

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jmoneco

Will the AI revolution cause a great divergence?☆

Cristian Alonso, Andrew Berg*, Siddharth Kothari, Chris Papageorgiou, Sidra Rehman

International Monetary Fund, United States



ARTICLE INFO

Article history:

Received 24 June 2020
 Revised 23 January 2022
 Accepted 26 January 2022
 Available online 4 February 2022

JEL classification:

E23
 O11
 O30
 O41

Keywords:

Automation
 Robots
 Divergence
 Development
 Technological change

ABSTRACT

Implications of a new wave of technological change that substitutes pervasively for labor are examined with particular focus on developing countries. While the model considered is minimalist by design, the resulting conclusions are powerful: improvements in the productivity of “robots” drive divergence, as advanced countries differentially benefit from their initially higher robot intensity, driven by their endogenously higher wages and stock of complementary traditional capital. Capital—if internationally mobile—is pulled “uphill”, resulting in a transitional GDP decline in the developing country. When robots substitute only for unskilled labor, the terms of trade, and hence GDP, may decline permanently.

© 2022 International Monetary Fund. Published by Elsevier B.V.

1. Introduction

A new wave of technological change involving artificial intelligence (AI) and machine learning, robotics, big data, and networks has led to renewed interest in the impact of pervasive automation on growth, wages, and inequality. While the economic landscape everywhere may shift radically, thus far the literature and the policy debate have focused almost exclusively on advanced economies. Narratives about the impact of a new round of automation on developing economies abound, both optimistic and pessimistic, but there has been very little systematic formal analysis.¹ This paper fills this gap by employing a simple but rich conceptual framework to explore the potential impact of the AI revolution on developing economies.

☆ We thank Aidar Abdychev, Emre Alper, Ed Buffie, Dominique Desruelle, Anton Korinek, Axel Schimmelpfennig, Preya Sharma, Felipe Zanna, and other colleagues at the IMF, and participants at the INET/IMF conference “Macroeconomics in the age of Artificial Intelligence” for valuable comments and suggestions. We acknowledge funding from the U.K. Foreign Commonwealth and Development Office (FCDO). The views expressed in this paper are the sole responsibility of the authors and should not be attributed to the International Monetary Fund, its Executive Board, or its management.

* Corresponding author:..

E-mail addresses: CALonso@imf.org (C. Alonso), ABerg@imf.org (A. Berg), SKothari@imf.org (S. Kothari), CPapageorgiou@imf.org (C. Papageorgiou), SRehman@imf.org (S. Rehman).

¹ The otherwise comprehensive set of contributions in [Agrawal et al. \(2019\)](#) contains no chapter focusing on international dimensions.

According to many—though by no means all—technologists, advances in AI and related technologies will allow machines to substitute for human labor across a much wider range of tasks than earlier waves of automation.² Ten or fifteen years ago it was widely accepted that machines could substitute for human labor in ‘routine’ tasks only, i.e. those typically middle-wage tasks involving neither on the one hand creativity and analytical skills nor on the other manual dexterity, and where the job could be explained step-by-step and hence programmed into a computer (Autor et al., 2003). Since then, however, advances in machine learning have led to machines with at least human-level perception and to AI programs with human or above-human-level capabilities in a broad range of tasks previously considered well out of reach.³

A burgeoning literature has begun analyzing the growth, labor market, and distributional implications for advanced countries. From an economic point of view, a key feature of this new wave of technologies is that it is likely to substitute more closely for labor, perhaps especially for unskilled labor. A general lesson is that automation that substitutes closely with workers will tend to increase incomes but also increase income inequality, at least during the transition and possibly in the long-run for some groups of workers.⁴ Our approach here follows Berg et al. (2018), who model the AI revolution as a reduction in the price of “robot capital”, which may substitute directly for labor and which complements traditional capital.⁵ Focusing on advanced economies, they find that the more easily robots substitute for workers, the higher the increase in GDP per capita and the greater the decrease in labor share, leading to a richer economy, but with more inequality. During a long transition, real wages may fall.⁶

Many observers argue that the current wave of automation will have significant effects on developing countries, in a literature that is largely qualitative and descriptive. Sachs (2019) and Yusuf (2017) suggest profound implications for development pathways and strategies, along with reductions in demand for unskilled labor. Nedelkoska and Quintini (2018) finds that developing countries are more vulnerable to automation, based on differences in industrial structure and, more importantly, in the way work is organized in these countries, notably a greater dependence on unskilled labor.⁷

This paper draws some simple but robust implications from a minimal set of assumptions about what this wave of technology may mean for countries at different levels of development. We employ a two-country one-sector neoclassical model with three factors of production: labor, capital, and “robots”, where “robots” are assumed to be close substitutes with labor. The AI or robot revolution is captured as an increase in the productivity of “robots”. We then extend the model to allow for two types of labor, with the developed country relatively well endowed in skilled labor. We assume that all differences between the advanced and the developing country derive just from differences in total factor productivity and/or the endowment of skilled labor, and that labor cannot move between countries. We can then examine the implications of an increase in robot productivity for inequality within and between each region, both in the long run and during the transition. Even this limited experiment turns out to make some powerful points about the likelihood of divergence arising from this wave of technology, as well as disentangling and clarifying many of the stories in the qualitative literature.⁸

We find grounds for concern in the form of three distinct channels through which developing economies could diverge further from advanced economies after the robot revolution: a share-in-production channel, a capital-flows channel, and a terms-of-trade channel. First, we find that just the addition of a highly-substitutable robot capital to the model, combined with high total factor productivity (TFP) in the advanced country, implies that an increase in productivity (or a reduction in cost) of robot capital results in further divergence of GDP levels between advanced and developing economies. In advanced economies, wages are higher because TFP is higher. These higher wages translate into more intense robot use in the advanced economy, resulting in higher robot shares in income. Higher robot shares in turn lead to much higher GDP growth in the advanced economy than in the developing economy when robot productivity increases.

There are also potential divergent transitional effects. While per capita GDP always increases in the long-run following an increase in robot productivity in a one-sector model, during the transition the robot revolution can reduce the level of per capita GDP in the developing economy through a capital-flows channel. The increase in productivity of robots induces a strong demand for additional resources in the advanced economy to finance investment in robots and in physical capital

² While in this paper we use the terms “AI”, “robots”, “automation”, and “technology” interchangeably, we also recognize that there is a challenge in specifying what each of these terms precisely mean. The discussions in Susskind (2019) and Grace et al. (2018) are perhaps closest to the spirit of this paper. The argument in the former and the forecast in the latter is that machine learning and related technologies are leading to a more-or-less gradual but pervasive growth in the scope of “AI”, broadly speaking, to substitute for human labor.

³ For general overviews, see Brynjolfsson and McAfee (2014), Ford (2015), and Susskind (2020) and, for a view that there is little new to see here, Mokyr et al. (2015) and Shiller (2019). As Susskind (2019) argues, the ability of various specific, i.e. “narrow”, AIs to substitute for human labor is more relevant to the labor market, and a much more immediate prospect, than the possible role of artificial general intelligence.

⁴ See e.g. Acemoglu (2021) and Benzell et al. (2021) which reinforce the unintended consequences of AI on inequality and also look for policy solutions.

⁵ “Robots” here and below, unless specifically noted, stands for the full-range of new technologies mentioned above.

⁶ The theoretical literature generally takes two approaches to modeling automation. An aggregate production function approach, similar to the model in this paper, is taken in Sachs (2018); Sachs and Kotlikoff (2012), Nordhaus (2015), Bessen (2017), Korinek and Stiglitz (2019), Caselli and Manning (2019), and Berg et al. (2018). In an influential series of papers, Acemoglu and Restrepo (2018a, 2018c, 2018d, 2019a) take a different approach and employ task-based models in the spirit of Zeira (1998) to examine the implications of task automation, and the creation of new tasks, for wages and output and the labor market. The substitution structure embedded in the production function in our model is also similar to the way Autor and Dorn (2013) model routine labor, abstract labor, and computer capital.

⁷ Rodrik (2016) discusses developmental implications of earlier waves of automation.

⁸ The model, of course, cannot directly speak to all the channels that are discussed in the qualitative literature. For example, the idea of leapfrogging is often raised as a possible benefit for developing countries, while the impact of automation on the viability of a manufacturing-led development strategy is viewed as a challenge. See Abdychev et al. (2018) and Bank (2019) for a broader discussion of these channels.

(which is assumed to be complementary to the robot capital as well as unskilled labor); as a result, capital flows “uphill” from developing countries to finance this capital accumulation.⁹

A third effect involves a potential reduction in the terms of trade for developing countries, leading potentially to a decline in the level of per capita GDP even in the long run. To capture this channel, we extend the model to allow for two goods/sectors and two types of labor, with one sector relatively intensive in unskilled labor. By assumption the developing country has relatively more unskilled labor, and therefore specializes in the sector intensive in that factor. We also assume—critically for this channel—that robots substitute more closely with unskilled than skilled labor.¹⁰

As in the one-sector model, an increase in robot productivity provides incentives to accumulate more robots and complementary physical capital—a direct effect which tends to increase income levels in the long run.¹¹ However, the two-sector model has a countervailing force acting through changes in relative prices. Because by assumption robots are strong substitutes for unskilled labor, an increase in robot productivity leads to a decline in demand for unskilled workers, thus reducing unskilled wages. This leads to a decline in the relative price of the good that uses unskilled labor more intensively, thus reducing the incentive to invest in the sector. As the developing region specializes in unskilled-intensive goods, and as the direct effect of an increase in robot productivity is small in the region given low robot shares, this lower price can result in income levels declining in developing countries in the long run.

The key assumption in this paper is that the current technological revolution is bringing a type of capital that is more substitutable with labor than previous rounds of technology. One line of evidence for this proposition can be found in the various technological studies that find, or more commonly predict, that AI and related technologies will be better than humans in many or most tasks within the foreseeable future. Frey and Osborne (2017) examines currently available technology and concludes that some 47 percent of jobs are subject to replacement by AI. Looking at prospective evolution of these rapidly-evolving technologies, AI researchers surveyed in Grace et al. (2018) on average expect to see AI outperforming humans at translating languages by 2024, driving a truck by 2027, and working in retail by 2031, with considerable variation around these estimates.¹² Empirical evidence that speaks directly to the sort of macro model we employ here is scarce but suggestive. Eden and Gaggl (2018) distinguish between traditional capital and the subset that embodies information and communication technology (ICT). Calibrating to data for 1950–2013 for the United States, they conclude that ICT capital is more highly substitutable with labor than traditional capital, with some evidence that it has been increasingly so over time.

The rest of the paper is organized as follows. Section 2 presents some stylized facts about automation and robot adoption in advanced and developing economies. It uses data on the distribution of industrial robots, meant literally, as a proxy for the more general concept of new nontraditional capital that we have in mind, simply because it is available in a comparable way for many countries. We establish that developing countries indeed are less robot-intensive, and that robot use is negatively correlated with wages. Section 3 presents our basic two-country one-sector model to study the impact of automation in the short and long run in advanced and developing countries. Section 4 extends the model to allow for two goods and two types of labor (skilled and unskilled). Section 5 discusses our main results and policy options to prevent or mitigate the negative impact of automation on developing economies.

2. The international distribution of robots

A key mechanism for the effects of automation on the international distribution of income in this paper is, as we will see below, that labor-substituting automation endogenously takes place more intensively in advanced countries, because wages are higher. In this section, we document this phenomenon, examining patterns in automation in the manufacturing sector across advanced and developing economies. We use data on the stock of robots from the International Federation of Robots (IFR).¹³ The IFR conducts annual surveys to collect data on the use of industrial robots in several countries and industries. As per the International Organization for Standardization (ISO), an industrial robot is defined as “an automatically controlled, reprogrammable, multipurpose manipulator programmable in three or more axes, which may be either fixed in place or mobile for use in industrial automation applications” (IFR, 2012).

While much of the debate on the impact of automation on labor market outcomes has been centered on advanced economies, robot adoption has been rising rapidly in developing countries too. Fig. 1 plots robot intensity (as measured by robots per worker in the manufacturing sector) for different country groups. The broad-based rise in robot intensity suggests

⁹ If, in contrast, the financial account is closed, then capital does not flow “uphill”; on the other hand the developing region loses the main potential benefit from the increases in robot productivity, which is the boost to long-run consumption that follows from the opportunity to accumulate claims on highly-productive robot capital in developed countries during the transition.

¹⁰ See Akerman et al. (2015) on broadband internet and Frey (2019) for a general discussion.

¹¹ Along somewhat similar lines, Eden and Gaggl (2019) argue that developing countries adopt less IT-intensive technologies because they are less-well-endowed in complementary skilled labor.

¹² This estimate is controversial. See the discussion in Susskind (2020), for example, for a review.

¹³ While we have data on robots, the broader concept of automation is a rapidly growing phenomenon, with the automation-related share of patents of physical and cognitive inventions having risen from only 23 percent in the late 1970s to 60 percent by 2014 (Mann and Püttmann, 2018). To be clear, the models in the paper consider a much broader range of technologies, including AI and machine learning algorithms and ever-faster related hardware, pervasive data and networks, and robotics per se. Indeed, one of the features of this new technological revolution, as argued for example in Susskind (2020), is that it extends well beyond manufacturing. However, data limitations force us to focus on this narrower concept in this section. Along similar lines, Acemoglu and Restrepo (2019b) uses similar data for the US to address more general questions.




Fig. 1. Adoption of robots in the manufacturing sector. Note: Data on number of robots from IFR. Data on manufacturing sector employment from various sources. See Section A in the Supplementary Material for details. HICS=High-Income Countries; UMICS=Upper Middle-Income Countries; LMICS=Lower Middle-Income Countries.




Fig. 2. Robot stock trends across income groups in the manufacturing sector. Note: Data on stock of robots from IFR. MICS=Middle-Income Countries.

that automation can have important ramifications for the developing world as well. In fact, robot adoption has been rising faster in developing countries than in advanced economies. In 2010, middle-income countries only accounted for about 5 percent of the total operational stock of robots. However, this number has grown dramatically to almost 24 percent by 2016. As shown in Fig. 2, while this is largely driven by a rapid pace of adoption in China, the trend is broad-based across other regions as well.

While robot use in low- and middle-income countries has increased significantly, robot density (number of robots per workers) still remains much higher in advanced economies. In fact, a high elasticity of substitution between robots and labor would imply that robot density varies more than proportionately compared to wages, all else equal (see equation 11 below). While it is beyond the scope of this paper to estimate the elasticity of substitution between robots and labor, here we combine data on real wages with the data on robots and employment to show that a positive correlation holds across countries. Section A in the Supplementary Material has details on data sources.¹⁴

Fig. 3 shows the relation between log of real wages and log of robots to employment ratio for a cross-section of 43 countries in 2010. Countries with higher wages have substantially higher robot intensity, with a slope of about 1.65, indicating that robot intensity varies more than proportionally compared to wages.¹⁵

¹⁴ All appendices appear under *Supplementary Material* available in the website of this journal.

¹⁵ Using the same robots data for advanced economies, [Graetz and Michaels \(2018\)](#) find that while increased robot use contributed positively to productivity growth and lowered prices, it only reduced low-skilled workers' employment share, leaving overall employment relatively unchanged. The slopes are similar if we use wages converted to USD using market and PPP exchange rates and when excluding commodity exporters.




Fig. 3. Real wages and robot density in the manufacturing sector. Note: Data on robots from IFR. Data on employment and wages from multiple sources. See Section A in the Supplementary Materials for details. Venezuela is excluded from the figure.

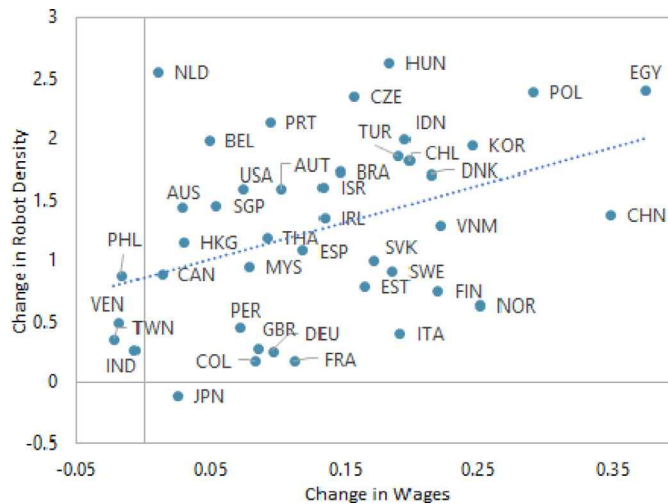


Fig. 4. Percent change in real wages and robot density in the manufacturing sector. Note: Data on robots from IFR. Data on employment and wages from multiple sources. See Section A in the Supplementary Material for details. Greece, Iran and Argentina are excluded from this figure.

Fig. 4 shows the relation between the change in log of wages and the change in log of robot density over the period for which data on robots is available.¹⁶ This evidence is suggestive of a positive relationship between real wages and robot density over time, with a slope greater than 1, plausibly in line with a relatively high elasticity of substitution between robots and labor.¹⁷

3. A two-region model with robots

In this section, we develop a two-region model to illustrate the possible impact of advances in AI and robotics on income gaps between countries. The model builds on Berg et al. (2018) by adding a developing region that has lower aggregate

¹⁶ Here, we disregard the first two years of reporting for the robot stock data for each country to account for a compliance bias. This bias can be inferred from the fact that some countries witnessed a very rapid increase in robot adoption in the first two years in which data is available, which may be a result of improved data reporting by the country rather than an actual increases in the underlying stock of robots. The slopes are similar if we use wages converted to USD using market and PPP exchange rates and when excluding commodity exporters.

¹⁷ Fig. 4 excludes Iran, Argentina and Greece, three outliers where real wages fell significantly, potentially due to sanctions in Iran's case and the severe economic crises in the others. Excluding these outliers increases the slope of the fitted line significantly to approximately 3. Fig. 1 in Section B in the Supplementary Material includes all countries in our sample. Here, the slope of the fitted line is positive, but only marginally greater than 1.

productivity than the advanced region.¹⁸ The two regions can trade with each other and financial assets can flow freely from one to the other. We begin with a model which features just one sector and one type of labor, with an extension to two types of labor discussed in Section 3.7. As wages are completely flexible in our model, lower labor demand following a robot revolution does not translate into unemployment, but into lower real wages.

This is the simplest model that illustrates the key channel that can lead to divergence in income levels between advanced and developing countries in response to an increase in robot productivity. In particular, when robots and labor are substitutable, higher wages result in robots being used more intensively in advanced economies in the initial steady state. An increase in robot productivity results in greater incentive to invest in robots (and complementary physical capital) in advanced economies where wages are high and robots are a more important component of the production process to begin with. This leads to higher GDP growth in advanced economies than in developing regions, thus leading to divergence.

3.1. Households

There are two regions indexed by i , representing an advanced economy ($i = A$) and a developing economy ($i = D$).¹⁹ Each region is populated by a household that lives forever and owns the three factors of production: labor (L), capital (K) and robots (Z). The household owns the firms operating the production technology and a financial asset, which allows it to borrow or save against the other region.

Household preferences are given by the utility function:

$$\sum_t \beta^t \frac{C_{i,t}^{1-\frac{1}{\tau}}}{1-\frac{1}{\tau}}$$

where $C_{i,t}$ is consumption of household i (with $i = A, D$) in period t , β is the discount factor, and τ determines the inter-temporal elasticity of substitution.

Household i seeks to maximize utility given its budget constraint:

$$C_{i,t} + I_{i,t}^K + I_{i,t}^Z + (B_{-i,t+1} - B_{i,t+1}) = r_t^K K_{i,t} + r_t^Z Z_{i,t} + w_{i,t} \bar{L}_i + (1 + r_t^B)(B_{-i,t} - B_{i,t}) + \Pi_{i,t} \quad (1)$$

where $I_{i,t}^K$ and $I_{i,t}^Z$ are investment in capital and robots, \bar{L}_i is the endowment of labor, and $(B_{-i,t} - B_{i,t})$ is the net financial asset holding for region i . Rates of return on capital, robots, and financial assets are given by r_t^K , r_t^Z , and r_t^B respectively. The wage rate is given by $w_{i,t}$. Finally, $\Pi_{i,t}$ represents the profits of firms operating the production technology in the country. The price of the final good is normalized to 1 in every period.

The laws of motion for accumulation of capital and robots are given by:²⁰

$$K_{i,t+1} = (1 - \delta_K)K_{i,t} + I_{i,t}^K \quad (2)$$

$$Z_{i,t+1} = (1 - \delta_Z)Z_{i,t} + I_{i,t}^Z \quad (3)$$

where δ_K and δ_Z are depreciation rates of capital and robot stock.

Maximizing utility subject to the budget constraint and laws of motion for capital and robots yields the standard Euler equation

$$\left[\frac{C_{i,t+1}}{C_{i,t}} \right]^{\frac{1}{\tau}} = \beta(1 + r_t^B) \quad (4)$$

and a no arbitrage condition that implies the equalization of net rates of return (after accounting for depreciation) for capital, robots and the financial asset²¹

$$r_t^K - \delta_K = r_t^Z - \delta_Z = r_t^B. \quad (5)$$

3.2. Firms

A representative firm operates the production technology in each region. Inputs are hired in competitive markets. Labor and robots are combined using a CES technology, with the composite then combined with capital using a Cobb-Douglas

¹⁸ An early version of the one-sector model of this paper appeared in [Abdychiev et al. \(2018\)](#)

¹⁹ We use region and economy interchangeably henceforth.

²⁰ The model is laid out and solved in discrete time. Derivations in continuous time using differentials would yield the same results.

²¹ Since there are no adjustment costs to capital or robots stocks and the net rental rate of capital, robots, and financial asset is the same, the holdings of these assets are indeterminate for the household. We assume that households hold all the capital and robots in the country, with the remaining wealth being held as a financial asset. Thus, capital and robots are not mobile across countries. However, this is not a restrictive assumption because households in one region can still invest in robots in the other region by lending resources through financial assets, which can in turn finance the capital and robot investment. In all cases, the return would be the same across assets.

function to obtain the final output. The production function is given by

$$Y_{i,t} = A_i (K_{i,t}^d)^\alpha \left[e^{\frac{1}{\sigma}} L_{i,t}^{\frac{\sigma-1}{\sigma}} + (1-e)^{\frac{1}{\sigma}} (b_t Z_{i,t}^d)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{(1-\alpha)\sigma}{\sigma-1}} \tag{6}$$

where $K_{i,t}^d$, $Z_{i,t}^d$, and $L_{i,t}$ are the quantity of capital, robots, and labor demanded by the firm. The level of total factor productivity (A_i) is the only parameter allowed to vary across regions. The cost-share parameters α and e , and the elasticity of substitution between labor and robots (σ) are assumed to be the same across the two regions.

The CES technology allows for flexibility, as different values of σ will correspond to different degrees of substitutability between labor and robots. Using this production function, the latest wave of technological innovation can be modeled as an increase in the productivity of robots, b_t .

Solving the profit maximization problem yields the standard first order conditions equating marginal products to factor prices. Dividing the first order condition for robots and capital with that for labor yields

$$\frac{Z_{i,t}^d}{L_{i,t}} = \frac{1-e}{e} \cdot b_t^{\sigma-1} \cdot \left(\frac{W_{i,t}}{r_t^Z} \right)^\sigma \tag{7}$$

$$\frac{L_{i,t}}{K_{i,t}^d} = \frac{1-\alpha}{\alpha} \cdot \frac{r_t^K}{w_t^i} \cdot \frac{1}{1 + \left(\frac{1-e}{e} \right)^{\frac{1}{\sigma}} \left(\frac{b_t Z_{i,t}^d}{L_{i,t}} \right)^{\frac{\sigma-1}{\sigma}}} \tag{8}$$

3.3. Equilibrium

A market equilibrium for this model is a set of prices and allocations such that:

1. Households choose consumption and holdings of robots, capital, and the financial asset to maximize utility given their budget constraint and the laws of motion for capital and robots.
2. Firms maximize profits by choosing the optimal combination of labor, capital, and robots.
3. All markets clear.

(a) Inputs Market. In each country i with $i = A, D$ and for every period t :

$$\text{(Labor)} \quad \bar{L}_i = L_{i,t} \quad \text{(Capital)} \quad K_{i,t} = K_{i,t}^d \quad \text{(Robots)} \quad Z_{i,t} = Z_{i,t}^d$$

(b) Good Market. For every period t :

$$\sum_i (C_{i,t} + I_{i,t}^K + I_{i,t}^Z) = \sum_i Y_{i,t} \tag{9}$$

(c) Financial Asset Market. For every period t :

$$B_{A,t} + B_{D,t} = 0 \tag{10}$$

3.4. Calibration of the initial steady states

We calibrate our model to study the differential response across regions of an increase in robot productivity for different levels of substitutability between labor and robots. We consider different values for σ , ranging from 1 (i.e., standard Cobb-Douglas production function with three factors of production) to 3 (i.e., robots substitute for labor). Rigorously estimating σ is beyond the scope of this paper. However, putting aside concerns regarding endogeneity and transition dynamics, the data shown in Fig. 3 and 4 suggest that σ is greater than one. The slope of the fitted line in these figures can be viewed as a rough estimate for σ , indicating that the elasticity of substitution may lie somewhere between 1.5 and 3. Edén and Gaggi (2018), using a somewhat different nesting than that employed here, conclude that the elasticity of substitution between labor and information-and-communication-technology (ICT) capital has increased rapidly since the late 90s, rising from 2.5 to 3.3. Calibrating their U.S. data to a production function similar to that employed here, Berg et al. (2018) find an elasticity between ICT capital and unskilled labor of 2.1.

For each value of σ , we calibrate a separate initial steady state where we choose e and A_D to match two moments: the relative GDP per capita between the two regions and the robot share in income in the advanced economy $\left(\frac{r^Z Z_A}{Y_A} \right)$. The relative GDP per capita is calibrated to be 15, which is the ratio of GDP per capita between the US and sub-Saharan African countries, as per the 5-year average from the Penn World Tables. The robot share in income in the advanced economy is calibrated to be 4 percent following Berg et al. (2018). The remaining parameters are standard values in the literature (Table 1).

Fig. 5 shows the resulting parameters for different values of the elasticity of substitution. The calibrated value of the labor share parameter (e) is higher for larger σ , so as to maintain a robot share in income of 4 percent in the advanced economy in the initial steady state.

Table 1
Calibration of one-sector model.

Parameter	Description	Value	Source/Reason
β	Discount factor	0.96	Long-run net return rate on capital around 4 percent
τ	Inter-temporal elasticity of substitution	1	Cobb-Douglas utility function
δ_K	Depreciation rate for capital	0.05	Berg et al. (2018)
δ_Z	Depreciation rate for robots	0.05	Berg et al. (2018)
α	Share of capital in production function	0.35	Berg et al. (2018)
\bar{L}_D	Stock of labor in the developing economy	1	Normalization
\bar{L}_A	Stock of labor in the advanced economy	1	Normalization
A_A	Total factor productivity in the advanced economy	1	Normalization
b_0	Initial robot productivity	0.1	Normalization
B_0	Initial asset holdings	0	Symmetry in the initial steady state




Fig. 5. Parameters Note: The figure plots the calibrated values of total factor productivity in the developing economy (A_D) and the labor share parameter (e) of the production function in the initial steady states. These parameters are chosen such that for each σ , GDP per capita in the advanced economy is 15 times that of the developing economy, and the robot share in income in the advanced economy is 4 percent.

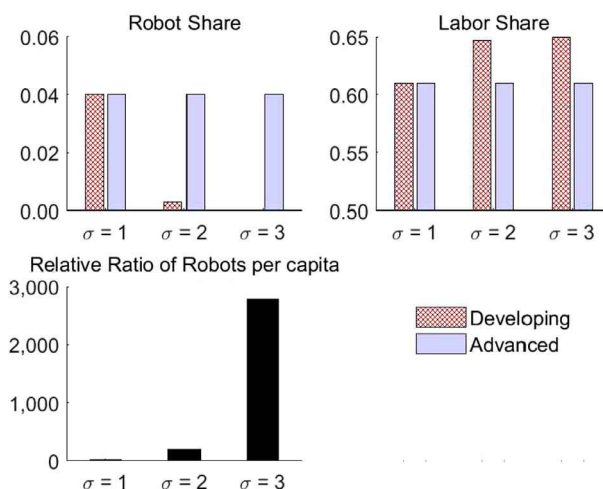


Fig. 6. Moments in the initial steady state. Note: The top panel of the figure plots the labor share ($\frac{wL}{Y}$) and the robot share ($\frac{rZ}{Y}$) for each steady state, that is, for each value of σ . The bottom left panel plots the robot intensity in the advanced economy relative to the developing economy ($\frac{Z_A/L_A}{Z_D/L_D}$).

However, the robot share in the developing region declines as σ increases in the initial steady state (Fig. 6). For the Cobb-Douglas case, with $\sigma = 1$, the robot share in the developing economy is the same as in the advanced economy. For higher values of σ , the robot share is lower in the developing region in the initial steady state because the higher elasticity of substitution magnifies the divergence in robot intensity (Z/L ratio) emerging from different wages. To get intuition, dividing equation 7 for the advanced economy with the same equation for the developing economy yields

$$\frac{Z_A/L_A}{Z_D/L_D} = \left(\frac{w_A}{w_D}\right)^\sigma \tag{11}$$

The higher wages in the advanced economy translate into higher robot-to-labor ratios. Furthermore, a higher elasticity of substitution amplifies the effect of wages on the robot-to-labor ratio. For a Cobb-Douglas production function, the robot-to-labor ratio between the two economies is proportional to the wage ratio, which results in the same labor share in the two economies. However, when robots and labor are substitutes ($\sigma > 1$), then the Z/L ratio in the advanced economy is more

than proportionally higher than the wage differential. The high elasticity of substitution implies that the high wages in the advanced economy result in greater substitution of labor for robots, leading to a higher robot share in output. Thus, when $\sigma > 1$, the robot share is lower in the developing economy than in the advanced economy and it declines with σ . Given that the capital share is the same across the two regions and does not vary with σ , the opposite is true for the labor share. The labor share is higher in the developing economy than in the advanced economy when σ is greater than 1, and it increases with σ .

3.5. Long-run impact of the robot revolution

3.5.1. Analytical results for a simplified one-region model

Before showing the quantitative results for the full model, we first derive some analytical results for the steady-state of a simplified one-region version of our model. This provides general intuition for our main mechanism. The one-region model is essentially identical to the model described above, but with no index i denoting region, the goods market clearing simplifying to $C_t + I_t^K + I_t^Z = Y_t$, and financial market clearing implying 0 holding of bonds.

The steady-state of this simplified model can be solved analytically. In steady-state, all interest rates are pinned down by the discount factor: $r^K - \delta_K = r^Z - \delta_Z = r^B = \beta^{-1} - 1$. The rest of the model can be simplified into solving for the 4 unknowns w , Z , K , and Q using steady state versions of Eqs. 6, 7, 8 and a zero profit condition given by

$$1 = \frac{(r^K)^\alpha \left(e(w)^{1-\sigma} + (1-e) \left(\frac{r^Z}{b} \right)^{1-\sigma} \right)^{\frac{1-\alpha}{1-\sigma}}}{A\alpha^\alpha (1-\alpha)^{1-\alpha}}$$

As shown in Section C in the Supplementary Material, log-linearizing these four equations around the initial steady-state, and imposing that deviations in interest rates will be 0 across steady-states, we can show that

$$\begin{aligned} \hat{w} &= \frac{\theta_z}{\theta_l} \hat{b} \\ \hat{z} &= \left(\sigma \frac{\theta_z + \theta_l}{\theta_l} - 1 \right) \hat{b} \\ \hat{k} &= \sigma \frac{\theta_z}{\theta_l} \hat{b} \\ \hat{q} &= \sigma \frac{\theta_z}{\theta_l} \hat{b} \end{aligned}$$

where \hat{w} , \hat{z} , \hat{k} , and \hat{q} are log deviation in wages, robot stock, capital stock, and GDP between the initial and final steady-state; θ_z and θ_l represent the robot and labor share in income in the initial steady-state; and \hat{b} represents a small change in robot productivity (again in log deviation from initial steady-state) which is how we model a robot revolution.

As these equations show, for a given change in robot productivity, the extent to which wages and GDP increase depends crucially on robot and the labor shares in the initial steady state. Two regions with the same robot and labor share will see the same increase in wages and GDP following an increase in robot productivity. On the other hand, if one region has a higher robot share, then it will see a bigger increase in wages and GDP.

Furthermore, Section 3.4 shows that when $\sigma > 1$, then the Z/L ratio in the advanced economy is more than proportionally higher than the wage differential. This endogenously leads to higher robot share (relative to labor share) in production in the advanced economy in the initial steady state, which would translate into a larger increase in wages and GDP in the advanced economy.

3.5.2. Quantitative results for the full two-region model

We now show quantitative results for our full two-region model. We model the robot revolution as a doubling of robot productivity (b) in both regions, as in Berg et al. (2018). The increase in robot productivity leads to higher GDP in both regions in the long run as households invest more in robots and in capital (which complements robots). However, which region grows more, and whether the developing economy falls further behind the advanced economy, depends crucially on the elasticity of substitution between robots and labor (Fig. 7).²²

In the Cobb-Douglas case, outcomes are symmetric in the two regions, with the change in GDP being the same, as both regions have the same robot share in output in the initial steady state.

However, when the elasticity of substitution is greater than 1, the developing economy benefits less. In this case, the robot share in output is larger in the advanced economy in the initial steady state (robots are a more important input in production), and so, the associated increase in GDP following a doubling of robot productivity is also larger. Investing in robots, and in complementary traditional capital, following an increase in robot productivity is most profitable where wages

²² For $\sigma = 3$, the increase in GDP per capita for the developing region in Fig. 7 is not noticeable due to the scale, but it is positive. This small, but positive, increase is apparent in Fig. 2 in Section D in the Supplementary Material.

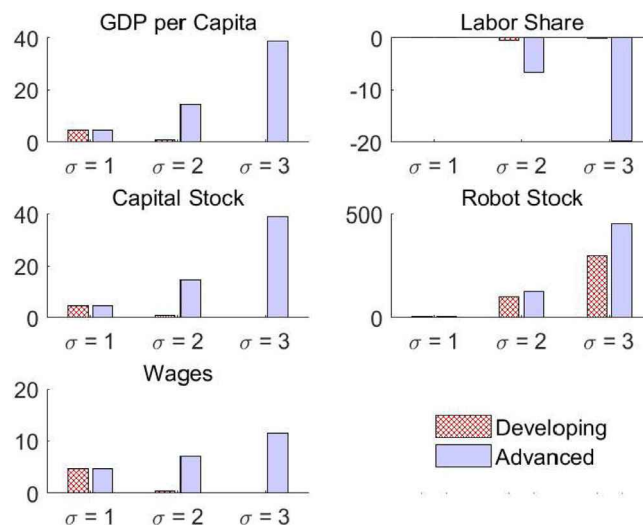


Fig. 7. Steady state comparison: percent changes with respect to initial steady state for different σ 's Note: The figure plots the percent change in various variables between the initial steady state and the final steady state following a doubling of robot productivity. Note that each σ (plotted on the x-axis) is associated with a different initial steady state calibration as described in Section 3.4.

are high because they save on the cost of employing expensive workers. Thus, the developing region falls further behind, diverging from the advanced region in GDP and (to a lesser extent) consumption.²³

While real wages increase in the long run in both regions, the change in labor share in output is more pronounced in the advanced economy. When robots easily replace workers, the robot and capital stocks increase by more than wages, leading to a fall in the labor share in both regions. The increase in real wages is stronger in the advanced region, but the increase in per capita GDP is even larger (due to faster robot and capital accumulation), so that the fall in labor share in output is more pronounced than in the developing region. Relatively higher growth in the advanced region is then associated with higher inequality as well.

Thus, the robot revolution may exacerbate income differences between advanced and developing economies if robots substitute workers because robots are used more intensively in advanced economies in the initial steady state. In Section D in the Supplementary Material, we show that this divergence result is robust to alternative calibrations of the TFP differential, so that substantial divergence would emerge even when considering a relatively-rich developing region such as China.

3.6. Short-run impact of the robot revolution

In this subsection, we explore the implications of the robot revolution during the transition. We assume that robot productivity doubles over four periods (with equal-sized increases each time) and remains constant thereafter.²⁴

While in the steady state we saw that GDP in the developing economy may fall back *relative* to the advanced economy when σ is greater than one, in the transition there may be a drop in the *absolute level* of GDP in the developing region as a result of the robot revolution. Fig. 8 shows the evolution of key variables along the transition path for different values of σ .

When the elasticity of substitution is 3, GDP falls in the developing region in the short-run. While the stock of robots increases in both regions, the capital stock in the developing economy falls for an extended period of time. This is because resources are channeled out of the developing region and transferred to the advanced economy to meet the stronger demand for capital and robots created there by the increase in robot productivity. The developing region acts as a lender, running a current account surplus and financing a stronger response to technological change in the advanced region. As a consequence, the divergence in terms of gross national product is less pronounced than the divergence in terms of GDP (Section F in the Supplementary Material).²⁵

Consumption also falls in both regions in the first few years as resources are diverted to investment, but the fall is larger and more persistent for the developing economy, again reflecting the outflow of resources and the build up of a net asset position with respect to the rest of the world. In the final steady state, consumption in the developing region is greater than output, the difference financed by the interest income on the accumulated assets.

The transition is particularly painful and long for workers in the developing region in terms of wages (Fig. 9). For an elasticity of substitution of 3, wages drop after the robot revolution in both regions as firms substitute away from workers

²³ As we show in the next section, with open capital accounts divergence in consumption is mitigated, because the developing region invests in some of the advanced-country robot stock during the transition.

²⁴ Section E in the Supplementary Material provides details on solving for the transition.

²⁵ Gross national product is defined as $wL + r^K K + r^Z Z + r^B (B_A - B_D)$.

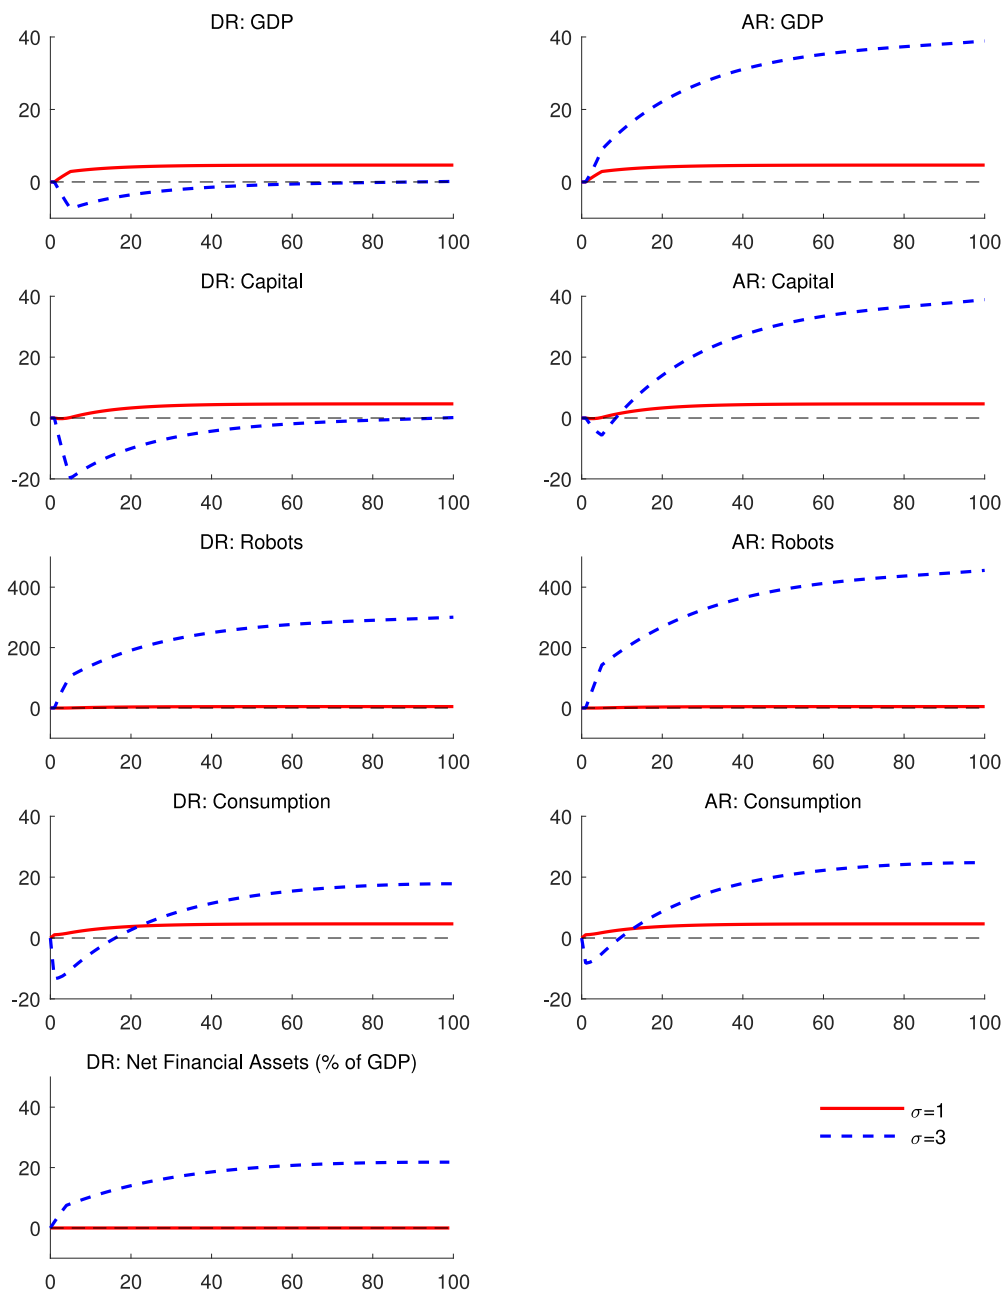


Fig. 8. Transition: GDP per capita, capital, consumption, and savings. Note: DR=Developing Region; AR=Advanced Region. All charts plot the transition path, showing the percent changes with respect to initial steady state, except “Net Financial Assets” which are shown as a percent of world GDP.

and into robots. However, the decline is larger in the developing region, where GDP is falling as capital flow out of the region. Eventually, the capital stock grows enough to compensate for the negative substitution effect and raises wages, but that takes longer in the developing region. In terms of labor income as a share of output, the decline is quite small for the developing region but quite large for the advanced economy. However, even in the developing economy, labor income as a share of gross national product falls substantially as the build up of foreign assets yields large interest income, indicating that inequality might increase in the developing economy as well.

Most of the transitional effects on the developing country are due to capital account implications of robot adoption in advanced countries (see Fig. 10). In the case of closed economies, the shape of the response is similar across regions but greatly amplified in the advanced region, where there are greater incentives to take advantage of the more productive robots. In the developing country, with low wages and low share of robots in output, the increase in demand for robots is smaller. This leads to a small increase in the interest rate in the developing country (barely noticeable in the plot), in contrast to

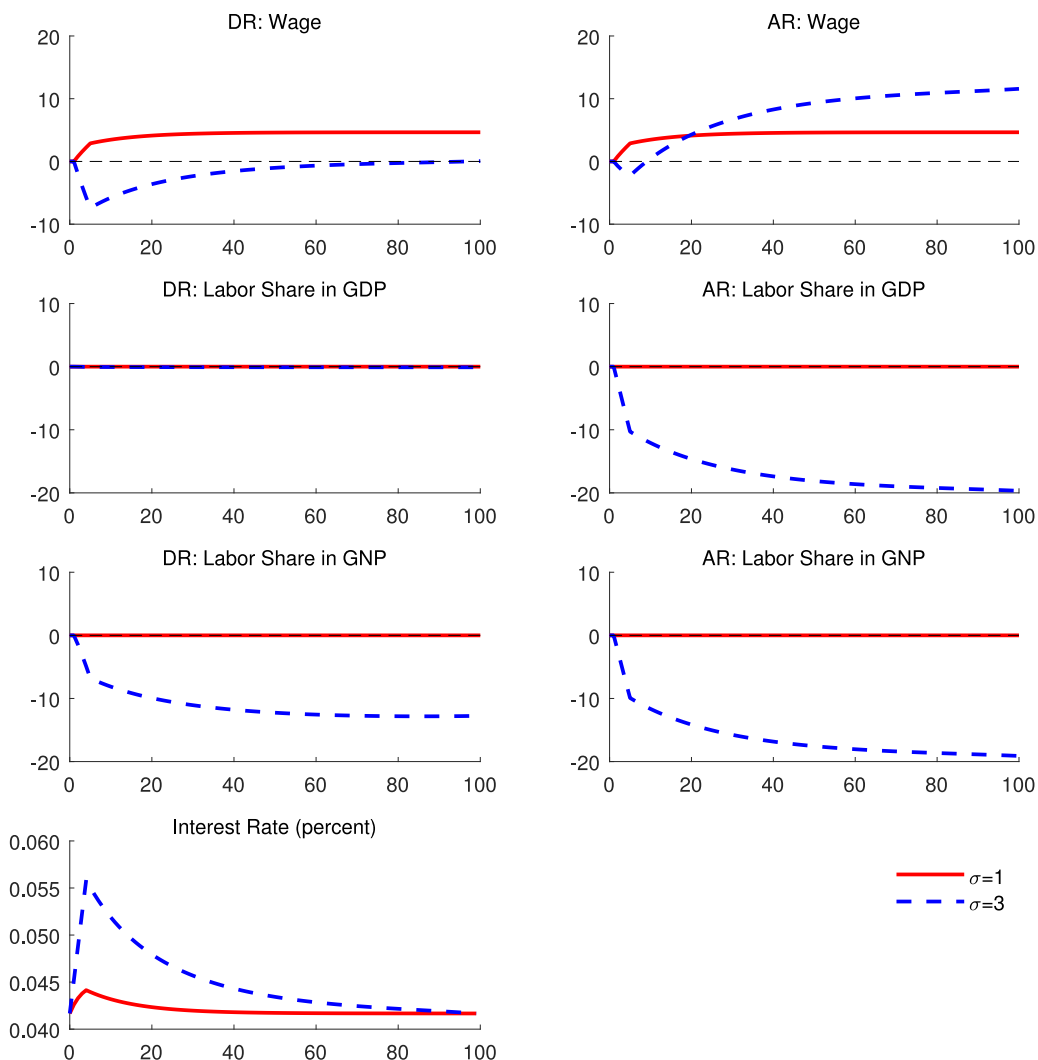


Fig. 9. Transition. Factor prices and labor shares. Note: DR=Developing Region; AR=Advanced Region. All charts plot the transition path, showing the percent changes with respect to initial steady state, except “Interest Rate” which is shown as a percent.

the large spike in the advanced country. In fact, this spike in the advanced country is only slightly larger than what would have happened under an open capital account, reflecting the small mitigating impact of capital flows from the developing region. A closed capital account thus insulates the developing country from most of the impact of the robot revolution. The developing country misses out on the chance to own some of the advanced-country robot capital stock, with consumers in the developing country losing almost all the long-run consumption benefits from the increase in the productivity of robot capital as they do not get to benefit from the high interest rate

The overall welfare impact of the robot revolution is positive for both countries, with the magnitude of welfare benefit dependent on whether capital accounts are open or not, especially for developing economies. We compute the increase in permanent consumption that would yield the same increase in welfare following the doubling of robot productivity, for each country (Fig. 11). For the advanced region, higher elasticity of substitution yields higher welfare following a robot revolution, reflecting the higher GDP increase in the long run. An open capital account leads to only marginally higher welfare than a closed one, mirroring the GDP dynamics of Fig. 10 in this case. In contrast, for the developing region, the relationship between welfare and elasticity of substitution is non-monotonic when the capital account is open. For low levels of the elasticity of substitution, welfare declines with the elasticity of substitution following the lower GDP increase in the long run. But for high levels of the elasticity of substitution, welfare increases with the elasticity of substitution as the developing region benefits more from the higher global interest rate at which it lends to the advanced region during the transition. This non-monotonic relation is naturally absent when the capital account is closed. Welfare in the developing country declines with the elasticity of substitution, mirroring the pattern of GDP increase in the long run (Fig. 7).

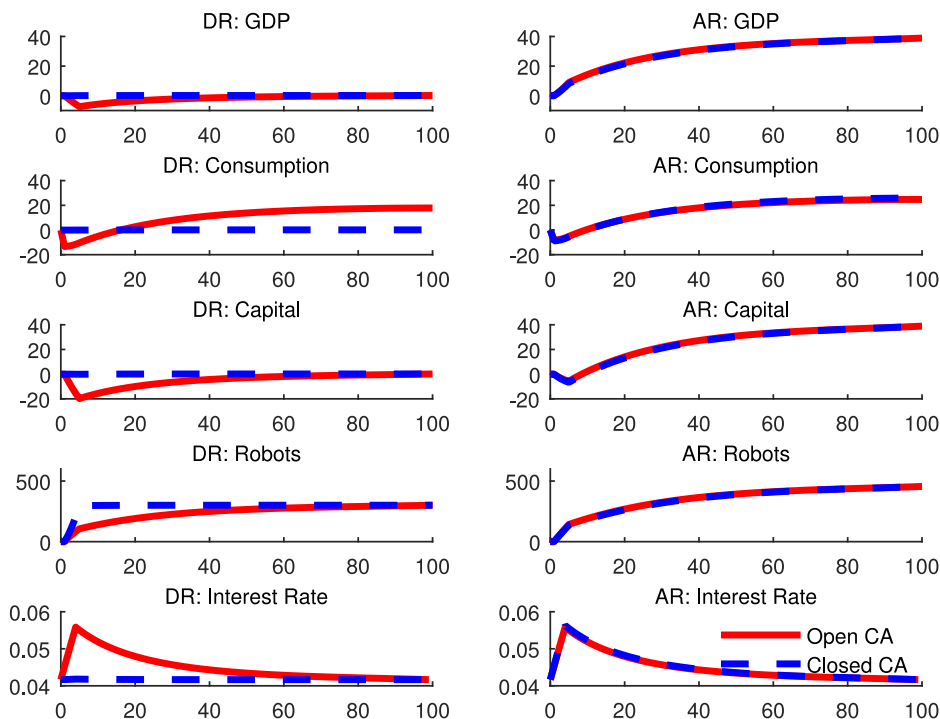


Fig. 10. Transition with closed and open capital accounts ($\sigma = 3$). Note: DR=Developing Region; AR=Advanced Region. CA = Capital Account. All charts plot the transition path, showing the percent changes with respect to initial steady state, except “Net Financial Assets” which are shown as a percent of world GDP.




Fig. 11. Equivalent percent increase in permanent consumption. Note: The charts show the equivalent percent increase in permanent consumption that follows after a doubling of robot productivity taking into account the transition dynamics, for different levels of σ . The increase is expressed as percent of consumption in the initial steady state.

3.7. Adding two skill levels

In this subsection, we extend our model to include two types of workers, skilled and unskilled. This extension has two key benefits compared to the baseline model. First, in the baseline model, all income differences between the two regions arise due to differences in aggregate productivity (TFP), which then drive the magnitude of the divergence result through the gap in wages. The two-skill-level model allows for a more realistic calibration, where differences in human capital endowments account for part of the income differences between the two regions, and so, the gap in TFP is mitigated. Second, and perhaps more important, robots may not substitute for all types of workers in the future. The two-skill-level model allows us to see how results differ when robots are complementary to skilled labor while they substitute for unskilled labor. This assumption, which is critical to the main results of this section, is a plausible and widespread view, as briefly discussed in the introduction (particularly footnote 10).

While the basic divergence result remains qualitatively unchanged, quantitatively the extent of divergence is smaller when robots only substitute for a part of the labor force. Furthermore, this richer model has implications for labor income inequality. For high substitutability between unskilled labor and robots, we find that labor income inequality rises and that the increase is higher in the advanced region.

Table 2
Calibration for two-labor model.

Parameter	Description	Value
α_K	Share of capital in production function	0.35
α_S	Share of skill in production function	0.30
\bar{L}_D	Stock of unskilled labor in the developing economy	0.98
\bar{S}_D	Stock of skilled labor in the developing economy	0.02
\bar{L}_A	Stock of unskilled labor in the advanced economy	0.70
\bar{S}_A	Stock of skilled labor in the advanced economy	0.30

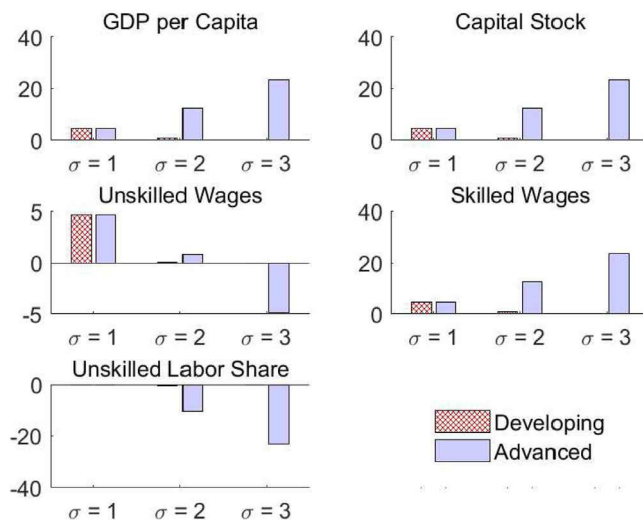


Fig. 12. Steady state comparison. Note: The figure plots the percent change in various variables between the initial steady state and the final steady state following a doubling of robot productivity. Note that each σ (plotted on the x-axis) is associated with a different initial steady state calibration as described in Section 3.4.

We consider a production function in region i given by

$$Y_{i,t} = A_i (K_{i,t}^d)^{\alpha_K} (S_{i,t})^{\alpha_S} \left[e^{\frac{1}{\sigma}} (L_{i,t})^{\frac{\sigma-1}{\sigma}} + (1-e)^{\frac{1}{\sigma}} (b_t Z_{i,t}^d)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{(1-\alpha_K-\alpha_S)\sigma}{\sigma-1}}$$

where $S_{i,t}$ and $L_{i,t}$ represent skilled and unskilled labor, respectively.

Our calibration strategy remains broadly unchanged. We calibrate the endowments of skilled workers in the two regions to match the share of workers with greater than secondary education in the US and in low income countries. Of course, the advanced region has relatively more skilled workers compared to the developing region. For the share of skilled labor in the production function, we follow Berg et al. (2018). We continue to calibrate aggregate productivity in the developing region A_D and unskilled labor share parameter e to match relative GDP in the two regions and robot share in income in advanced economies. Our calibration strategy implies that the capital and skilled labor shares are the same in both regions, and the robot share in the initial steady state in the advanced region is maintained at 4 percent for all σ . The new parameters are summarized in Table 4. Other parameters remain as in Table 1.

Accounting for differences in skill across countries leads to lower differences in calibrated TFP. In particular, while A_D was calibrated to a value of about 0.19 in the baseline model, the calibrated value in the two-skill-level model is almost twice as high at about 0.38 (both with respect to a TFP normalized at 1 for the advanced economy).

As in the baseline model with only one type of labor, the divergence effect emerges in the long run (Fig. 12). The advanced region experiences a much larger increase in per capita GDP for large values of σ than the developing region.

Quantitatively, the divergence effect is smaller than in the one-skill-level model reflecting a lower TFP gap across countries and a lower share of unskilled labor/robots. For example, when $\sigma = 3$, the advanced economy grows by about 39 percent in the baseline model while only growing by about 23 percent in the two-skill-level model (Table 3). The reason is that robots only substitute for a subset of the labor force and total factor productivity is not so different. An increase in robot productivity leads to greater investment in robots (and complementary physical capital). However, as robots complement skilled labor which is in fixed supply, this greater investment also raises skilled wages which reduces the incentive to invest. A similar effect emerges in the developing economy, but is weaker in magnitude, leading to a milder divergence. This dampening force did not exist in the baseline model.

This version of the model also predicts increases in labor income inequality for large values of σ . Skilled wages in both economies increase in line with per-capita GDP as skilled labor complements robots. However, the absolute level of unskilled

Table 3
Percent change in per-capita GDP following Increase in robot productivity.

	Model: 1 Sector and 1 Skill Level		Model: 1 Sector and 2 Skill Levels		Model: 2 Sectors and 2 Skill Levels	
	DR	AR	DR	AR	DR	AR
$\sigma = 1$ (Cobb-Douglas)	4.7	4.7	4.7	4.7	4.7	4.7
$\sigma = 2$	0.9	14.5	0.7	12.6	-0.8	11.2
$\sigma = 3$	0.1	38.9	0.1	23.5	-4.0	19.5

Note: DR = Developing Region; AR = Advanced Region. Table shows the percent change in per capita GDP following a doubling of robot productivity for different σ for the three models. Columns 2 and 3 report results for the baseline model described in Sections 3.1 through 3.5. Columns 4 and 5 report results for the model described in Section 3.7 while Columns 6 and 7 report results for the model described in Section 4.

wages can fall following an increase in robot productivity for high values of σ . This is because robots substitute for unskilled workers, and the increase in robot investment following an increase in robot productivity reduces the demand for unskilled workers. The fall in unskilled wages is larger in advanced economies as there is greater investment in robots in this region, thus reducing unskilled labor demand by more.

4. Two-Sector model with two skill levels

Until now we have assumed only one good. However, the output composition of the developing region is likely to reflect its relatively greater abundance of unskilled labor. In this section we extend our baseline model to include two sectors that are distinguished by the intensity with which they use skilled and unskilled labor. With only this addition, and again the assumptions that robots substitute for—and low-income countries are relatively well-endowed in—unskilled labor, we find a terms-of-trade channel that tends to amplify the divergence effect.²⁶

The model continues to feature two regions indexed by i , representing an advanced economy ($i = A$) and a developing economy ($i = D$). In addition, the model now has two types of goods indexed by j ($j = T1, T2$), with both goods being tradable. The household's utility function is given by:

$$\sum_t \beta^t \frac{\left[(C_{i,t}^{T1})^t (C_{i,t}^{T2})^{1-t} \right]^{1-\frac{1}{\sigma}}}{1 - \frac{1}{\sigma}}$$

where $C_{i,t}^{T1}$ and $C_{i,t}^{T2}$ is consumption of good T1 and T2 respectively, by household i (with $i = A, D$) in period t .

Household i seeks to maximize utility given its budget constraint:

$$C_{i,t}^{T1} + p_t^{T2} C_{i,t}^{T2} + I_{i,t}^K + I_{i,t}^Z + (B_{-i,t+1} - B_{i,t+1}) = r_t^K K_{i,t} + r_t^Z Z_{i,t} + w_{i,t} \bar{L}_i + w_{i,t}^S \bar{S}_i + (1 + r_t^B)(B_{-i,t} - B_{i,t}) + \Pi_{i,t} \quad (12)$$

where p_t^{T2} is the relative price of the second good, $w_{i,t}^S$ is the wage of skilled workers, \bar{S}_i is the endowment of skilled labor, and $w_{i,t}$ and \bar{L}_i are the wages and endowment of unskilled workers. Note that good T1 is assumed to be the numeraire, and we implicitly assume that only good T1 is used for capital and robot accumulation.²⁷

A representative firm operates in each sector in each region. The production function for sector j in region i is given by

$$Y_{i,t}^j = A_i (K_{i,t}^{j,d})^{\alpha_{K,j}} (S_{i,t}^j)^{\alpha_{S,j}} \left[(e_j)^{\frac{1}{\sigma_j}} e_j (L_{i,t}^j)^{\frac{\sigma_j-1}{\sigma_j}} + (1 - e_j)^{\frac{1}{\sigma_j}} (b_t Z_{i,t}^{j,d})^{\frac{\sigma_j-1}{\sigma_j}} \right]^{\frac{(1-\alpha_{K,j}-\alpha_{S,j})\sigma_j}{\sigma_j-1}}$$

Note that A_i is still the only parameter that varies across regions and is assumed to be the same across the two sectors. The share parameters for capital, skilled labor, and unskilled labor ($\alpha_{K,j}$, $\alpha_{S,j}$ and e_j), and the elasticity of substitution between unskilled labor and robots (σ_j), can vary across sectors but are the same across the two regions.

The equilibrium definition is the same as used in Section 3.3, with the market clearing conditions now given by:

$$\bar{L}_i = \sum_{j=\{T1,T2\}} L_{i,t}^j \quad \forall i$$

$$K_{i,t} = \sum_j K_{i,t}^{j,d}$$

$$Z_{i,t} = \sum_j Z_{i,t}^{j,d}$$

²⁶ To be somewhat more concrete, examples of low-skill industries are: food and beverage manufacturing, and textiles and apparel manufacturing; and for high-skill industries are: manufacturing of aircraft and manufacturing of drugs and medicines (Wörz, 2004).

²⁷ Results are similar if we assume that robots and capital are produced using a Cobb–Douglas aggregate of the two goods.

Table 4
Calibration for two-sector model.

Parameter	Description	Value
$\alpha_{K,T1} = \alpha_{K,T2}$	Share of capital in production function	0.35
$\alpha_{S,T1}$	Share of skill in T1 production function	0.35
$\alpha_{S,T2}$	Share of skill in T2 production function	0.25
\bar{L}_D	Stock of unskilled labor in the developing economy	0.98
\bar{S}_D	Stock of skilled labor in the developing economy	0.02
\bar{L}_A	Stock of unskilled labor in the advanced economy	0.70
\bar{S}_A	Stock of skilled labor in the advanced economy	0.30

$$\sum_i (C_{i,t}^{T1} + I_{i,t}^K + I_{i,t}^Z) = \sum_i Y_{i,t}^{T1}$$

$$\sum_i C_{i,t}^{T2} = \sum_i Y_{i,t}^{T2}$$

Note that as the production functions in both sectors exhibit constant returns to scale, and as both goods are tradable, the model can have a corner equilibrium in which one of the goods is only produced in one region.²⁸

4.1. Calibration of the initial steady states of two-sector model

Our calibration strategy follows the approach in [subsection 3.4](#). For each level of σ , ranging from 1 to 3, we calibrate a separate initial steady state where we choose e_{T1} , e_{T2} and A_D to match three moments: the relative GDP per capita between the two regions and a robot share of 4 percent in the advanced economy in each sectors. Both sectors are assumed to have the same elasticity of substitution between unskilled workers and robots.

As in [subsection 3.7](#), we calibrate the endowments of skilled workers in the two regions to match the share of workers with greater than secondary education in the US and in low income countries. Furthermore, we assume that the main difference between the two sectors is the relative importance of skilled and unskilled labor in the production function. In particular, the T1 sector has a skilled share which is 10 percentage points higher compared to the T2 sector. [Table 4](#) summarizes the values of the additional parameters we calibrate relating to technology in the two sectors and the endowment of the two types of labor in each region. We maintain the same calibrated values as in [Table 1](#) for common parameters across the two models.

As in the one-sector model, the robot share varies endogenously in the developing economy ([Fig. 13](#)). In particular, the robot share is lower in the developing region for values of σ greater than 1, while the unskilled labor share is higher. On the other hand, the robot share in output is the same in the advanced economy for different values of σ because that is one of the targets of our calibration strategy.

The skilled labor share in this model is the weighted average of the skilled labor shares in both sectors and does not depend on σ . The skilled labor share in the developing region is 0.25, which is the calibrated value of $\alpha_{S,T2}$ because the region only produces the unskilled-labor intensive good (T2) in equilibrium for our baseline calibration. Since the advanced region produces both goods, its skilled labor share lies between the values of $\alpha_{K,T1}$ and $\alpha_{K,T2}$.

4.2. Long-run impact of the robot revolution in the two-sector model

In this model, a combination of two forces determine the relative impact of an increase in robot productivity on the two regions—a direct effect (akin to the one-sector model) and a second effect due to changes in relative prices.

After a doubling of robot productivity, the direct effect is an incentive to accumulate more robots and complementary physical capital in both sectors, similar to the one-sector model of [Section 3](#).²⁹

However, in addition to this direct effect, the two-sector model also has an indirect effect working through changes in relative prices. The increase in robot productivity decreases the demand for unskilled labor, especially when robots and unskilled labor are highly substitutable (i.e., when σ is large), thus lowering unskilled wages. As good T2 uses unskilled

²⁸ Our solution algorithm first solves the unconstrained problem, and then checks whether the non-negativity constraint on the solution holds. If the non-negativity constraint does not hold because the quantity produced of one of the goods is negative, the algorithm assumes that production of that good is zero and allocates all the inputs to the production of the other good. For example, with our calibration for this section, we find that T1 is not produced in the developing region ($i = D$) in equilibrium. For this allocation to be an equilibrium, it must be the case that given factor prices, the marginal/average cost of producing T1 in the developing region is higher than the price of T1 in equilibrium in the world market. This inequality holds in our equilibrium, indicating that there is, in fact, no incentive to produce T1 in the developing region, as doing so would lead to losses.

²⁹ Furthermore, within each region, the direct effect is larger for good T2 as this sector is more intensive in the robots and unskilled labor composite ($1 - \alpha_{K,T2} - \alpha_{S,T2} > 1 - \alpha_{K,T1} - \alpha_{S,T1}$). Thus, a doubling in robot productivity will lead to a larger increase in investment in robots and capital in the T2 sector, all else equal.

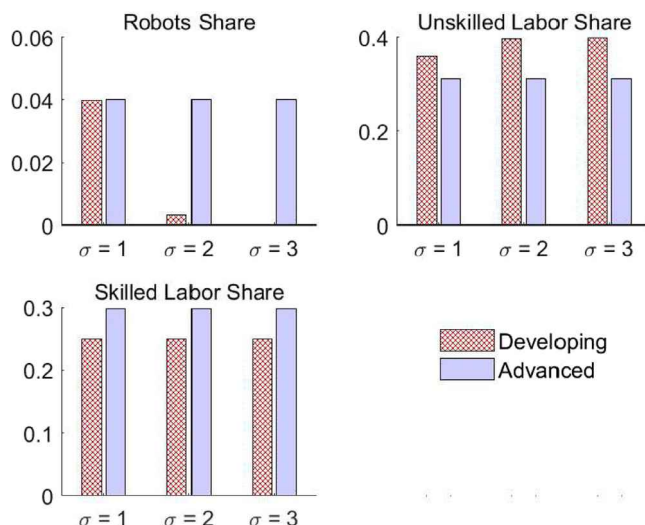


Fig. 13. Moments in the initial steady state. Note: The top panel of the figure plots the robot share and unskilled labor share for each steady state, that is, for each value of σ . The bottom left panel plots the skilled labor share.

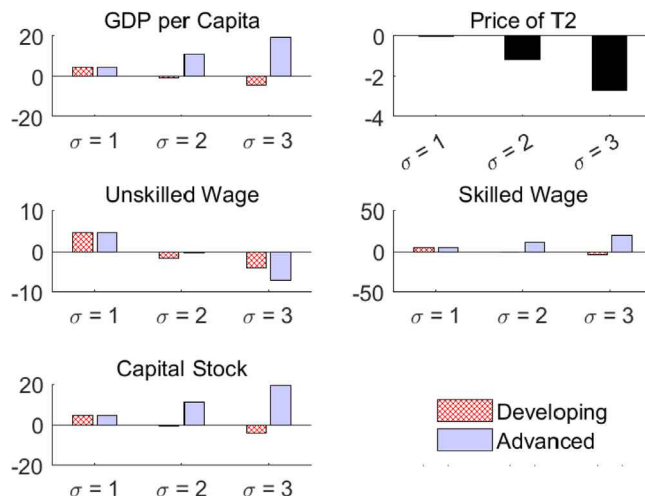


Fig. 14. Steady state comparison (percent changes with respect to initial steady state). Note: The figure plots the percent change in various variables between the initial steady state and the final steady state following a doubling of robot productivity. Note that each σ (plotted on the x-axis) is associated with a different initial steady state calibration as described in Section 4.1.

labor (and robots) more intensively, this results in a decline in the relative price of good T2.³⁰ This decline in relative price implies less incentive to invest in robots and capital in the T2 sector, thus countervailing the direct effect.³¹

Following a doubling of robot productivity, the same divergence effect seen in the one-sector model emerges when $\sigma > 1$, with the increase in per capita GDP being much larger in the advanced region (Figure 14). This is driven by the direct effect which is larger in the advanced economy when $\sigma > 1$ —the intuition for this result is the same as in the one sector model, as the robot share in output is larger in the initial steady state in the advanced region.

However, while per capita GDP always increased in the long run in the one-sector model following an increase in robot productivity, in the two sector-model *income levels* can actually decline in the developing economy for large values of σ due to the effects arising from changes in relative prices.

³⁰ On the other hand, if robots were to substitute for skilled workers (i.e., if skilled labor were to be the input combined with robots using a CES technology, instead of unskilled labor), then the terms of trade effect would be reversed with the relative price of T2 increasing. See Acemoglu and Restrepo (2018b) for a model where automation can impact high and low skilled workers.

³¹ While we assume that capital and robots are produced using good T1 only, the qualitative result does not change if we allow capital to be produced using a Cobb-Douglas aggregate of the two goods. In particular, the price of T2 relative to that of investment will fall as long as investment uses T1 good to some extent.




Fig. 15. Decomposing steady state change (percent changes with respect to initial steady state). Note: The figure decomposes the percent change in GDP between the initial and final steady state following a doubling of robot productivity, using an intermediate step in which the price of the second good is kept constant at its initial level. The dotted red bar shows the increase in GDP that would have happened if the price of the second good had been kept constant at the level of the initial steady state. Note that this is not an equilibrium outcome and, in fact, for $\sigma = 3$ is associated with zero production of good 1 in both countries. The solid blue bar shows the change between steady states and so, it incorporates the decline in the price of the second good needed to reach equilibrium. Note that each σ (plotted on the x-axis) is associated with a different initial steady state calibration as described in Section 4.1.

To get a sense of why per capita GDP can fall in the two sector model, Fig. 15 decomposes the two channels (i.e., share-in-production or direct effect, and terms-of-trade or indirect effect) by considering an intermediate step in which robot productivity increases but the price of the second good is kept constant at the level of the initial steady state. The change between the initial steady state and the intermediate step (in red bars) resembles the divergence results produced by the share-in-production channel with the developing region benefiting less from the robot revolution than the advanced region as robot utilization is lower. Through this channel, GDP grows less in the developing region, but does not decline. Instead, the decline in GDP for the developing region emerges when the price of the second good declines to restore the equilibrium and reach the final steady state (blue bars). It is then the terms-of-trade channel that pushes GDP to drop in the final steady state in the developing region.

Considering how output in both the sectors is affected can provide further intuition for this result:

- In the advanced economy, output of both goods increases in equilibrium. For good $T1$ this is simply reflecting the direct effect. For good $T2$, there are two forces. While the direct effect provides incentive to invest in more robots and capital, the decline in relative price of $T2$ reduces the incentive to invest in this sector. In equilibrium, the direct effect is larger in the advanced economy, thus resulting in higher output of good $T2$ in equilibrium.³²
- In the developing region, output of good $T2$ declines in equilibrium. As in the advanced economy, the direct effect provides incentive to invest in more robots. However, the direct effect is comparatively small in the developing region when $\sigma > 1$ because robot share in output is small in the initial steady state (as in the one-sector model). Thus, in the developing region, the effect from a decline in relative price of $T2$ dominates, and output of good $T2$ falls. Furthermore, as the developing region only produces good $T2$, GDP in the region falls.

Income inequality also increases substantially in both regions when robots and unskilled labor are highly substitutable. The greater the σ , the greater the increase in the robot share at the expense of unskilled labor share. But also, the skill premium expands as skilled wages increase by more than unskilled wages, which even drops for large degrees of substitutability.

5. Conclusion

Developing economies face tremendous challenges in their attempt to converge to the income levels of emerging and advanced economies. Of course, opportunities emerge and disappear as the global landscape changes. The economic environment and potential sources of growth that, for example, the US and China faced during their early stages of economic development are remarkably different from what Cambodia and Tanzania are currently facing. And specifically, automation over the last few decades has been rapidly transforming the global economic landscape for all countries, including developing economies.

This paper considers the implications for developing countries of machines that substitute pervasively for labor, a topic that has generated a burgeoning literature focused on advanced countries. It makes simple and plausible assumptions: (1) the AI revolution can be modeled as an increase in productivity (or reduction in cost) of a distinct type of capital—dubbed “robots”—that substitutes closely with labor; (2) the only difference between the advanced and developing country is the level of TFP; and (3) labor is immobile across regions.³³

This setup is minimalist, but the resulting conclusions are powerful and general: improvements in the productivity of “robots” drive divergence between advanced and developing countries. Advanced countries will make greater use of such

³² This result has a flavor of “re-shoring”, in that the higher robot productivity drives some of the production of the low-skill-intensive good to the advanced region.

³³ Allowing for migration of unskilled labor from developing to advanced countries could help mitigate the divergence results.

machines, since they will have higher wages, and they will thus differentially benefit from a reduction in their cost. And in the transition, if capital is mobile, the high profitability of both robots (because of the increase in productivity) and of traditional capital (which complements robots) pulls capital from the developing to the developed country, resulting in a transitional decline in GDP in the developing country. If, instead, capital is immobile, there is no “uphill” capital flow, but the developing region loses the long-run increase in consumption associated with its transitional accumulation of high-yielding advanced-country robot capital. It is worth underscoring that none of these results hold in the textbook Cobb-Douglas world, where each region benefits equally from improvements in technology.

We also consider an extended model with two types of labor, with “robots” substituting for one type (“unskilled”) and complementing the other (“skilled”). In this case, there is a permanent decline in the terms of trade in the developing region, insofar as it is relatively rich in unskilled labor. With this additional channel, the developing country could observe a fall not just in relative but in absolute GDP.

Our framework is simple and the key assumptions plausible; the resulting mechanisms thus seem fundamental. With this simplicity, we have ignored many important considerations, of which two off-setting ones deserve mention here, both involving the way in which the new wave of technologies may change the dynamics of development. First, the new wave of technologies may differentially benefit low-income countries by allowing them to “leap-frog” earlier hurdles to development. Cloud technologies and mobile phones may obviate the need for the construction of extensive on-the-ground infrastructure, for example. Or, new technologies may allow global supply chains to extend to services bringing poor regions more quickly into the global economy.

On the other hand, these new technologies may also have implications for the viability of rapid catch-up through industrialization—e.g. the “flying geese” model by which poor countries grow rapidly by moving up the value chain, learning-by-doing along the way. Our model captures some of the flavor of the problem: as mentioned in footnote 32, we observe something like “re-shoring” in the way an increase in robot productivity pulls some low-skill-intensive production to the advanced country, and also in the way the advent of low-cost robots draws capital from poor to rich countries, with capital mobility. However, in our setup there is nothing special for growth about the production of any good, and in particular no learning-by-doing that would make anything like reshoring per se detrimental to growth.

There is no silver bullet for averting divergence.³⁴ Developing countries, more urgently than ever before, need to invest in raising aggregate productivity and skill levels so that the labor force be complemented rather than substituted by robots, but of course this is easier said than done. In our baseline model, increases in total factor productivity—which are a proxy for the many institutional and other fundamental differences between developing and advanced countries not captured by labor and capital inputs—are especially beneficial by incentivizing more robots and physical capital accumulation. Of course, such improvements are always beneficial, but the gains are stronger in the context of the AI revolution. Our two-sector, two-skill model also underscores the importance of human capital accumulation to prevent divergence, and points to potentially heterogeneous growth dynamics among developing economies with different skill levels.

That said, continued advances in automation and the explosive use of robots in virtually all sectors of the economy threatens the complementarity that currently exists in these countries between humans and machine. The landscape is likely going to be much more challenging for developing countries which have hoped for high dividends from a much-anticipated demographic transition. By 2030, more than half of the increase in the global labor force is expected to come from the African continent. This was hailed by policymakers as possibly the continent’s big chance to benefit from China’s graduating middle-income status (de Carvalho Chamon and Kremer, 2006). Our findings show that robots may steal these jobs from Africa and unless a drastic shift to productivity gains and education investment is put in place rapidly, Africa’s much anticipated demographic transition could yield negative not positive dividends.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:[10.1016/j.jmoneco.2022.01.004](https://doi.org/10.1016/j.jmoneco.2022.01.004).

References

- Abdychev, A., Alonso, C., Alper, E., Desruelle, D., Kothari, S., Liu, Y., Perinet, M., Rehman, S., Schimmelpennig, A., Sharma, P., 2018. The Future of Work in Sub-Saharan Africa. IMF Departmental Paper Series (18/18).
- Acemoglu, D., 2021. Harms of AI. Technical Report 29247. National Bureau of Economic Research.
- Acemoglu, D., Restrepo, P., 2018. Artificial Intelligence, Automation and Work. Technical Report. National Bureau of Economic Research.
- Acemoglu, D., Restrepo, P., 2018. Low-skill and high-skill automation. *J Hum Cap* 12 (2), 204–232.
- Acemoglu, D., Restrepo, P., 2018. Modeling automation. In: *AEA Papers and Proceedings*, Vol. 108, pp. 48–53.
- Acemoglu, D., Restrepo, P., 2018. The race between man and machine: implications of technology for growth, factor shares, and employment. *American Economic Review* 108 (6), 1488–1542.
- Acemoglu, D., Restrepo, P., 2019. Automation and new tasks: how technology displaces and reinstates labor. *Journal of Economic Perspectives* 33 (2), 3–30.

³⁴ In a sense the results in this paper are an example of the general phenomenon underscored in Korinek and Stiglitz (2019) and Korinek and Stiglitz (2021), which emphasizes the need for redistribution to make everyone better off in the face of technical progress, in general. A key feature in our setting is that large-scale international redistribution is much less plausible than in the domestic context as it would require reforms at the supra-national level (Korinek and Stiglitz, 2021).

- Acemoglu, D., Restrepo, P., 2019. Robots and jobs: evidence from US labor markets. *Journal of Political Economy* 0 (ja), null.
- , 2019. In: Agrawal, A., Gans, J., Goldfarb, A. (Eds.), *The Economics of Artificial Intelligence: An Agenda*. The University of Chicago Press, Chicago.
- Akerman, A., Gaarder, I., Mogstad, M., 2015. The skill complementarity of broadband internet *. *Q J Econ* 130 (4), 1781–1824. doi:10.1093/qje/qjv028.
- Autor, D.H., Dorn, D., 2013. The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review* 103 (5), 1553–1597. doi:10.1257/aer.103.5.1553.
- Autor, D.H., Levy, F., Murnane, R.J., 2003. The skill content of recent technological change: an empirical exploration. *Q J Econ* 118 (4), 1279–1333.
- Bank, W., 2019. *World Development Report: The Changing Nature of Work*. Technical Report. World Bank.
- Benzell, S.C., Laurence, J.K., Guillermo, L., Victor, Y.Y., 2021. Simulating Endogenous Global Automation. Technical Report 29220. National Bureau of Economic Research.
- Berg, A., Buffie, E.F., Zanna, L.-F., 2018. Should we fear the robot revolution?(the correct answer is yes). *J Monet Econ* 97, 117–148.
- Bessen, J.E., 2017. Automation and jobs: when technology boosts employment. Boston Univ. School of Law, Law and Economics Research Paper (17-09).
- Brynjolfsson, E., McAfee, A., 2014. *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.
- de Carvalho Chamon, M., Kremer, M.R., 2006. Asian growth and african development. *American Economic Review* 96 (2), 400–404.
- Caselli, F., Manning, A., 2019. Robot arithmetic: new technology and wages. *American Economic Review: Insights* 1 (1), 1–12.
- Eden, M., Gaggl, P., 2018. On the welfare implications of automation. *Rev Econ Dyn* 29, 15–43. doi:10.1016/j.red.2017.12.004.
- Eden, M., Gaggl, P., 2019. Do poor countries really need more IT? *World Bank Econ Rev* 34 (1), 48–62.
- Ford, M., 2015. *The rise of the robots: Technology and the threat of mass unemployment*. Oneworld publications.
- Frey, C.B., 2019. *The technology trap - Capital, labor and power in the age of automation*. Princeton University Press.
- Frey, C.B., Osborne, M.A., 2017. The future of employment: how susceptible are jobs to computerisation? *Technol Forecast Soc Change* 114, 254–280.
- Grace, K., Salvatier, J., Dafoe, A., Zhang, B., Evans, O., 2018. When will AI exceed human performance? evidence from AI experts. *Journal of Artificial Intelligence Research* 62, 729–754.
- Graetz, G., Michaels, G., 2018. Robots at work. *Review of Economics and Statistics* 100 (5), 753–768.
- Korinek, A., Stiglitz, J. E., 2019. Artificial intelligence and its implications for income distribution and unemployment, In: Agrawal et al. (2019). pp. 349–390.
- Korinek, A., Stiglitz, J.E., 2021. Artificial Intelligence, Globalization, and Strategies for Economic Development. Working Paper 28453. National Bureau of Economic Research doi:10.3386/w28453.
- Mann, K., Püttmann, L., 2018. Benign effects of automation: new evidence from patent texts. Available at SSRN 2959584.
- Mokyr, J., Vickers, C., Ziebarth, N.L., 2015. The history of technological anxiety and the future of economic growth: is this time different? *Journal of Economic Perspectives* 29 (3), 31–50.
- Nedelkoska, L., Quintini, G., 2018. Automation, Skills Use and Training. Technical Report 202. OECD Social, Employment and Migration Working Papers.
- Nordhaus, W.D., 2015. Are we Approaching an Economic Singularity? Information Technology and the Future of Economic Growth. Technical Report. National Bureau of Economic Research.
- of Robotics, I.F., 2012. *World Robotics Industrial Robots*. Technical Report. IFR.
- Rodrik, D., 2016. Premature deindustrialization. *Journal of Economic Growth* 21 (1), 1–33.
- Sachs, J.D., 2018. R&D, structural transformation, and the distribution of income. In: *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- Sachs, J.D., 2019. Some brief reflections on digital technologies and economic development. *Ethics & International Affairs* 33 (2), 159–167. doi:10.1017/S0892679419000133.
- Sachs, J.D., Kotlikoff, L.J., 2012. *Smart Machines and Long-term Misery*. Technical Report. National Bureau of Economic Research.
- Shiller, R.J., 2019. Narratives about Technology-Induced Job Degradations Then and Now. Technical Report. National Bureau of Economic Research.
- Susskind, D., 2019. Re-thinking the capabilities of technology in economics. *Economics Bulletin* 39 (1), 280–288.
- Susskind, D., 2020. *A World Without Work*. Macmillan.
- Wörz, J., 2004. Skill intensity in foreign trade and economic growth. The Vienna Institute for International Economic Studies. WIIW working papers series.
- Yusuf, S., 2017. Automation, AI, and the Emerging Economies. Technical Report. Center for Global Development.
- Zeira, J., 1998. Workers, machines, and economic growth. *Q J Econ* 113 (4), 1091–1117.