Our model emphasizes that the intensity of automation and the technological possibilities for automation are different concepts, since the intensity of automation does not depend only on its cost but also on its benefit. Hence, in the last phase of our simulation, middle-skill products get automated more intensively than low-skill products, even though automating low-skill products is less costly. Yet, some papers argue that the technological opportunities for automation are today lower for low-skill than for middle-skill workers. This is easy to reconcile with our model if we assume that for both low-skill and middle-skill products, a common fixed share can never be automated. After the start of Phase 2, the share of low-skill workers hired in products that can never be automated will be larger than the corresponding share for middle-skill workers (since

[21] Empirically, the polarization looks more like a J curve than a U curve as the difference in wage growth rates between the bottom and the middle of the income distribution is modest. This pattern is well-captured here where high-skill wages grow considerably faster than both low-skill and middle-skill wages from the time of initial automation.

a higher share of low-skill products will have been automated), and as a result, it will be on average easier to automate a middle-skill product than a low-skill one.

Naturally, the phase of intense automation of middle-skill products may occur sooner than that of low-skill products (for instance if the supply of middle-skill workers is low enough to generate a large middle-skill over low-skill wage ratio). This may be what happened in the 19^{th} century when the tasks of (middle-skilled) artisans got automated, as their wages were high relative to that of unskilled workers.[22]

5 Extensions

In this section, we extend our model to allow first, for an endogenous supply response in the skill distribution, second, for different technologies in the production of machines and the consumption good, and third for machines that do not depreciate immediately.

5.1 Supply Response in the Skill Distribution

We now allow the labor supply to respond to changes in factor rewards. We adopt a specification of heterogeneous productivity endowments of low-skill and high-skill workers. Let there be a unit mass of heterogeneous individuals, indexed by $j \in [0,1]$ each with 2-dimensional endowment. Each is endowed with the ability to supply $l\bar{H}$ unit of low-skill labor and $\Gamma(j) = \bar{H}\frac{(1+q)}{q}j^{1/q}$ units of high-skill labor (the exact distribution of high-skill abilities is of no crucial importance). Workers are thereby ranked in increasing order of their endowment of high-skill abilities on $[0, \bar{H}(1+q)/q]$. The parameter $q > 0$ governs the shape of the ability distribution with $q \to \infty$ implying equal distribution of skills. The following proposition generalizes Proposition 2.

Proposition 3. *Consider three processes $[N_t]_{t=0}^{\infty}$, $[G_t]_{t=0}^{\infty}$ and $[H_t^P]_{t=0}^{\infty}$ where $(N_t, G_t, H_t^P) \in (0,\infty) \times [0,1] \times (0,\bar{H}]$ for all t. Assume that G_t, g_t^N and H_t^P all admit strictly positive*

[22] These unskilled workers were often used to operate the machines which replaced the artisans. Though our model does not directly account for this since low-skill workers and middle-skill workers are hired in different firms, it would in the simple extension where all products use the three types of workers.

limits, with $G_\infty < 1$. Then the asymptotic growth rate of w_t and GDP_t are:[23]

$$g_\infty^w = \frac{1+q}{1+q+\beta(\sigma-1)} g_\infty^{GDP},\ g_\infty^v = g_\infty^{GDP} = \psi g_\infty^N. \tag{22}$$

Proof. See Appendix 8.6. □

At all points in time there exists an indifferent worker $(\bar{j}_t)$ where $w_t = (1+q)/q(\bar{j}_t)^{1/q} v_t$, with all $j \leq \bar{j}_t$ working as low-skill workers and all $j > \bar{j}_t$ working as high-skill workers. This introduces an endogenous supply response as the diverging wages for low- and high-skill workers encourage shifts from low-skill to high-skill jobs. The gradual reduction in supply dampens the relative decline in low-skill wages. Hence, besides securing themselves a higher future wage growth, low-skill workers who switch to high-skill occupation also benefit the remaining low-skill workers. Although, the supply effect is important in the low-skill sector, the increase in supply of high-skill workers is asymptotically irrelevant, and it continues to hold that $g_\infty^{GDP} = \psi g_\infty^N$ (where g_∞^N will depend on the asymptotic stock of high-skill workers $\bar{H} \equiv \int_0^1 \Gamma(j)dj$). In fact, the steady-state values for $(G^*, \hat{h}^{A*}, g^{N*})$ are identical to the case where the supply of high-skill workers is fixed at $\bar{H}$.

Note that as all changes in the stock of labor are driven by demand-side effects, wages and employment will move in the same direction. Extending our analysis of middle-skill workers to allow for switches between sector of employment would therefore reproduce the employment patterns of 'job polarization' in addition to 'wage polarization'.

5.2 Alternative Production Technology for Machines

The assumption of identical production technologies for consumption and machines imposes a constant real price of machines once they are introduced. As shown in Nordhaus (2007) the price of computing power has dropped dramatically over the past 50 years and the declining real price of computers/capital is central to the theories of Autor and Dorn (2013) and Karabarbounis and Neiman (2013). As explained in section 2, it is possible to interpret automation as a decline of the price of a specific equipment from infinity (the machine does not exist) to 1. Yet, our assumption that once a machine is invented, its price is constant, is crucial for deriving the general conditions under which

[23]If $G_t$tends towards 1 sufficiently fast such that $lim_{t\to\infty}(1-G_t)N_t^{\frac{\psi(1-\mu)(\epsilon-1)(1-q)}{q+\epsilon}}$ is finite then $g_\infty^w = \frac{1+q}{q+\epsilon} g_\infty^{GDP}$for ϵ finite and w_t is asymptotically finite (bounded below by $\widetilde{\varphi}^{-1}$) in the perfect substitute case.

the real wages of low-skill workers must increase asymptotically in Proposition 2. We generalize this in what follows.

Let there be two final good sectors, both perfectly competitive employing CES production technology with identical elasticity of substitution, σ. The output of sector 1, Y is used for consumption. The output of sector 2, X, is used solely for machines. The two final good sectors use distinct versions of the same set of intermediate inputs, where we denote the use of intermediate input for sector 1 as $y_1(i)$ and that in sector 2 as $y_2(i)$ with $i \in [0, N]$. The two versions of intermediate input i are produced by the same intermediate input supplier using production technologies that differ only in the weight on high-skill labor:

$$y_k(i) = \left[l_k(i)^{\frac{\epsilon-1}{\epsilon}} + \alpha(i)(\tilde{\varphi}x_k(i))^{\frac{\epsilon-1}{\epsilon}} \right]^{\frac{\epsilon\beta_k}{\epsilon-1}} h_k(i)^{1-\beta_k},$$

where a subscript, $k = 1, 2$, refers to the sector where the input is used. Importantly, we assume $\beta_2 \geq \beta_1$, such that the production of machines relies more heavily on machines as inputs than the production of the consumption good. Continuing to normalize the price of final good Y to 1, such that the real price of machines is p_t^x, and allowing for the natural extensions of market clearing conditions, we can derive the following generalization of Proposition 2 (where $\psi_k = (\sigma - 1)^{-1}(1 - \beta_k)^{-1}$).

Proposition 4. *Consider three processes $[N_t]_{t=0}^{\infty}$, $[G_t]_{t=0}^{\infty}$ and $[H_t^P]_{t=0}^{\infty}$ where $(N_t, G_t, H_t^P) \in (0, \infty) \times [0, 1] \times (0, H]$ for all t. Assume that G_t, g_t^N and H_t^P all admit strictly positive limits, then:*

$$g_{\infty}^{p^x} = -\psi_2 (\beta_2 - \beta_1) g_{\infty}^N$$
$$g_{\infty}^{GDP} = \left[\psi_1 + \psi_1 \frac{\beta_1 (\beta_2 - \beta_1)}{1 - \beta_2} \right] g_{\infty}^N, \tag{23}$$

and if $G_\infty < 1$ then the asymptotic growth rate of w_t is[24]

$$g_{\infty}^w = \frac{1}{1 + \beta_1(\sigma - 1)} \frac{1 - \beta_2 + \beta_1 (\beta_2 - \beta_1) \left(1 - \psi_1^{-1}\right)}{1 - \beta_2 + \beta_1 (\beta_2 - \beta_1)} g_{\infty}^{GDP}. \tag{24}$$

Proof. See Appendix 8.5. □

[24] If G_t tends towards 1 sufficiently fast such that $lim_{t\to\infty}(1 - G_t)N_t^{\psi_2(1-\mu_1)\frac{\epsilon-1}{\epsilon}}$ is finite, then $g_{\infty}^w = \frac{1}{\varepsilon}\left(1 - \frac{(\beta_2-\beta_1)(\varepsilon-1)}{(1-\beta_2+\beta_1)}\right) g_{\infty}^{GDP} \geq g_{\infty}^{p^x}$ whether ϵ is finite or not. It is clear that there always exists an ϵ sufficiently high for the real wage of low-skill workers to decline asymptotically.

Proposition 4 naturally reduces to Proposition 2 for the special case of $\beta_2 = \beta_1$. When $\beta_2 > \beta_1$, the productivity of machine production increases faster than that of the production of Y, implying a gradual decline in the real price of machines. For given g_∞^N, a faster growth in the supply of machines increases the (positive) growth in the relative price of low-skill workers compared with machines, w/p^x, but simultaneously, it reduces the real price of machines, p^x. The combination of these two effects always implies that low-skill workers capture a lower fraction of the growth in Y. Low-skill wages are more likely to fall asymptotically for higher values of the elasticity of substitution between varieties, σ, as this implies a more rapid substitution away from non-automated products.

5.3 Machines as a capital stock

We assumed so far that machines were an intermediate input that depreciates immediately. In practice, 'machines' often take the form of equipment capital, software, etc. which are durable (although their depreciation rate is typically higher than that of structures and housing). Appendix 8.8 derives the equilibrium under the alternative assumption that intermediate inputs producers rent machines from a capital stock. Capital increases with investment and depreciates at a fixed rate. The investment good is produced with the same technology as the consumption good.

Proposition 1 still holds—with the same sufficient condition (18)—but the system of differential equations must now involve three control variables and three state variables, Proposition 2 also carries through. Moreover, the asymptotic steady-state values for the growth rate of the number of products $\left(g^N\right)^*$—and therefore the growth rate of the economy, and the growth rate of low-skill wages—the share of automated products G^*, and the normalized mass of high-skill workers in automation $\widehat{h}^{A*}$ are the same as in the baseline case. The transitional dynamics look similar to the baseline case, but the central variable which determines whether automation is intensive or not, is now the ratio of low-skill wages to the gross rental rate of capital. In addition, since the expenditures on machines now correspond to capital income, the decline in the labor share tends to be more pronounced in this case, and high-skill workers' income need not become a larger share of GDP.

6 Quantitative Exercise

In this section, we conduct a quantitative exercise to compare empirical trends for the United States for the past 50 years with the predictions of our model using Bayesian techniques. We use the baseline model introduced in section 2, but we depart from the assumption of constant stock of each skill-type. As argued in Goldin and Katz (2008), during this period the relative supply of skilled workers increased dramatically and consequently, we allow workers to switch between skill-types as described in section 5.1.

Our model has predictions both for the labor market and the overall economy. In the following we match the skill-premium, the ratio of skilled to non-skilled workers, the growth rate of real GDP/employment, the share of labor in total GDP and we associate the use of machines with private equipment (real private non-residential equipment, 'Table 2.2. Chain-type Quantity Indexes' from NIPA). The values for the skill-premium and the ratio of skilled to non-skilled workers follows the methodology of Acemoglu and Autor (2011). The data for labor's share of GDP, the growth rate of real GDP/employment and the real stock of GDP is taken directly from the National Income and Products Accounts. All time series start in 1963 from when the skill-premium and skill-ratio are available and until 2007 to avoid the Financial Crisis. We match the accumulated growth rate of private equipment by indexing both X and real private equipment to 100 in 1963.[25]

6.1 Bayesian estimation

Due to the relatively small sample size we use Bayesian techniques to estimate our model, though little would change if we instead employed Maximum Likelihood procedures (in fact since we choose a uniform prior the maximum likelihood point estimate is equal to the mode of the Bayesian estimator). The model presented until now is not inherently stochastic, and in order to bring it to the data, we add auto-correlated measurement errors in the sense that we consider an economy where the underlying structure is described deterministically by our model above, but the econometrician only observes variables of

[25] The use of machines, X, has no natural units and we can therefore not match the level of X. Alternatively, we could normalize X by GDP, but we do no think of equipment as the direct empirical counter-part of X. First, equipment is a stock, whereas X is better thought of as a flow variable in our model. Second many aspects of automation might not be directly captured in equipment. Hence, equipment is better thought of as a proxy for X that grows in proportion to X. Empirically, equipment/GDP is about twice that of our predicted value of X/GDP.

interest with measurement error. Here, we briefly describe the estimation procedure and refer the reader to the appendix for a more complete treatment.

For time period $t = 1, ...T$, let $(Y_1^t, ..., Y_M^t) \in R^{M\times T}$ be a vector of M predicted variables with a time path of $Y_m^t = (Y_{m,s})_{s=1}^t$ for $m \in \{1, ..., M\}$ and $Y^t = (Y_1^t, ..., Y_M^t)$. Let the complete set of parameters in the deterministic model be $b_P \in \mathcal{B}_P \subset R^K$. We can then write the predicted values as a correspondence $\Omega : \mathcal{B}_P \rightarrow R^{M\times T}$ with:

$$Y^T = \Omega(b_P).$$

We add measurement errors, ϵ_m^T for $m \in \{1, ..., M\}$ to get observed predicted variables:

$$\hat{Y}_m^T = Y_m^T + \epsilon_m^T,\ m \in \{1, ..., M\},$$

where we assume normally distributed measurement errors $\epsilon_m^T \sim N(0, \Sigma_m)$. The errors are independent across types, $E[\epsilon_m^{T\prime}\epsilon_n^T] = 0$, for $m \neq n$, but potentially auto-correlated: the elements of Σ_m are such that the t, t'-element of Σ_m is given by $\sigma_m^2 \rho_m^{|t-t'|}$, where $\sigma_m^2 > 0$ and $-1 < \rho_m < 1$. Hence, σ_m^2 is the unconditional variance of a measurement error for variable m and ρ_m is its auto-correlation. This gives a total of $2M$ stochastic parameters and we label the combined set of these and b_P as $b \in \mathcal{B} \subset R^{2M+K}$.

This leads to a joint probability density for $\hat{Y}^T$ of $f(\hat{Y}^T|b) = \Pi_{m=1}^M f_m(\hat{Y}_m^T|b)$, with $\hat{Y}_m^T \sim N(Y_m^T, \Sigma_m)$. For a given prior on parameters $\pi(b)$, the Bayesian posterior distribution can then be written as:

$$f(b|\hat{Y}^T) = \frac{f(\hat{Y}^T|b)\pi(b)}{\int_{b'\in\mathcal{B}} f(\hat{Y}^T|b')d\pi(b')}.$$

With a full parametrization of the model the parameters are not uniquely identified and we restrict $\bar{H} = 1$ without loss of generality. Therefore, our deterministic model has 14 parameters including n_0 and G_0. Including parameters for the five measurement errors, this leaves us with 24 parameters. Table 2 gives our chosen prior distribution on parameters, all uniformly and independently distributed. The domain of the prior is deliberately kept wide for parameters not easily recovered from other studies such as the characteristics of the automation technology.Alternat

Table 2: Prior Distribution on Parameters: Uniformly and independently distributed on $[min, max]$.

	σ	μ	β	l	γ	$\tilde{\kappa}$	θ	η	κ	ρ	φ	q	n_0	G_0
min	3	1/2	0.1	0.5	0.15	0.01	1.5	0.2	0.01	0.015	5	0.2	0.1	0.1
max	6	0.99	0.95	1	0.5	0.99	2.5	0.3	0.99	0.07	15	0.99	10	0.8

6.2 Posterior Distribution and Mode

The posterior distribution, $f(b|\hat{Y}^T)$ cannot be evaluated analytically so we employ Monte Carlos simulation methods with the Geweke (1989) weighting procedure to get the mode and the posterior unconditional distributions for each parameter (details in Appendix 8.9). Table 3 shows the mode of the posterior distribution. The unconditional posterior distribution of each parameter is shown in Figure 22 in the appendix which demonstrates that variance for the posterior unconditional distribution is small.

Table 3: The mode of the posterior distribution.

	σ	μ	β	l	γ	$\tilde{\kappa}$	θ	η	κ	ρ	φ	q
Mode	4.27	0.623	0.764	0.881	0.200	0.225	2.10	0.250	0.680	0.060	9.85	0.83
	N_0	G_0	ρ_1	σ_1	ρ_2	σ_2	ρ_3	σ_3	ρ_4	σ_4	ρ_5	σ_5
Mode	1.89	0.56	0.94	0.99	0.91	0.09	0.94	0.011	0.74	0.0019	0.94	0.0012

Three parameter estimates are worth noting. First, the parameter of the automation externality, $\tilde{\kappa}$, is centered around 0.225 implying a substantial automation externality, a force for an accelerated Phase 2. Second, G_0 is centered around 0.56 implying that Phase 2 was already well underway in the early 1960s. Finally, the estimate of β—the factor share to machines/low-skill workers—is centered around 0.76 implying substantial room for automation.

Figures 8 and 9 further show the predicted path of the matched data series along with their empirical counterparts at the mode of the posterior distribution. Panel A of Figure 8 demonstrates that the model matches the rise in the skill-premium from the late 1970s onwards reasonably well, but misses the flat skill-premium in the period before. As argued in Goldin and Katz (2008), the flat skill-premium in this period is best understood as the consequence of a large increase in the stock of college-educated workers caused by other factors than an increase in the skill-premium (the Vietnam war and the increase in female college enrollment). Correspondingly, our model, which only allows relative supply to respond to relative factor rewards, fails to capture a substantial

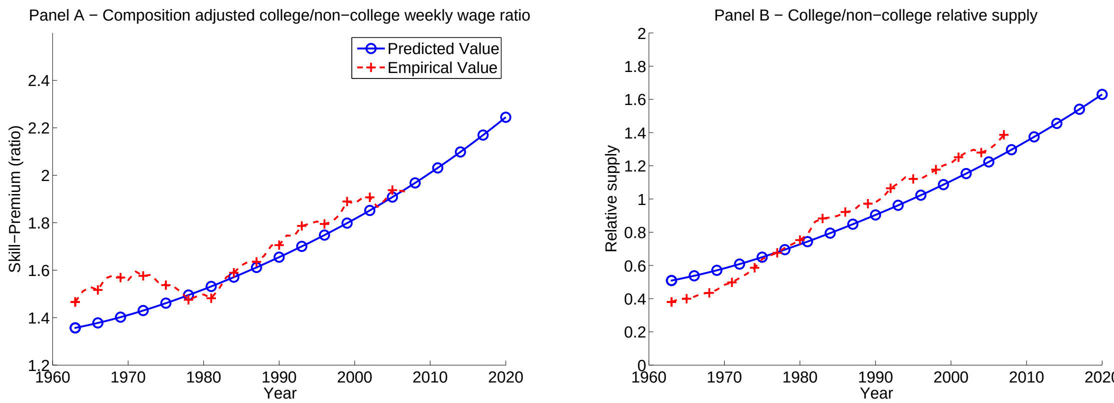

Figure 8: Predicted and Empirical time paths - Skill-premium and Skill-ratio

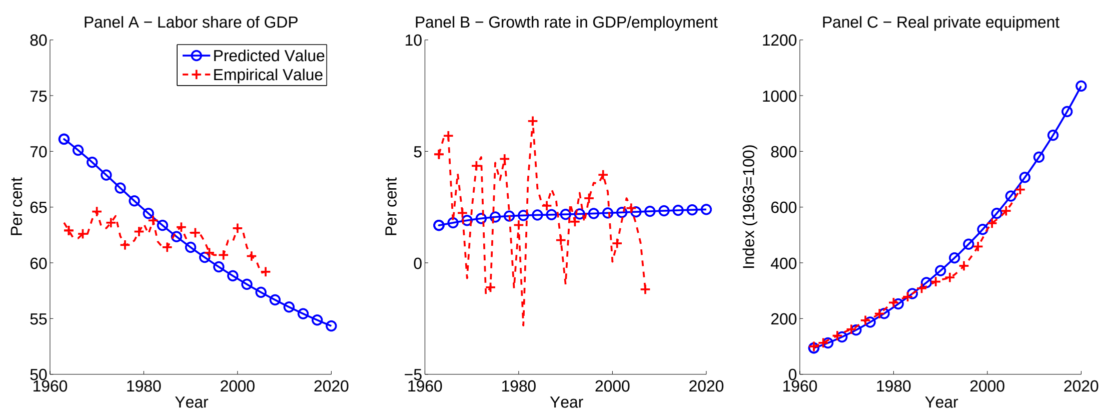

Figure 9: Predicted and Empirical time paths - Labor share of GDP, GDP growth and Private Equipment

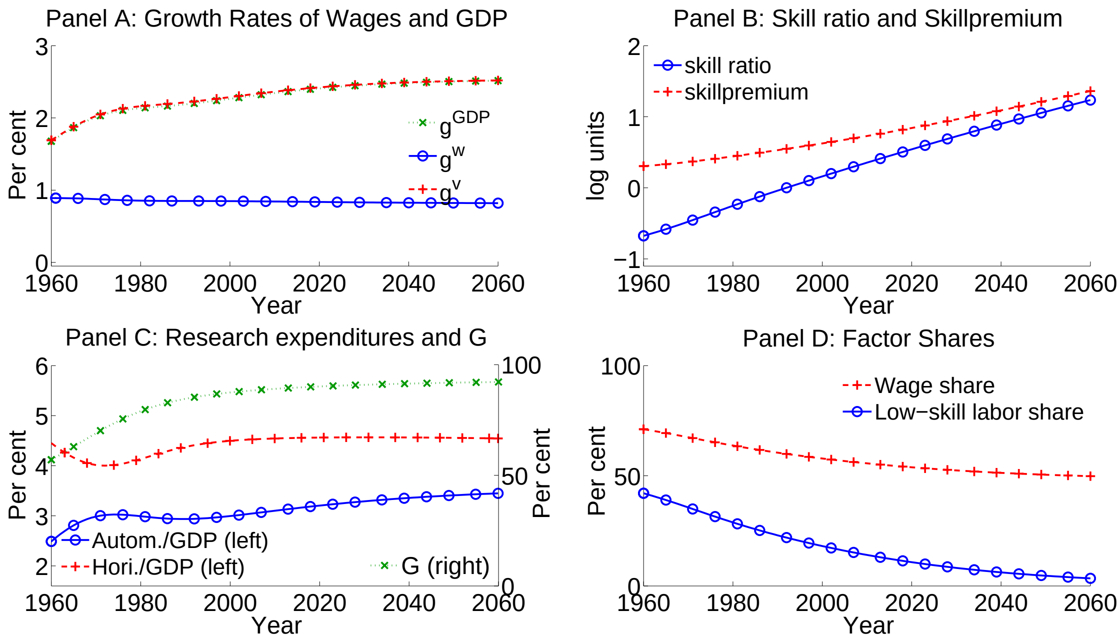

Figure 10: Transitional Dynamics with calibrated parameters

increase in the relative stock of skilled labor in the 1960s and early 1970s (Panel B of Figure 8). More importantly, the model predicts a substantially higher drop in the labor share of GDP than what has been observed empirically (11 versus 5 percentage points). The simple structure of the model forces any increase in the use of machines to be reflected in a drop in the labor share of GDP. A number of extensions would allow for more flexibility. For example, one could allow either for physical capital such as buildings or land, for an elasticity of substitution lower than 1 between high-skill labor and the low-skill - machines aggregate in production, or explicitly model several sectors as discussed in Section 3.7. As demonstrated in Figure 9 Panel B, the model matches the average growth rate of GDP/employment, but as a long-run growth model, is obviously not capable of matching the short-run fluctuations around trend. Panel C shows that the model captures the exponential growth in private equipment very well. Note that our exercise is qualitatively different from that of Krusell et al. (2000). Whereas they take time paths of factors of production as given (labor inputs, structures and equipment), we consider them to be endogenous and restricted to obey the structure of the model.

In Figure 10, we plot the transitional dynamics from 1960 to 2060. Panel A shows that low-skill wages are growing at half the rate of high-skill wages over the entire period.

Panel B shows that the skill ratio and the skill premium are predicted to keep growing at nearly constant rates. Panel C suggests that the share of automated products today is not far from its asymptotic steady-state value. Finally, Panel D shows that the labor share is predicted to stabilize at a slightly lower level than today.

7 Conclusion

In this paper, we introduced automation in a horizontal innovation growth model. We showed that in such a framework, the economy will undertake a structural break. After an initial phase with stable income inequality and stable factor shares, automation picks up. During this second phase, the skill premium increases, low-skill wages stagnate and possibly decline, the labor share drops—all consistent with the US experience in the last 50 years—and growth starts relying increasingly on automation. In a third phase, the share of automated products stabilizes, but the economy still features a constant shift of low-skill employment from recently automated firms to as of yet non-automated firms. With a constant and finite aggregate elasticity of substitution between low-skill workers and machines in the long-run, low-skill wages grow in the long-run as long as the machines production technology does not differ too much from the consumption good production technology. When the supply-side of the economy is allowed to respond, similar results carry through, but the rise in the skill premium is associated with a rise in the skill ratio which partly mitigates it. Wage polarization can be accounted for once the model is extended to include middle-skill workers.

In a critique of Piketty (2014), Ray (2014) formulates 'a Fourth Fundamental Law of Capitalism': "Uneven growth or not, there is invariably a long run tendency for technical progress to displace labor." Our framework distinguishes between production tasks that can potentially be automated and those that cannot—the ones done by high-skill workers in the baseline version. This allows us to relate the displacement of some workers by technology with the growing labor income inequality. Further, the model shows that there is indeed a long-run tendency for technical progress to displace labor, but this only occurs if the wages of the workers which can be substituted for are sufficiently large relative to the price of machines. This in turn can only happen under three scenarios: either automation must itself increase the wages of these workers (the scale effect dominates the substitution effect), or there is another source of technological progress (here, horizontal innovation), or technological progress allows a reduction in the price of machines relative to the consumption good. Finally, when machines are produced with a technology similar to the consumption good, automation can only reduce wages temporarily, as in Figure 4: a prolonged drop in wages would end the incentives to automate in the first place.

Fundamentally, the economy in our model undertakes an endogenous structural change when low-skill wages become sufficiently high. This distinguishes our paper from most of the literature, which seeks to explain changes in the distribution of income inequality through exogenous changes: an exogenous increase in the stock of equipment as per Krusell et al. (2000), a change in the relative supply of skills, as per Acemoglu (1998), or the arrival of a general purpose technology as in the associated literature. This makes our paper closer in spirit to the work of Buera and Kaboski (2012), who argue that the increase in income inequality is linked to the increase in the demand for high-skill intensive services, which results from non-homotheticities in consumption.

The present paper is a first step towards a better understanding of the links between automation, growth and income inequality. In future research, we will extend it to consider policy implications. The simple sensitivity analysis on the automation technology (section 3.5) suggests that capital taxation will have non-trivial implications in this context. Automation and technological development are also intrinsically linked to the international economy. Our framework could be used to study the recent phenomenon of 'reshoring', where US companies that had offshored their low-skill intensive activities to China, now start repatriating their production to the US after having automated to more advanced production technologies. Finally, our framework could also be used to study the impact of automation along the business cycle: Jaimovich and Siu (2012) suggests that the destruction of the 'routine' jobs happen during recessions, which raises the question of whether automation is responsible for the recent 'jobless recovery'.

References

Acemoglu, D. (1998). Why do new technologies complement skills? directed technical change and wage inequality. *Quarterly Journal of Economics*, 113(4):1055–1089.

Acemoglu, D. (2002a). Directed technical change. *Review of Economic Studies*, 69(4):871–809.

Acemoglu, D. (2002b). Technical change, inequality, and the labor market. *Journal of Economic Literature*, 40(1):7–72.

Acemoglu, D. (2007). Equilibrium bias of technology. *Econometrica*, 75(5):1371–1409.

Acemoglu, D. (2010). When does labor scarcity encourage innovation? *Journal of Political Economy*, 118(6):1037–1078.

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, volume 4, pages 1043–1171. Elsevier.

Acemoglu, D. and Guerrieri, V. (2008). Capital deepening and nonbalanced economic growth. *Journal of Political Economy*, 116(3):467–498.

Aghion, P. M. and Howitt, P. (1996). Research and development in the growth process. *Journal of Economic Growth*, 1(1):49–73.

Aghion, P. M. and Howitt, P. (1997). *Endogenous Growth Theory*. MIT Press.

Aghion, P. M., Howitt, P., and Violante, G. L. (2002). General purpose technology and wage inequality. *Journal of Economic Growth*, 7(4):315–345.

Autor, D. H. (2013). The "task approach" to labor markets: An overview. *Journal for Labour Market Research*, 46(3):185–199.

Autor, D. H. and Dorn, D. (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, D. H., Katz, L. F., and Kearney, M. S. (2006). The polarization of the u.s. labor market. *American Economic Review*, 96(2):189–194.

Autor, D. H., Katz, L. F., and Kearney, M. S. (2008). Trends in u.s. wage inequality: Revising the revisionists. *The Review of Economics and Statistics*, 90(2):300–323.

Autor, D. H., Katz, L. F., and Krueger, A. B. (1998). Computing inequality: Have computers changed the labor market? *Quarterly Journal of Economics*, 113(4):1169–1213.

Bárány, Z. and Siegel, C. (2014). Job polarization and structural change. In *Jahrestagung des Vereins für Socialpolitik 2014: Evidenzbasierte Wirtschaftspolitik*, number D14-V1.

Bartel, A. P., Ichniowski, C., and Shaw, K. (2007). How does information technology really affect productivity? plant-level comparisons of product innovation, process improvement and worker skills. *Quarterly Journal of Economics*, 122(4):1721–1758.

Blackorby, C. and R. Russel, R. (1981). The morishima elasticity of substitution; symmetry, constancy, separability, and its relationship to the hicks and allen elasticities. *Review of Economic Studies*, 48(1):147–158.

Buera, F. J. and Kaboski, J. P. (2012). The rise of the service economy. *American Economic Review*, 102(6):2540–2569.

Burstein, A., Morales, E., and Vogel, J. (2014). Accounting for changes in between-group inequality.

Caselli, F. (1999). Technological revolutions. *American Economic Review*, 89(1):78–102.

Costinot, A. and Vogel, J. (2010). Matching and inequality in the world economy. *Journal of Political Economy*, 118(4):747–786.

Elsby, M. W. L., Hobijn, B., and Sahin, A. (2013). The decline of the u.s. labor share. *Brookings Papers on Economic Activity*, 47(2):1–63.

Galor, O. and Moav, O. (2000). Ability-biased technological transition, wage inequality, and economic growth. *Quarterly Journal of Economics*, 115(2):469–497.

Geweke, J. (1989). Bayesian inference in econometric models using monte carlo integration. *Econometrica*, 57(6):1317–1339.

Goldin, C. and Katz, L. F. (2008). *The Race between Educaiton and Technology*. Harvard University Press.

Goos, M. and Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in britain. *Review of Economics and Statistics*, 89(1):118–133.

Goos, M., Manning, A., and Salomons, A. (2009). Job polarization in europe. *American Economic Review*, 99(2):58–63.

Hornstein, A., Krusell, P., and Violante, G. L. (2005). The effects of technical change on labor market inequalities. In Aghion, P. M. and Durlauf, S., editors, *Handbook of Economic Growth*, volume 1, chapter 20, pages 1275–1370. 1 edition.

Jaimovich, N. and Siu, H. E. (2012). The trend is the cycle: Job polarization and jobless recoveries. Working Paper 18334, National Bureau of Economic Research. NBER wp18334.

Jones, C. (2002). Sources of u.s. economic growth in a world of ideas. *American Economic Review*, 92(1):220–239.

Karabarbounis, L. and Neiman, B. (2013). The global decline of the labor share. *Quarterly Journal of Economics*, 129(1):61–103.

Katz, L. F. and Murphy, K. M. (1992). Changes in relative wages, 1963-1987: Supply and demand factors. *The Quarterly Journal of Economics*, 107(1):35–78.

Krusell, P., Ohanian, L. E., Ríos-Rull, J.-V., and Violante, G. L. (2000). Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica*, 68(5):1029–1053.

Lloyd-Ellis, H. (1999). Endogenous technological change and wage inequality. *American Economic Review*, 89(1):47–77.

Mazzolari, F. and Ragusa, G. (2013). Spill-overs from high-skill consumption to low-skill labor markets. *Review of Economics and Statistics*, 95(1):74–86.

Nelson, R. R. and Phelps, E. S. (1966). Investment in humans, technological diffusion, and economic growth. *The American Economic Review*, 56(1/2):69–75.

Nordhaus, W. D. (2007). Two centuries of productivity growth in computing. *Journal of Economic History*, 67(1):128–159.

Peretto, P. F. and Seater, J. J. (2013). Factor-eliminating technical change. *Journal of Monetary Economics*, 60(4):459–473.

Piketty, T. (2014). *Capital in the Twenty-First Century*. Harvard University Press.

Piketty, T. and Zucman, G. (2014). Capital is back: Wealth-income ratios in rich countries 1700-2010. *Quarterly Journal of Economics*, 129(3):1255–1310.

Ray, D. (2014). Nit-piketty: A comment on thomas piketty's capital in the twenty first century.

Rodriguez, F. and Jayadev, A. (2010). The declining labor share of income. Research Paper 2010/36, United Nations Development Programme.

Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy*, 98(5):71–S102.

Spitz-Oener, A. (2006). Technical change, job tasks, and rising educational demands: Looking outside the wage structure. *Journal of Labor Economics*, 24(2):235–270.

Trimborn, T., Koch, K.-J., and Steger, T. M. (2008). Multi-dimensional transitional dynamics: A simple numerical procedure. *Macroeconomic Dynamics*, 12(3):301–319.

Zeira, J. (1998). Workers, machines, and economic growth. *Quarterly Journal of Economics*, 113(4):1091–1117.

Zeira, J. (2013). Globalization and divergence. Working paper, The Hebrew University of Jerusalem.