

DISCUSSION PAPER SERIES

DP15772

Artificial Intelligence, Globalization, and Strategies for Economic Development

Anton Korinek and Joseph E Stiglitz

DEVELOPMENT ECONOMICS

INTERNATIONAL MACROECONOMICS AND FINANCE

INTERNATIONAL TRADE AND REGIONAL ECONOMICS

MACROECONOMICS AND GROWTH

CEPR

Artificial Intelligence, Globalization, and Strategies for Economic Development

Anton Korinek and Joseph E Stiglitz

Discussion Paper DP15772
Published 06 February 2021
Submitted 04 February 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Development Economics
- International Macroeconomics and Finance
- International Trade and Regional Economics
- Macroeconomics and Growth

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Anton Korinek and Joseph E Stiglitz

Artificial Intelligence, Globalization, and Strategies for Economic Development

Abstract

Progress in artificial intelligence and related forms of automation technologies threatens to reverse the gains that developing countries and emerging markets have experienced from integrating into the world economy over the past half century, aggravating poverty and inequality. The new technologies have the tendency to be labor-saving, resource-saving, and to give rise to winner-takes-all dynamics that advantage developed countries. We analyze the economic forces behind these developments and describe economic policies that would mitigate the adverse effects on developing and emerging economies while leveraging the potential gains from technological advances. We also describe reforms to our global system of economic governance that would share the benefits of AI more widely with developing countries.

JEL Classification: F63, O32, O25, D63

Keywords: artificial intelligence, labor-saving progress, inequality, terms-of-trade losses

Anton Korinek - anton@korinek.com

University of Virginia, Economics and Darden School of Business and CEPR

Joseph E Stiglitz - jes322@columbia.edu

Columbia University, University Professor

Acknowledgements

We thank Avital Balwit, Katya Klinova, Rafael Proença, Martin Schindler, Don Suh and participants at the 2020 IMF/INET Conference on Macroeconomics in the Age of AI for insightful comments and suggestions. Financial support from the Institute for New Economic Thinking is gratefully acknowledged.

1 Introduction

All around the world, there are fears of job losses and increasing inequality resulting from AI and related forms of automation technologies. Developing countries and emerging market economies have even more reason to be concerned than high-income countries, as their comparative advantage in the world economy relies on abundant labor and natural resources.² Declining returns to labor and natural resources as well as the winner-takes-all dynamics brought on by new information technologies could lead to further immiseration in the developing world. This would undermine the rapid gains that have been the hallmark of success in development over the past fifty years, and threaten the progress made in reducing poverty and inequality.

For much of the past half-century, there was a presumption that advances in technology would benefit all—embodied by the trickle-down dogma that characterized neoliberalism. However, there was never any economic theory that said that advances in technology would *necessarily* benefit all; and much economic research cautioned that this might not be the case and that technological progress may well create both winners and losers – see Korinek and Stiglitz (2019) for a review. As long as the winners and losers from technological progress are located within the same country, there is at least the possibility that domestic policy measures can compensate the losers. However, when technological progress deteriorates the terms of trade and thus undermines the comparative advantage of entire countries, domestic policy measures are insufficient to provide compensation to the losers of progress, and entire nations may be worse off.

This paper argues that concerns about the impact of technological progress on developing countries may be justified – developing countries may face a stark new set of challenges going forward. The paper also shows, however, that there are policy measures that mitigate the adverse effects so that advances in technology may lead to a world with greater *shared* prosperity. These encompass both domestic policies and development strategies as well as international cooperation and a rewriting of the global rules governing the information economy.

We start by laying out two simple models that capture the key properties of AI and related technologies that underlie the concerns about recent technological progress. AI is likely to be resource-saving and labor-saving, devaluing the sources of comparative advantage of many developing countries, deteriorating their terms of trade, and potentially making them worse off in absolute terms. We show that innovations may be labor-saving even in the long run when the capital stock in the economy has adjusted to the novel technological developments. Some countries (e.g. China) have experience a sufficiently deep economic transformation that their comparative advantage has shifted, so that they may be among the winners of the AI revolution. For other countries, however, AI may arrest or reverse the convergence towards the standards of living of rich countries that they had experienced for much of the past century.

Information technologies such as AI also tend to give rise to natural monopolies, creating a small set of so-called superstar firms that are located in a few powerful countries but serve the entire world economy. The rate and direction of technological progress chosen by the market are generally suboptimal. This creates the possibility of steering innovation in AI and other technologies in directions that increase global social welfare, for example by preserving the planet or creating satisfying employment opportunities

² Our analysis and findings in this paper apply equally to most developing countries and emerging market economies. For succinctness of exposition, we will only refer to developing countries in the remainder of the paper.

rather than substituting for labor and creating more unemployment and inequality. With the direction of innovation largely set by policies in the advanced countries, there is no presumption that the course of technology will take such forms. However, even if advanced countries determine the broad thrust of innovation, there is scope for adaptation towards more labor-using technologies in developing countries.

Taking a step back, we evaluate to what extent the discussed concerns about technological progress are justified, given what we know at present. We observe that there is vast uncertainty about the impact of artificial intelligence, even among experts in the field. Some argue that AI is less important than the big innovations of the 20th century and will have rather limited impact on the economy, whereas others go as far as predicting that AI will lead to more rapid technological progress than mankind has ever seen before. We also discuss how to reconcile the buzz among technologists over the past decade with economic data that suggests rather modest productivity increases over the period – encapsulated by the so-called productivity puzzle. Finally, we analyze how the forces generated by progress in AI interact with other recent developments, in particular with the COVID-19 pandemic, with secular population dynamics, and with the need for a Green Transition.

Given all the uncertainties surrounding AI, and applying our insights on steering innovation to economic research, we observe that it would be welfare-maximizing to steer our own research in directions where the expected social value added of economic analysis is greatest, and to think particularly hard about potential events that would be highly disruptive to our society.

To grasp the historical nature of what is going on, we look at the broader history of technological progress. Humanity spent much of its history at a Malthusian stage in which the vast majority of the population lived at subsistence levels. The Industrial Revolution that lifted living standards started a bit over two centuries ago and was but a blip in the history of humanity. For developing countries, the era of manufacturing-based export-led growth that enabled the East Asian Miracle stretched over the past half-century – only one quarter of the history of the Industrial Revolution. It is easily conceivable that we are now going into another era. There is even a risk that the terms-of-trade losses generated by progress in AI may push developing countries back into the Malthusian dynamics that have characterized much of our history. But the Industrial Revolution also offers ample lessons on how to manage innovation in a positive way: technological revolutions are very disruptive, but collective action can mitigate the adverse effects and generate an environment in which the gains are shared broadly. The labor-using nature of the Industrial Revolution ushered in an Age of Labor in which the economic gains of workers also shifted political dynamics in their favor, but there is a risk that future labor-saving progress may do the opposite. The decline of manufacturing calls into question the manufacturing-based export-led growth model of the past, but we observe that a new development model that follows a more multi-pronged strategy may take its place.

The key policy question for developing countries is how they can improve the likelihood of benign outcomes from technological progress. We delineate a specific set of policies that ensure that technological advances lead to broadly shared prosperity and that mitigate their adverse effects on income distribution. Taxation and redistribution are a first line of measures to compensate the losers of progress, although the scope for redistribution may be limited in developing countries. Targeted expenditure policies can serve double duty by providing an income to workers and providing a valuable social return – for example, investments in education or infrastructure are both labor-intensive and enhance human capital and the physical infrastructure of countries, both of which serve to bridge the digital divide and to ensure that all citizens can participate in the opportunities afforded by digital

technologies. To replace the manufacturing-based export-led growth model, developing countries could steer technological progress and technology adoption in new directions, in part by leveraging the opportunities that modern AI and other digital technologies afford in agriculture and services

Finally, we describe a set of policies at the supra-national level to reform our global system of governance in a way that enables developing countries to benefit more from advances in AI and other information technologies while addressing the downsides of these new technologies. A global tax regime for the digital age would enable countries to raise taxes on transactions that occur within their borders. And a global minimum tax rate on capital might encourage capital-saving/labor-using innovation. Competition policy is also increasingly a policy area that transcends national borders as the footprint of the digital giants is global and authorities in their countries of origin do not face the incentives required to ensure a competitive marketplace. Current intellectual property regimes are not attuned to the needs and circumstances of developing countries and could be reformed. Moreover, addressing information policy including the regulation of data at the supra-national level would provide a voice to developing countries that could otherwise not influence the design of such policies.

The remainder of this article is organized as follows. In the second section, we introduce two simple models to provide an overview of the downside risks of automation and AI-induced disruptions in the economic sphere. In the third section, we discuss the uncertainties surrounding the nature and level of the impacts as well as the broader context. The fourth section reviews what we can learn from the bigger historical picture of technological progress. The fifth section distills the critical role of government policy in managing the effects of technological progress and enabling the benefits of innovation to be widely shared. The sixth section analyzes how our global system of governance needs to be updated to allow developing countries to maximize the benefits and minimize the costs of advances in AI and other digital technologies.

2 Downside Risks of Technological Progress

For much of the past half-century, there was a presumption that technological progress would benefit all — embodied by the trickle-down dogma that has characterized neoliberalism. This presumption was supported neither by theory nor evidence. Indeed, economic theory has always held that advances in technology may create winners and losers and do not necessarily benefit all. The experiences of the US and some other high-income countries over the past half-century have shown such worries to be justified: most of the benefits of growth have gone to those at the top, and large proportions of those at the bottom of the income distribution have seen their incomes decrease in real terms.

Over the following pages, we introduce a few analytic results that describe why advances in AI may be both labor-saving or resource-saving and may make some developing countries worse off. Some readers may wish to skip the math but will hopefully find the surrounding discussion useful.

In the context of a competitive economy, we can think of technological progress as shifting the isoquants that capture what inputs the economy needs to produce a given amount of output. In the left panel of Figure 1, we illustrate progress in an economy in which output is produced by combining capital K and




Figure 1: Isoquants and Factor Price Frontier

labor L to produce a given amount; we can either use more capital or more labor. Progress shifts the isoquants inwards, capturing that we need fewer inputs to produce a given output. (In an economy with multiple goods, we could also capture technological progress as moving out the production possibility frontier, which reflects how a given factor endowment can be used to produce different goods: progress means that one can get more of any output for given factor inputs.) But this increase in production capabilities does not in general tell us how the gains from progress will be distributed.

For the question of distribution, a concept that is very useful is the factor price frontier, which is a dual to the described isoquants: it specifies the highest possible return to one factor given the returns to other factors. Technological progress moves out the factor price frontier. Intuitively, this implies that every factor *could* be made better off. In the right panel of Figure 1, we illustrate the factor price frontier in a capital-labor economy with labor and capital that earn wage w and rental rate r . If the economy operates at high capital-labor ratios (top-left), wages are high and the rental rate is low; if the economy operates at lower capital-labor ratios (bottom-right), wages are low and the rental rate is high. Technological progress shifts out the factor price frontier, capturing that all factors could earn more.

On the factor price frontier, there is one particular point which represents the competitive equilibrium for given factor endowments, illustrated by point E_0 in the figure. If technological progress is Hicks-neutral, then the competitive equilibrium will make all factors better off in equal proportion, as in point E_1 . However, this is a knife-edge case. More broadly, the competitive equilibrium may well pick out a point on the factor price frontier where some factor owners gain more than others, or indeed one where all the gains plus some more go to one group, and some of the other groups are worse off, indicated by point E_2 . There is no assurance that technological progress will make everybody better off in the competitive equilibrium, even though there is sufficient output that everybody *could* be made better off. The appendix contains derivations for an example using a Cobb-Douglas production technology.

There is a simple criterion for evaluating whether innovation will result in wages going up (as in E_1) or down (as in E_2): whether at the original wage (w_0), the demand for labor increases or decreases. Much

of the analysis below is an enquiry into the kinds of changes in technology that at pre-existing wages increase or decrease labor demand. By the same token, the factor price frontier tells us that *if there is an unambiguous improvement in technology*, i.e. one in which the factor price frontier moves unambiguously out, then in the long run, workers are always better off, as long as the “natural rate of interest” is unchanged. (We will discuss the conditions under which the natural rate of interest remains unchanged below.)

The distributive effects of innovations can be viewed as generating quasi-rents, as we showed in Korinek and Stiglitz (2019): the winners of progress (e.g. capitalists or skilled workers) experience gains without necessarily having contributed to the innovation, whereas the losers experience losses. For example, automation may lower the demand for labor and wages but lead to a corresponding increase in the return to capital that is in the nature of a quasi-rent. It wasn't that capital owners did anything to receive these additional returns. That, in turn, has a big implication: governments can capture some of the quasi-rents by taxing the winners and redistributing the revenue; and given the nature of the gains, governments may even be able to raise taxes in ways which have no or limited distortionary effects. This makes it possible that the benefits of AI could lead to a Pareto improvement, i.e. that everybody could be better off with “managed” technological progress.

However, there is a big difference between looking at the impacts of AI within a single country and from a global perspective. When the benefits are experienced in one country and the costs are borne in another, a Pareto improvement would require that the winners compensate the losers across national boundaries. We don't have a global governance system that would make it possible to conduct such compensatory transfers and redistribute across borders, and even if we did, it would be difficult to muster the political will in a world where governments are primarily focused on the welfare of their own citizens.

As a result, technological progress will in general create winners and losers at the country level, reflected in improvements and deteriorations of countries' terms of trade. When innovations have significant effects on the economy, these cross-border effects may potentially be very large. In the following, we will analyze several of the specific forms of progress that the AI revolution and related automation technologies are likely to induce, with particular focus on how they may hurt developing countries.

2.1 Resource-Saving Technological Progress

A type of technological progress that is of great concern to some developing countries is natural resource-saving technological progress. We start our analysis with this type of progress because it allows us to lay out the basic structure of factor-saving progress in particularly simple terms. AI and other digital technologies have often been praised for their potential to produce more output with fewer natural resources. For instance, they may help us to reduce the demand for depletable natural resources and lower carbon emissions.³ Although it has gotten less attention than labor-saving progress so far, such resource-saving innovations may have particularly adverse distributional effects on those developing countries that have a strong comparative advantage in natural resources, and that have specialized in exporting them.

³ Examples include algorithms that optimize efficiency in data centers or that make transportation networks more efficient. Technologies that enable tele-work also reduce the carbon footprint of workers. On the other hand, the data centers running cutting-edge AI applications are quite energy hungry and may lead to increases in demand for electricity and depletable natural resources; still, on net, it is likely that the demand for energy and natural resources will decrease.

Consider a simple model of a resource-exporting country that has an endowment of N units of natural resources, which it exports in exchange for consumption goods. Advanced economies import the natural resources and combine them with skilled labor L to produce consumption goods according to the Cobb-Douglas production function

$$Y = F(L, N) = L^\gamma N^{1-\gamma}$$

where Y is world output. In competitive markets, this production function results in factor shares γ for skilled labor and $1 - \gamma$ for natural resources.⁴ One interpretation of the described production function – in the spirit of Zeira (1998) and Acemoglu and Autor (2011) – is that final goods are produced using a series of tasks that are combined in Cobb-Douglas fashion, where a fraction γ of the tasks employ skilled labor and the remaining fraction $1 - \gamma$ relies on natural resources.

In our simple model, the factor share $1 - \gamma$ of natural resources is also the income share of the resource-exporting country in world output. If we use the price of final output as the numeraire, then the marginal product of natural resources $F_N = (1 - \gamma) \cdot Y/N$ simultaneously represents the terms-of-trade (i.e. the price of export goods over the price of import goods) of the resource-exporting country.

Now consider a technological innovation that allows producers to economize on the use of natural resources and instead rely more on skilled labor, resulting in a change to the Cobb-Douglas coefficient in our production function of $d\gamma > 0$. In the task-based interpretation of our framework, we could view this as a fraction $d\gamma$ of the tasks relying on natural resources being replaced by skilled labor. Producers have a choice between the old and the new technology so $Y = \max \{F(N, L), F'(N, L)\}$ where we denote the new production function with a prime. They find it optimal to deploy the new technology when natural resources are sufficiently scarce, as captured by the following lemma:

Lemma (Use of Natural Resource-Saving Innovation): *Producers will find it optimal to use the new technology described by $d\gamma > 0$ if natural resources are sufficiently scarce, i.e. if the world economy's natural-resource/skilled-labor ratio is below a threshold,*

$$\frac{N}{L} < \hat{n}$$

Given the units in which the inputs in the production function above are denoted, the threshold is $\hat{n} = 1$ in our example.

Figure 2 illustrates the isoquants of the old and new technology F and F' for a discrete change $\Delta\gamma > 0$. The two isoquants intersect at a threshold point T , which corresponds to the natural resource/labor ratio \hat{n} at which a producer finds it optimal to switch technology. The dotted ray from the origin through point T separates the factor space into two areas: if the natural resource intensity of the world economy is greater than the threshold (corresponding to the area above the ray), the old production technology is more efficient – the corresponding isoquant reflects that within this area, the old technology requires fewer inputs. If the resource intensity is below the threshold, the new technology is more efficient and requires fewer inputs to produce a given amount of output. The economy's overall production possibilities are given by the envelope, i.e. the convex hull, of the two isoquants, indicated by a dashed line, since it is optimal for producers to use a mix of the two technologies over the indicated segment. That means that

⁴ It is straightforward to expand the model to capture imperfect competition.




Figure 2: Isoquants for change in technology

the new technology will begin to be used at a resource labor ratio that is somewhat greater than \hat{n} . (For infinitesimal changes in γ , the length of this dotted segment is, naturally, infinitesimally small.)

If resources are sufficiently scarce so that the new technology is used, $\frac{N}{L} < \hat{n}$, then overall output under the new technology rises by

$$\frac{dY}{d\gamma} = Y \cdot \ln \frac{L}{N} > 0$$

In that case, the income share of skilled labor rises by $d\gamma > 0$, and the overall returns to skilled labor $LF_L = \gamma Y$ increase since both the share earned by workers and overall output rises. However, the increase in world output does not necessarily imply that all factor owners are better off. The income share of natural resources declines by $d\gamma$, and the overall effect of the described technological progress on the marginal product of natural resources is

$$\frac{dF_N}{d\gamma} = \frac{d[(1-\gamma)Y/N]}{d\gamma} = -\frac{Y}{N} + \frac{1-\gamma}{N} \frac{dY}{d\gamma} = \left[(1-\gamma) \ln \frac{L}{N} - 1 \right] \frac{Y}{N}$$

The square brackets after the last equality reflect two competing effects: the first term captures what we may call a productivity effect – world output rises since the innovation relaxes the constraint posed by the limited availability of natural resources; the term is positive whenever the new technology is used, i.e. whenever natural resources are relatively scarce, as observed in the lemma above. The second term captures what we may call a displacement effect and is always negative – the relative share of world output earned by natural resources and thus by the developing country declines. The productivity effect will only dominate if the scarcity of natural resources is really severe, i.e. if $\frac{N}{L} < e^{-1/(1-\gamma)}$, otherwise the displacement effect dominates.

Hicks (1932) defined technological progress as factor-saving when it reduces demand for a factor at given market prices. In our setup, this is equivalent to reducing the marginal product of a factor given the available factor supplies. We can therefore summarize our results in the following proposition:

Proposition 1 (Natural Resource-Saving Innovation): (i) *An innovation captured by $d\gamma > 0$ is natural resource-saving if and only if*

$$\frac{N}{L} > n^* = e^{-\frac{1}{1-\gamma}}$$

(ii) *If the condition is met, the innovation reduces the terms-of-trade F_N and the total income $N \cdot F_N$ of the resource-exporting developing country, making the country worse off in absolute terms.*

A tangible example would be oil-exporting countries that rely on their export revenue to buy food and other basic essentials. If they suffer large terms-of-trade losses, the consequences could be dire. Many oil-exporting countries have already experienced developmental challenges while being resource-rich. Resource-saving AI, while saving the planet, would make them resource-poor countries that still experience the same developmental challenges. This would really test the global community. More generally, the impact on exporters of different types of natural resources may be quite different – for example, oil exporters will fare very differently from exporters of rare earth metals.

Let us now consider the effects of an innovation $d\gamma > 0$ on another country i with endowments of skilled labor and natural resources of (L^i, N^i) . We already know that the country will be better off if its natural-resource intensity corresponds to the global average N/L since the innovation makes the world as a whole better off. However, in the more general case that factor endowments are distributed asymmetrically around the world, the benefit of the innovation to each country will depend on its relative endowment of skilled labor. By evaluating the impact on the overall income of country i , $d[F_L L^i + N^i F_N]/d\gamma$, we find:

Proposition 2 (Threshold of Natural-Resource/Skilled-Labor Intensity): *A natural resource-saving innovation $d\gamma > 0$ will make country i worse off if and only if its natural-resource intensity is greater than a critical threshold \tilde{n} that depends on the world economy's natural-resource intensity $n = N/L$,*

$$\frac{N^i}{L^i} > \tilde{n} = n \cdot \frac{1 - \gamma \ln n}{1 + (1 - \gamma) \ln n}$$

Conversely, any country with lower natural-resource intensity will be better off from the innovation.

In the setup above, we have outlined the effects of natural resource-saving technological progress in a simple two-factor setting. Similar results hold for production functions with additional factors, e.g. the specification $Y = F(K, L, N)$.

2.2 Labor-Saving Technological Progress

Many are concerned that AI may be labor-saving, or at least unskilled-labor saving, at the global level. Labor-saving progress means that at existing factor prices, demand for labor will go down. If this occurs, equilibrium wages will go down and workers will be worse off. In the case of unskilled labor-saving progress, the same will be true for the equilibrium wages and incomes of unskilled workers.

Over the past half-century, the US and many other countries have experienced technological progress that was biased against labor and reduced the labor share of national income (Karbarbounis and Neiman, 2013), although we note that the decline in the labor share of some countries (including the US) was also heavily influenced by the weakening of the bargaining power of workers, e.g. by changes in labor legislation and rules, unionization and globalization. And there are indications that progress may even have been labor-saving for some, reducing the real incomes of workers with lower levels of education, in particular workers without college degrees. For example, Autor et al (2003) observe that machines are becoming more and more efficient at performing routine functions that have traditionally been performed by unskilled labor, and this has dragged down wages for unskilled workers. AI is likely to continue the trend (see e.g. Berg et al., 2018; Korinek and Stiglitz, 2019).

Although labor-saving technological progress would make the world as a whole richer, it would hit developing countries that have a comparative advantage in cheap labor especially hard. If worldwide demand for labor, or for unskilled labor, declines, such countries would experience a significant deterioration in their terms of trade and lose a substantial fraction of their export income, to the point that it may make entire countries on net worse off. See also Alonso et al. (2020) for a quantitative evaluation and Faber (2021) for empirical evidence.

Consider a model of a developing country that has an endowment of L^d units of raw labor, which it employs to produce labor-intensive intermediate goods that are exported in exchange for consumption goods. Advanced economies import the labor-intensive intermediate goods and combine them with capital K to produce consumption goods according to the Cobb-Douglas production function

$$Y = F(K, L) = K^\alpha L^{1-\alpha}$$

where capital K stands for a wide range of types of capital, including e.g. the human capital inherent in skilled labor. The rest of the world also has an endowment of labor L^r , and for simplicity we assume that raw labor is converted into the labor-intensive intermediate goods one-for-one, giving rise to a total supply of $L = L^r + L^d$. In competitive markets, this production function results in factor shares α for all types of capital and $1 - \alpha$ for raw labor. Using the price of final output as the numeraire, the marginal product of labor $F_L = (1 - \alpha) \cdot Y/L$ represents the terms-of-trade of the developing country.

Consider a technological innovation that increases the Cobb-Douglas coefficient on capital by $d\alpha$ and assume that capital is sufficiently abundant, i.e. $k > \hat{k} = 1$, so producers find it optimal to deploy the new technology. In that case, the innovation increases total output but reduces the share of income earned by labor. In analogy to the resource-saving innovation of the previous section, we find the following results:

Proposition 3 (Labor-Saving Innovation): (i) *An innovation captured by $d\alpha > 0$ is labor-saving if and only if*

$$\frac{K}{L} < \tilde{k} := e^{1/(1-\alpha)}$$

(ii) *A labor-saving innovation reduces the terms-of-trade F_L and the total income $L^d \cdot F_L$ of the developing country, making the country worse off in absolute terms.*

The condition in part (i) of the proposition is satisfied as long as the capital intensity of the world economy is not too high. Otherwise the productivity effect from saving on scarce labor would outweigh the displacement effect of labor in production.

A tangible example would be a country that owns little capital and exports simple manufacturing goods that are produced using largely unskilled labor, e.g. textiles, in exchange for imports of food. If the population of the country is living close to subsistence levels before the innovation, then the terms-of-trade losses associated with the innovation may well push the incomes of unskilled workers below subsistence levels, resulting in widespread misery.

As in the previous section, we can also find a threshold capital intensity below which a country i is worse off from the innovation.

Proposition 4 (Threshold for Capital Intensity): *A labor-saving innovation $d\alpha > 0$ will make country i worse off if and only if its capital intensity K^i/L^i is less than a critical threshold,*

$$\frac{K^i}{L^i} < \bar{k} = k \cdot \frac{1 - (1 - \alpha) \ln k}{1 + \alpha \cdot \ln k}$$

However, the results of our simple two-factor model only capture the short-run effects. The described innovation may be labor-saving, but there are two additional considerations that matter for the long run: First, in a dynamic setting, much depends on how technology will evolve in the future. If the new technology is currently more productive but improves at a lower rate in the future, the older technology will eventually predominate. If the new technology implies further progress that saves on labor, the effect on wages may be exacerbated. Second, the greater productivity of capital will induce capital owners to accumulate more of it, which may in turn raise wages. We analyze these latter effects in the following.

Long-run Effects of Labor Displacement

The long-run effects of technological change that displaces labor heavily depend on how much capital the economy accumulates, which depends on the consumption growth and time preferences of economic agents. Many of our intertemporal models assume additively separable time preferences. This assumption greatly simplifies the analysis of intertemporal problems but is by no means natural or general.⁵ With time separability, the long run interest rate is determined by $\rho + \eta g$, where ρ is the pure rate of time preference together, g is the rate of consumption growth, and η is the (absolute value) of the elasticity of marginal utility. Assume that the rate of technological change (and hence the rate of consumption growth) is the same for the new and old technologies. Then the long-run equilibrium return on capital (denoted by an asterisk) is also constant at $F_K^* = \rho + \delta + \eta g$, where δ is the depreciation rate, and is unaffected by the innovation. In that case, an analysis of the effect on the factor price frontier is convenient, for we know that from a long-run perspective, any innovation shifts out the factor price frontier. By implication, the innovation must shift out the factor price frontier at $F_K^* = \rho + \delta + \eta g$, increasing the income that is collectively earned by the other factors. In the analytic model above in which labor is the only factor that cannot be accumulated, this implies that labor will necessarily be better off in the long run (see also Caselli and Manning, 2019).⁶

⁵ Alternative representations include, for example, the recursive utility setup of Koopmans (1960) or, more recently, the non-homothetic preferences of Straub (2020).

⁶ In the case of non-homothetic/non-separable utility functions, it is possible that the long run interest rate might increase with new technology, and in that case, workers could be worse off *even in the long run* even in a model that encompasses only capital and labor. Similarly, even with time-separable utility functions with constant elasticity of marginal utility, it is possible that the new technology has a higher *rate* of technological change (making it all the more superior to the old technology), which would increase the rate of interest and the equilibrium return on capital. At the same time, however, this would also raise the rate of increase of wages.

When there are multiple factors in addition to capital, then the sum of the returns on these other factors will always increase in a model with fixed rate of time preference, but some factors may well be worse off, even in the long run. For example, in a model with skilled labor, unskilled labor and capital, skilled labor may appropriate all the gains from technological progress and capital accumulation plus some more, and unskilled labor may be worse off. Similarly, in a model with capital, labor, and a third scarce factor that cannot be accumulated, like land or natural resources, the third factor may appropriate the gains from technological progress and capital accumulation, and labor overall may be worse off.

Consider a world economy in which output is produced from three factors using a production function of the form

$$F(K, L, N) = AK^\alpha N^\gamma L^{1-\alpha-\gamma}$$

The competitive wage rate in this economy is given by $F_L = (1 - \alpha - \gamma)Y/L$.

The effects of a technological change that increases the use of capital at the expense of labor (“automation”) may be captured by an infinitesimal change $d\alpha$. As before, producers will only find it desirable to employ this new technology if the capital intensity of the economy is above a threshold $k > \hat{k} = 1$.

In the short run, the effect of a marginal change in α on labor income is

$$\frac{dF_L L}{d\alpha} = \left[(1 - \alpha - \gamma) \ln \frac{K}{L} - 1 \right] Y$$

The technological change is labor-saving if the capital/labor ratio or capital intensity $k = K/L$ in the world economy satisfies

$$k < \tilde{k} := e^{1/(1-\alpha-\gamma)}$$

In that case, developing countries that are only endowed with labor will be worse off from the innovation. Akin to our derivation above, we could also derive a threshold level \bar{k} such that countries with a capital intensity lower than \bar{k} would be worse off. Conversely, if the capital intensity of the world economy is above this threshold, capital is sufficiently abundant that the productivity effect from using the highly abundant resource dominates the displacement effect.

Assuming it is profitable to use the new technology (i.e. $k > \hat{k}$), it is always natural resource-using in the short run since natural resources benefit from the productivity effect and there is no displacement of natural resources when α changes,

$$\frac{dF_N N}{d\alpha} = \gamma \frac{dY}{d\alpha} = \gamma Y \ln \frac{K}{L}$$

By contrast, in the long run, the capital stock has time to adjust to its long-run equilibrium level at which the marginal return to capital is determined by the rate of time preference, the depreciation rate, and consumption growth so $F_K^* = \rho + \delta + \eta g$, where we denote long-run values by asterisks. Then the long-run capital stock and output satisfy

$$K^* = \left(\frac{\alpha A}{F_K^*} \right)^{\frac{1}{1-\alpha}} N^{\frac{\gamma}{1-\alpha}} L^{\frac{1-\alpha-\gamma}{1-\alpha}}$$

$$Y^* = F(K^*, L, N) = A \left(\frac{\alpha A}{F_K^*} \right)^{\frac{\alpha}{1-\alpha}} N^{\frac{\gamma}{1-\alpha}} L^{\frac{1-\alpha-\gamma}{1-\alpha}}$$

The long-run effect of a marginal change in α on labor income is

$$\frac{dF_L^* L}{d\alpha} = \left\{ \frac{1-\alpha-\gamma}{(1-\alpha)^2} \left[\ln \frac{\alpha A}{F_K^*} + \gamma \ln \frac{N}{L} - \frac{\gamma(1-\alpha)}{1-\alpha-\gamma} \right] \right\} Y^*$$

This immediately yields the following result:

Proposition 5 (Possibility of Labor-Saving Innovation in the Long Run): *An innovation captured by $d\alpha > 0$ is long-run labor-saving if the natural resource/labor ratio in the economy satisfies*

$$\frac{N}{L} < \tilde{n}^* := \left[\frac{(F_K^*) \cdot e^{\frac{\gamma(1-\alpha)}{1-\alpha-\gamma}}}{\alpha A} \right]^{1/\gamma}$$

Intuitively, the technological change replaces labor by capital, which is accumulated using a combination of natural resources and labor. If natural resources are sufficiently abundant, then the technological change is labor-using in the long run. To see this, it may also be helpful to consider the limit case that natural resources are infinitely abundant so that they do not constrain production: in that case, we return to a two-factor model, in which case we know that labor has to be better off. Proposition 5 shows that if natural resources are sufficiently abundant, the same result holds. Conversely, if the condition in the proposition is satisfied, natural resources are sufficiently scarce that the described technological change is labor-saving even in the long run.

For developing countries that specialize in exporting natural resources, an analogous condition can be derived under which technological progress that displaces natural resources with capital would be natural resource-saving in the long run.

Let us also consider three-factor production function that allow for more general elasticities of substitution than the unitary elasticity encapsulated by the Cobb-Douglas case. One extreme is that capital and labor may enter production as perfect substitutes, being combined with natural resources in CES fashion. This would capture fully worker-replacing technological change, as we analyze e.g. in Korinek and Stiglitz (2021). When capital and labor are perfect substitutes, the competitive return on capital pins down the returns to labor. In the short run, an increase in capital will then be clearly labor-saving. In the long-run, the return on labor is pinned down by the long-run return to capital F_K^* so that all the gains from productivity growth will go to natural resources. Conversely if capital and natural resources are perfect substitutes, short-run increases in capital will be natural resource-saving, and in the long run, the return to natural resources is pinned down by the long-run return to capital F_K^* , implying that all the long-run benefits from technological progress would go to labor. More general conditions can be derived on the basis of the degree of substitutability/complementarity of capital with different types of labor and natural resources.

More broadly, development economics has long debated whether there is convergence among countries, as suggested by neoclassical theory, i.e. whether poorer countries will eventually catch up to the standard of living of rich countries (see e.g. Mankiw et al., 1992; Barro and Sala-i-Martin, 1992). The argument is

typically based on the notion that over time, (i) countries will catch up to the world technological frontier and (ii) countries will accumulate capital (including human capital) to reach the same capital intensity as the most advanced countries, as suggested e.g. by the Solow growth model. While these forces toward convergence and equalization of incomes certainly exist and may be the dominant forces over some time periods, our analysis highlights an important opposing force: the factor endowments of developing countries, in particular endowments of unskilled labor and natural resources, may be devalued by future technological progress. The convergence in standards of living between developing countries and developed that has marked the past half century would then be arrested, perhaps even reversed. This would present great challenges to domestic policy within developing countries. In many parts of the world, inequalities within developing countries are greater than in developed. AI may exacerbate those inequalities – and developing countries typically lack the institutional capacities to counteract them.

Possibility of Labor-Using Innovation

However, while it is possible, perhaps likely, that technological progress may hurt labor, it is also possible that technological progress may enhance the prospects of labor, i.e. be labor-using. One example is what is called IA (Intelligence Assistance) rather than AI. Past innovations that fell into this category included telescopes and microscopes that enabled us to see what the naked eye could not; thermometers that enabled us to measure what we otherwise had not been able to detect precisely, or tests, that tell us whether we are carrying coronavirus. These IA innovations enable us to focus human attention where one should.

Frequently, automation technologies affect particular *tasks* but not jobs, which consist of multiple tasks (see e.g. Acemoglu and Autor, 2011). In the past, automation was associated with highly routine tasks. Going forward, automation is increasingly also affecting tasks that require higher levels of skill. For example, a doctor is engaged in diagnosis but also in explaining the diagnosis to the patient. AI may do a better job in diagnosis – for example, in radiology – but it may not quite replace the doctor in communicating with the patient yet. Of course, both AI and IA imply extensive restructuring of the economy.

Driverless trucks—always a few years in the future—provide another example. Truck driving provides significant employment opportunities for men with only a high school education. There is thus understandably concern for the disruption that driverless trucks might bring about. But truck drivers also perform a number of related tasks: they fill orders, load and unload, monitor the truck etc. Some of these activities may also be automated soon, but some of them may not. More generally, most jobs have multiple dimensions and consist of multiple tasks. Some parts will be automated soon but other parts won't for some time. Workers will be able to devote more attention to those parts that are not.

The central concern of this paper remains: there may be a reduction in the demand for labor, especially for unskilled labor, even in the long run. We will further evaluate whether or not these fears are justified below in Section 3.

2.3 Information, Digital Monopolies and Superstars

The rise of AI and other information technologies may also lead to greater and greater concentrations of market power. As a result, the economy may move to an equilibrium that is more distorted by market power, with greater rents for dominant firms. Actors with sufficient market power who are aware of the

impact of their decisions on market prices (including factor prices) are also likely to use that power deliberately to advantage themselves relative to others. The resulting distortions may offset part or all of the benefits of innovation. And this may exacerbate the adverse distributive effects of labor-saving or resource-saving innovation.

While economists begin with the competitive model in the back of our minds – this is what we have been indoctrinated in – that model becomes less and less appropriate as we move to an economy which is dominated by AI. It is hard to conceive of an AI economy being competitive, or at least well-described by the competitive equilibrium model.

There are several reasons why advances in AI intensify market power. First, AI is an information good, and information goods are different from other goods in that they are non-rivalrous – they can be used at close-to-zero marginal cost, implying that a single firm can serve a very large market. Moreover, the creation of AI programs typically involves high sunk costs and/or fixed costs – in a private market, firms need to earn monopoly rents to recoup these costs. Moreover, even small sunk costs may result in markets not being contestable, i.e. there could be sustained rents and profits. In addition, AI applications and platforms typically involve significant network externalities. Some of these arise because firms accumulate vast amounts of data that allow them to train their algorithms better than those of the competition. All of these effects create large barriers to entry and a tendency towards what is sometimes called “superstar” effects (see e.g. Korinek and Ng, 2019).

Speaking of such firms as “superstars” makes it sound very positive, but perhaps it would be more appropriate to refer to them as “super monopolies.” Some of the literature identifies a growing number of “superstar firms” in the economy that are “super profitable” (see e.g. Autor et al., 2020). However, rather than reflecting “super productive” technology, much of these profits may arise from the exercise of monopoly power that is derived from the nature of these information technologies. For example, in the US, a large fraction of the gains in the stock market over the past decade have been concentrated in digital giants, to an important extent driven by their market power. Moreover, algorithmic advances have also enabled digital firms to extract more consumer surplus through discriminatory pricing.

The superstar effects and the monopoly effects that we discussed here are likely to play out not only at a company level but also at a country level, and are likely to be particularly severe with AI. They may be exacerbated by agglomeration economies associated with R&D in AI. There is a risk that those countries that lead in the advancement in AI may reap all the benefits, becoming “superstar countries” and reaping all the rents associated with the development of AI. The rest of the world, and in particular most developing and emerging economies, may be left behind. Moreover, the monopolization of knowledge may also impede the catching up process. In the past, advances in technology were driven to an important extent by basic research that was financed by governments in high-income countries and that was freely available to all—including to developing countries. This may change with AI.

Data as a new resource for developing countries

Some suggest a silver lining for developing countries may be that present-day machine learning technologies are very data hungry, and more diverse data contain more information – but while that might generate some income, selling data seems an unlikely route for the developing countries to make up their lost income, and there is no evidence suggesting that it could. Indeed, the marginal return to more

diverse data may be limited. Future advances in machine learning algorithms may make them less reliant on large quantities of data; more specific, tailored data is what will be needed.

2.4 Misguided Technological Progress

Economic theory has illuminated why the nature (e.g. the factor bias) of innovation may not be welfare maximizing. Much of economics takes the factor bias of technological change as exogenously given, and the standard economic welfare theorems assert the efficiency of competitive market economies for a given level of technology. However, the direction and rate of technological progress are themselves economic decisions, as emphasized by an important literature that goes back half a century on induced innovation (e.g. Kennedy, 1964; von Weizsäcker, 1966; Samuelson, 1965; and more recently Acemoglu, 1998, 2002). There is no analogue of the welfare theorems for innovation: markets on their own will not in general be efficient either in the level or direction (nature) of innovative activity and technological change. This calls for policy to actively steer technological progress, as we will discuss further below.

One of the reasons for the inefficiency of innovation in a market economy is that competitive markets are inconsistent with innovation. Competitive markets cannot generate the rents necessary to finance research and development. Innovation always confers at least temporary monopoly power upon innovators. But monopoly power is inherently inconsistent with perfect competition, and leads to inefficiently low quantities of output.

Another important reason for the inefficiency of innovation in a market economy is the pervasive micro- and macroeconomic externalities associated with innovation (see e.g. Stiglitz and Greenwald, 2014a; Korinek and Stiglitz, 2019). In economies with incomplete markets or imperfect information, actions by one party have effects on others which they don't take into account in their decision making. There are also always some knowledge spill-overs.

Even more importantly, markets inherently do not care about income distribution. Market forces may drive economic decisions towards efficiency but will not give any consideration towards whether there are winners and losers. Decisions about what kinds of innovation and how much innovation to adopt have distributive effects, which are of direct concern themselves but also have efficiency implications.

There are some self-correcting forces: for example, if labor is getting cheaper, innovators face smaller incentives to save on labor, providing a corrective mechanism within the market economy to an ever-decreasing share of labor, but this mechanism no longer works when wages are set by efficiency wage considerations or reach subsistence levels.

In Korinek and Stiglitz (2020b), we show that in the presence of constraints on redistribution, policy can improve welfare by steering innovation to take into account its distributive implications.

What is most relevant for developing countries is that these distributive implications go across borders, and so decisions made in one country have effects on other countries that the innovating country and the innovators within that country have no incentive to take into account.

Even if markets were efficient in the choice of technology for the conditions of the country in which the innovation occurs, those conditions are markedly different from the conditions in developing countries. In developing countries, a key question is *adopting* appropriate technologies. But the same kind of analysis which argues for the need for government intervention in steering technological innovation also provides

arguments for intervention in steering technology adoption. This is especially so if, after the initial adoption of technology from abroad, there is adaptation to local circumstances, and the benefits and costs of the technological evolution are not fully appropriated, e.g. in the process of learning by doing. (See, e.g. Atkinson and Stiglitz, 1969). These concerns have long been at the center of concern of industrial policy.

2.5 Broader Harms Associated with AI

There are also a number of broader harms associated with AI that have recently received a lot of attention—the ways in which new technology can affect security (including cybersecurity), privacy, incitement to “bad” behavior, including through hate speech, political manipulation, and, in the economic arena, price discrimination, sometimes exacerbating pre-existing societal divides.

While these matters affect both high-income and developing countries, we are concerned that the international community may address them in a way that does not reflect the priorities and needs of developing countries. The world is beginning to discuss an appropriate regulatory regime and a global set of rules to address these potential harms. It is unclear whether developing countries and emerging markets will be sufficiently represented at the table when these discussions take place. In fact, many of the standards, rules and regulations are likely to be set by high-income countries and China, even though the impacts may be larger, and potentially different, on developing countries and emerging markets.

Moreover, the institutional capacity of developing countries to counter these harms may be more limited—especially when facing off against the technology giants. Weaker institutional foundations may make some countries more prone to abuses of autocratic and totalitarian leaders using mis-/disinformation and surveillance technologies. Less educated populations may suffer more from the consequences of mis-/disinformation, such as those associated with the anti-vax movement.

3 Evaluating the Downside Risks

3.1 Uncertainty about the Pace and Scale of Progress

The impact of technological change depends heavily on its pace and scale. If it occurs slowly, there is time to adjust. If automation is limited to a few tasks or sectors at a time, the impacts will be limited.

However, there is a lot of uncertainty about the pace of change and the magnitude of the coming disruption, even among experts in this area. There are some economists, for example Gordon (2016), who assert that we are not in an era of unprecedented innovation, and that economic growth will be less rapid in the future than it has been over the past century. Gordon argues that indoor toilets and electricity had far bigger consequences on people’s standards of living than more recent innovations.

Another view is that AI is a truly transformative technology, a *General-Purpose Technology* that has the potential to revolutionize every sector of the economy (see e.g. Trajtenberg, 2019). Like steam engines or electricity in previous technological revolutions, this view predicts that AI will lead to significant productivity gains and structural changes across the entire economy.

An even more radical perspective that goes back to John von Neumann is that Artificial Intelligence may eventually advance to the point where AIs can do research, design better versions of themselves and thereby recursively self-improve. This could give rise to accelerating technological progress and, in the

words of von Neumann, “the appearance of approaching some essential singularity in the history of the race beyond which human affairs, as we know them, could not continue” (see Ulam, 1958). This concept has been popularized by Good (1965), Vinge (1993) and Kurzweil (2005) and is now also discussed in a number of economics papers (see e.g. Nordhaus, 2015; Aghion et al., 2019; Trammell and Korinek, 2021).

This perspective emphasizes that AI-powered machines are not only physically stronger than humans, and they can not only process information better and faster than humans, but in an increasing number of domains, they can also learn better and faster than humans. From this perspective, AI may be much more disruptive than a “mere” general purpose technology; AI programs are already becoming agents of their own, may increasingly replace human jobs, and may ultimately make the factor labor redundant. If machines can engage in all the tasks that have traditionally been performed by labor, and if they can do so at ever lower cost, then traditional labor would become redundant, in the sense that the marginal product of (at least low skilled) human labor no longer covers the subsistence cost necessary to keep a human alive (Korinek and Stiglitz, 2020b). This would represent the extreme case of labor saving innovation: it is literally labor-replacing innovation – employing labor would become a strictly dominated technology.

There are lots of studies for how many jobs may be replaced by automation and AI in coming decades, typically based on job- or task-level data. The predictions in these studies vary widely, ranging from a relatively small percentage of 9% of all jobs (OECD, 2016) to an estimate of 20–25% (Bain and Company, 2018) and almost half by Frey and Osborne (2017) and McKinsey Global Institute (2017). Even the lower numbers suggest a significant effect, especially because the impact may be concentrated. Job losses may first be focused on unskilled and routine jobs. Knowing what fraction of all jobs will be lost to AI therefore doesn’t necessarily provide a good metric of the impact on income distribution.

Applying our earlier insights on steering innovation to economic research, we should also steer our own research in directions where the expected social value added of economic analysis is greatest, where it has the highest welfare impact. As Yogi Berra said, forecasting is very difficult, especially the future. What we will discover in the future is in a fundamental sense unknowable. Sitting here today we can’t be sure which of the described scenarios will play out.

Even if they have a relatively low probability, we need to think particularly hard about events that will be highly disruptive to our society, to think through the consequences now, and to prepare for how we might ameliorate some of the more adverse effects. Extensive labor replacing innovation would be such an event. We don’t think that it is a low probability event, but even if we took the view that there is just a 10 percent probability that we would have significant labor-saving innovation, it makes sense to focus attention on such an event. A technological change that, say, changes the redistribution of income by two percent is not likely to be a source of societal turmoil. We have those types of events all the time, and we feel fairly confident that our economic and social system can absorb such a change. We should be thinking hardest about scenarios that pose the most adverse social impacts, so that we are better prepared to deal with them when they occur. Moreover, thinking about how to prepare for low probability scenarios that are highly adverse – for example in the context of managing the impact of radical technological change on income distribution – also teaches us valuable lessons for scenarios in which the impact is less stark.

3.2 The Productivity Puzzle: Are we really in an era of unprecedented innovation?

In relating the debate about the economic significance of AI-based innovation to recent economic data, we encounter a well-known puzzle: if we are really living in an era of significant technological disruption, why aren't the increases in innovation that we hear about showing up in our GDP data? This is analogous to the puzzle of missing productivity growth from computerization that Bob Solow described in the 1980s when the general purpose technology of the time – computers – spread throughout the economy (Solow, 1987). At the time, it took until the following decade for US national accounts to show a pickup in productivity growth.

Part of the explanation for the productivity puzzle is that there are long lags, as was the case for computerization. At present, AI is influential in a limited number of sectors, like inventing better ways of advertising. Even if AI is transforming advertising, this will not transform our overall standard of living. (In this particular case, it may actually lower overall efficiency, as it may undermine the price system by enabling pervasive discriminatory pricing.) Going forward, many sectors of the economy will require complementary investments and changes in processes and organization as well as new skills among their workers to take full advantage of AI (see e.g. Brynjolfsson et al., 2019).

Another part of the explanation is that the productivity puzzle also reflects a measurement problem. Many recent technologies may have led to increases in societal welfare that are not captured by GDP (see e.g. Brynjolfsson, 2020), for example because services are exchanged against “eyeballs,” i.e. users are exposed to advertisements instead of paying for services.

As this discussion illustrates, it is difficult to predict what impact AI will have going forward.

3.3 Putting AI in the Broader Context of Development

There are a number of other important factors that are relevant when it comes to managing the potential adverse effects of AI on developing countries in coming decades.

COVID-19

The COVID-19 pandemic has imposed an extra shadow cost on physical interaction with humans, and this provides strong incentives to accelerate the automation of jobs that require physical interaction (see e.g. Korinek and Stiglitz, 2020b). The resulting changes will have long-lasting effects on the economy, even after the pandemic is overcome. The new technologies that are introduced now will reduce the demand for labor worldwide for some time to come.

Population dynamics

Second, population dynamics will interact in important ways with the impact of labor-saving or -replacing technology (see e.g. Varian, 2020). In countries in which the working age population is growing rapidly, such as in many African countries, lots of new jobs will have to be created to maintain a given employment rate; however, young populations also generate significant demand for education, which in turn creates jobs.

Conversely, in countries in which the working age population is declining, such as China, the impact of job automation on the workforce is mitigated as workers that are replaced by technological progress can simply retire. Moreover, aging populations create large service sector needs, particularly in healthcare. Many of these service sector jobs are unlikely to be replaced by automation or AI in the near future.

The Green Transition

Third, an important force affecting developing countries in coming decades is the threat of global warming, which calls for significant public policy interventions to facilitate the Green Transition, i.e. the transition away from an economy that is dependent on fossil fuels to one that is more environmentally sustainable and relies more on renewable energy. We will focus on the close similarities between technological change and the Green Transition in the following section.

3.4 Technological Change and the Green Transition

The Green Transition will require large investments to move our economies away from fossil fuels. Renewable energies require new infrastructure. Moreover, energy efficiency requires significant investment in housing and new transportation systems. AI won't replace workers required for much of those investments in the near future.

There are many similarities between the effects of AI and the Green Transition. Both involve large changes in relative prices and generate significant redistributions, and many developing countries will be strongly affected. On the one hand, without the investments necessary to save our planet, developing countries will experience some of the largest losses from global warming. On the other hand, the Green Transition is similar to resource-saving innovation and risks undermining the standard of living of oil-exporting countries.

There is an important complementarity between the Green Transition and AI: the Green Transition is likely to increase the demand for labor so as to offset some of the negative effects on labor demand of AI. Indeed, one might say that it would be a good thing that labor is replaced in many activities, including manufacturing, because this frees up labor to meet the demands of the Green Transition, enabling us to better address the challenges of climate change. We note that one can't simultaneously claim "our economy can't afford to mitigate climate change," i.e. that there are insufficient resources, and "we face a potential crisis with a surplus of labor and a surplus of savings." (Already before the 2008 crisis, many, including the Fed Chair Ben Bernanke, claimed that there was a savings glut. In the years after the 2008 crisis, many economists claimed that even at a zero-interest rate, there was an insufficiency of aggregate demand, i.e. a surplus of savings.)

However, we do face challenges in how to channel surplus resources into what is required for the Green Transition. Some of the investment will have to be public, and this will require additional fiscal resources. There are significant deficiencies in current tax systems around the world, ranging from insufficient progressivity or outright regressivity to pervasive tax loopholes, inadequate taxation of multinationals, lack of taxes on negative externalities such as environmental taxes, and lack of taxation of rents from fixed factors such as land and natural resources. These deficiencies are particularly stark in developing countries. However, they also mean that reforms to our tax systems should make it possible to raise the revenue required for public investments.

There may be other institutional constraints that make it difficult to raise the required resources, and to reallocate capital towards green investment in both the private and the public sector. While many sources of savings are long term (pension funds and sovereign wealth funds) and the investments needed for the Green Transition are long term, standing between are short term financial markets. We may need, for instance, new Green Development Banks to help finance private green transition, combined with better disclosure of risks associated with “brown” investments (i.e. ones that contribute to pollution) and changes in fiduciary standards, to move resources out of brown investments.

4 Lessons from Past Technological Transformations

To grasp the historical nature of what is going on, we need to place the advent of AI and related technologies into the broader history of technological progress. Humanity spent much of its history at a Malthusian stage. The Industrial Revolution started a bit over two centuries ago, and was but a blip in the history of mankind. The era of manufacturing-based export-led growth that enabled the East Asian Miracle stretched over the past half-century – only one quarter of the history of the Industrial Revolution. It is easily conceivable that we are now going into another era.

Many are far more sanguine than we are about the disruptive potential of AI. They point to the automobile and other innovations at the end of the nineteenth century. Jobs were lost, making buggy whips and horse carriages obsolete, but overall, labor demand increased and more jobs were created. Our analytical discussion made clear that there is no inherent reason that innovation has these effects. This time could well be different. Looking at the time before the Industrial Revolution and the early decades of the revolution itself serves as a reminder.

4.1 Pre-Industrial Revolution

Before the Industrial Revolution, innovation proceeded at a far slower pace than today. There were still many innovations, but the actual living standard of the vast majority of people was stagnant (Maddison, 2003). The interpretation of Malthus (1798) was that every time an innovation took place, the population started to grow and absorbed the surplus that was generated.

This pre-industrial state of affairs may be still relevant in the least developed countries and is particularly problematic in some African countries, where the death rate has been greatly reduced by medical innovations but reproductive rates have continued to be very high. The affected countries have been slow to go through the demographic transition that marked the rise of living standards in Asia. As a result, a number of countries are facing a difficult-to-manage explosion in population combined with stagnant living standards.

There is a risk that poor countries may see a return to Malthusian dynamics if technological progress undermines the source of their comparative advantage. Consider a country that exports manufacturing goods produced using cheap labor but is not very productive in agriculture, for example because of a shortage of land and a high population density. The country uses its export revenues to import food for workers in the manufacturing sector, granting them a living standard that is above subsistence levels. If a new technology produces the manufacturing goods more cheaply, the wages of the manufacturing workers will fall, and they may well fall below the subsistence cost of workers. If that is the case, the country may return to a Malthusian state of affairs in which an important part of the population suffers from hunger and deprivation. Increasing agricultural productivity may mitigate this dire state of affairs but

the question is, would they be sufficient to support a population that was previously supported by imported food?

4.2 Industrial Revolution

The Industrial Revolution marked the beginning of rapid growth in high-income countries, but not everyone in the affected countries participated in the resulting gains. At first, there was an increase in inequality. The Industrial Revolution thus provides us with a number of lessons that are very relevant today:

Innovation can be very disruptive

Even when an innovation ultimately proves to be beneficial for society at large, it can give rise to very large disruptions. The Industrial Revolution in high-income countries obviously transformed the world. Standards of living had been stagnant everywhere for thousands of years, and then suddenly they followed a J curve, dipping but ultimately going up by an order of magnitude. In the short run, there was a lot of social upheaval—Dickens' novels make it clear that not everyone prospered. In the UK, there were people who were living under much worse conditions in the cities of the mid-19th century than they had been in the rural areas prior to that. Looking at those people, the Industrial Revolution was not a Pareto improvement.

Collective action can mitigate the adverse effects

The onset of the industrial revolution posed many challenges that required collective action. However, it took time for societies to put in place the collective mechanisms to respond to these challenges. This is why the industrial revolution had significant negative effects on the masses for some time – even indicators such as life expectancy initially went down. Eventually, governments played a very important role in mitigating the adverse effects. One example is the problems posed by urbanization, including challenges in sanitation, environmental degradation, public health, infrastructure, and congestion. Another example is that education was a very important element in creating a productive workforce – education was therefore also in the interests of capitalists and received broad public support.

In areas such as labor legislation, unionization, and the creation of a social safety net, capitalists resisted changes. In high-income countries, these institutions were not created until the end of the 19th century and beginning of the 20th century. In the United States, the ready availability of land implied that labor was relatively scarce, limiting the extent to which labor could be exploited. Nonetheless, in the early years of the 20th century, labor was not doing very well. It was only dramatic events like the 1911 Triangle Shirtwaist Factory fire in New York City that led to labor legislation that really protected workers. In most high-income countries, labor legislation has been taken for granted in recent decades, but in 1900, it was not obvious if meaningful labor legislation would ever be enacted, and even today, in the United States there has been an erosion of protections provided by minimum wage and health and safety standards, either as a result of a failure to adapt the legislation to changing circumstances or inadequate enforcement (as in the context of Covid-19).

These reforms helped support the structural transformation that occurred with the rise of manufacturing, and they showed that equality and growth are complementary (see Ostry et al., 2019). At a basic level,

they were necessary to sustain social peace and democracy. And they ushered in what might be called an “Age of Labor.”

Most developing countries have not gone through this process yet. Moreover, this Age of Labor may not last forever. In the US, minimum wages have declined in real terms in recent decades, below the level of sixty years ago. Labor protections have come under attack, and many protections on hours and working conditions have been eviscerated. Advances in AI may further contribute to undermining labor’s bargaining position and thus these social protections. And in developing countries, they may do so before workers have ever acquired similar levels of rights and protections as they have in high-income countries.

Politics and political economy

The Age of Labor conferred not only unprecedented economic returns upon workers but also, in parallel, unprecedented political power. However, this power has been eroded more recently (see e.g. Boix, 2019). In democracies, the political question is, if workers are the majority, how is it possible that there are so many rules that disadvantage them? In simple models of democracy, it should be the median voter (or more broadly, the “majority”) that dominates. The majority of voters want a more egalitarian society. But in recent decades in many countries the political and economic rules have evolved in the opposite direction.

Hacker and Pierson (2020), attempt to explain this “democratic paradox.” In their view, Western Europe and the United States have taken two different courses. Most of Western Europe has followed the path of social democracy where there is a broad social consensus about egalitarianism, with strong institutions of social protection; political debates entail fine-tuning the system. The United States took another approach, in which a coalition between the rich and extremist groups (including religious groups and the gun lobby) holds a significant amount of power. This unstable coalition of rent seekers with extremist groups has gone in the opposite direction of egalitarianism, but to do that, it had to undermine democratic institutions, e.g. through disenfranchisement (e.g. voter suppression), disempowerment (e.g. gerrymandering), and “putting democracy in chains” (MacLean, 2017), thereby restraining what the majority can do. The European path offers a model for developing countries that is more likely to lead to shared prosperity (see also Stiglitz, 2019).

4.3 Manufacturing-Based Export-Led Growth

In developing countries, there has been a single model of development that has proved enormously successful over the past fifty years: manufacturing-based export-led growth (see Stiglitz, 2018a). It enabled many East Asian countries to close the gap between themselves and high-income countries, increasing per capita incomes in these countries manifold.

One big change inherent in this development strategy was moving from discussions of static comparative advantage to more dynamic comparative advantage. This was central to the East Asia “Miracle.” Half a century ago, South Korea was advised that its comparative advantage was growing rice. It rejected that advice – South Korea did not want to guide its policies based on its static comparative advantage at that time, growing rice, but on its dynamic comparative advantage. It observed that no one ever became rich by growing rice. Instead, it pursued an industrial policy that led it towards industrialization. That model served most of East Asia remarkably well, in a way few had anticipated. In 1969, Gunnar Myrdal wrote

Asian Drama, predicting that Asia will never develop – a reminder that economists should be careful in the forecasts they make.

The path to development in East Asia has been via exports of cheap labor-intensive manufactured goods. This development strategy combined learning, the provision of employment opportunities, foreign exchange, tax revenue—everything that was needed for a quick developmental transition. Their development trajectory began with taking advantage of their static comparative advantage in cheap labor, and especially cheap unskilled labor. But the countries of East Asia didn't let matters rest there. Over time, countries worked their way up the “value” chain, producing higher value added and more complex products and developing their dynamic comparative advantage.

Earlier advances in technology have already reduced the importance of cheap labor; but now advances in AI may erode it further still. Going forward, growth led solely by exports of labor-intensive manufacturing goods will no longer be available as a strategy of development. Indeed, the share of manufacturing employment is decreasing globally. Moreover, the jobs that can be outsourced may be more easily automated. There may be reshoring of production that had previously been outsourced, using highly automated production processes, and the process is accelerated by the Covid pandemic. The forces that facilitated the development in East Asia may thus be going in reverse, making it difficult for other developing countries to follow the strategy.

One of the critical reasons for the success of the export-led growth model based on manufacturing goods was that it enabled developing countries to catch up in multiple domains. Developing countries are poorer than developed countries not only because there is a gap in material resources but also because of a gap in knowledge (World Bank, 1998). The World Bank thought of itself as a knowledge bank, not only helping countries to catch up in resources but also to catch up in knowledge.

AI may have characteristics that will actually increase the gap in knowledge and make it more difficult to catch up. Cutting-edge AI technology is highly specialized. A disproportionate share of the people working in AI are in private companies, and a significant share of the knowledge is not in the public domain and therefore not easily accessible to developing and emerging economies. (This contrasts with many past technologies, when publicly financed knowledge production was more central, so access to knowledge was more easily available to developing and emerging economies.) Moreover, an important resource input to AI is data, and access to data is very concentrated and not globally public. The implication is that the nature of AI technology and how these advances are generated will make it more difficult to catch up than in the past.

4.4 What is Different This Time

Not only may the AI revolution make it more difficult for developing countries to catch up, it may well be that the AI revolution will be more difficult to manage for economic policymakers than earlier technological transitions. The structural transformation from an agrarian rural economy to an industrial urban economy eventually led to a more egalitarian society. As we have noted, the reasons included that innovation associated with that transition overall was unskilled-biased, i.e. it increased the relative productivity of unskilled labor. Moreover, industrial production provided a strong force towards mass education. Furthermore, industrial production typically involved large establishments that could be unionized relatively easily, and the unions advocated for wage compression. All these forces led to greater equality.

In the current transition, our destination may involve significantly more disruption and inequality, unless countervailing policy interventions are made. AI may be labor-saving and resource-saving, and it is likely more biased towards ever-higher skills so that general education becomes less important. This may reduce support for equality-enhancing public education, which has been one of the strong forces for more equalitarian outcomes in the past. Moreover, the service sector which is becoming an increasingly important part of the economy is marked by smaller establishments and, in addition, worker tenure has declined, making it harder to unionize the workforce. Digital technologies are likely to create more barriers to entry and give rise to more monopoly power and winner-takes-all dynamics, with rents going to a small number of extremely wealthy individuals and enterprises, disproportionately located in high-income countries.

Although for many developing countries, average income per capita may increase, large fractions of society may be left behind. Moreover, some developing countries may experience declines in income per capita as innovation erodes their comparative advantage. As a result, declines in incomes per capita would be compounded by technological dynamics that generate further redistributions away from workers towards those who benefit from the new technologies within a given country.

Although greater inequality would increase the need for social protection, it may result in a less egalitarian politico-economic equilibrium, as the new concentrations of economic and political power may reduce support for the critical role of government in mitigating the adverse distributional consequences of technological change.

5 Economic Policy Responses

We have seen how economic policy played a critical role in shaping economic outcomes in previous eras of innovation; the same will be true in the case of AI. In this section, we discuss what policy levers developing countries can employ to address the effects of technological disruption. Some of these are similar to what worked in earlier periods of technological change; some are attuned to the special problems posed by AI and labor-replacing innovation. In section 6, we will discuss changes in global policies, norms and rules that would assist developing countries in their response to technological change. In this short paper, we can only touch on a few of the more salient policies.

5.1 Taxation and Redistribution

Among the critical policies to combat rising inequality are those of taxation and redistribution, with a particularly important role for progressive taxation. The bias of technological progress against unskilled labor and the winner-takes-all effects described earlier, combined with the increases in monopoly power, have contributed to rising levels of income inequality, and even greater levels of wealth inequality *within countries* around the world. This has increased the importance of progressive taxation, yet in recent years a number of countries have actually made their tax systems more and more regressive. For example, many countries tax the returns to capital and rents (such as land rents, monopoly rents, and other forms of exploitation rents) at lower rates than workers. In the US, the rich pay a lower fraction of their income in taxes than the majority of the population (Saez and Zucman, 2019).

Raising taxes is a particular challenge for developing countries, in which the informal sector is typically much larger than in high-income economies. Yet new digital tools and new data may actually give governments new policy tools to increase tax compliance. For example, when an activity becomes

intermediated via centralized digital platforms, it becomes easier for government to access business transactions and levy taxes on them. For example, governments have long found it difficult to monitor and tax the earnings of taxi drivers. But if driving is intermediated via digital platforms, all their earnings – including most tips – are recorded.⁷

One of the dilemmas when it comes to taxation and redistribution is that labor-saving technological progress reduces tax revenue from labor – traditionally the most highly taxed factor in the economy – precisely at the time when the need for redistribution rises (see e.g. Korinek, 2020). This necessitates that taxation increasingly shifts towards other factors and rents. From the perspective of efficiency, the taxation of rents is particularly desirable (George, 1879). Imposing taxes on fixed factors like land acts like a lump sum tax. And taxing rents may actually discourage rent-seeking, enhancing efficiency.

What we have observed earlier is that technological progress always creates winners and losers, and the gains of the winners are quasi-rents that governments may be able to tax without introducing distortions. In particular, some of the monopoly rents of digital giants can be taxed without introducing major distortions into the economy.

In designing tax systems, an important concern is about incidence: the possibility that general equilibrium effects imply that taxes are ultimately borne by other factors and agents than those on whom they are levied, undermining the desired redistributive objectives. For example, a common result in simple models is that capital taxation discourages capital accumulation by capitalists. However, the adverse effects may be more than offset by public investments in human and physical capital (see e.g. Stiglitz, 2018b).

Taxes on “bads” rather than goods, i.e. Pigouvian taxes on activities and goods that create negative externalities (for example polluting or carbon-emitting goods), would contribute to the Green Transition in a dual way: they would not only provide tax revenue for public investments but also correct market prices to reflect the negative externalities.

Social Protection

If individuals could obtain insurance against the adverse effects of disruptive innovations, then it would be more likely that these innovations would be Pareto improvements (Korinek and Stiglitz, 2019). One of the functions of social insurance is to socialize these risks that otherwise would have been borne by individuals. But in developing countries, systems of social protection are typically less developed, making it even more likely that there be significant groups that are worse off. While a simple prescription might be to provide better systems of social protection, resource and institutional constraints make that difficult. Countries worry that money so spent is money taken away from developmental investments.

Universal Basic Income

Many in the tech community have responded to concerns about the impact of technological progress on employment by advocating a universal basic income. We offer three thoughts on such proposals.

⁷ Some are justifiably concerned that digital platforms are in fact very efficient at exploiting workers, e.g. that the hourly earnings of drivers may not exceed minimum wages. But platforms can also provide information on whether workers are exploited and with proper regulation, make it easier to address such exploitation than it used to be before the digital age.

First, we would greatly welcome if proponents of a Universal Basic Income were true to the meaning of the word “universal,” proposing that the basic income covers all the citizens of the world equally, and not just people who were lucky to be born in specific locations that have the fiscal capacity to afford such programs (e.g. in Alaska where oil revenue is collected in the Alaska Permanent Fund and distributed to the residents of the state). Given the general difficulty of providing transfers across national borders, this would be a major step forward.

Second, UBI programs might make sense in a future in which labor truly becomes redundant (Korinek and Juelfs, 2020). But for now, it may be more desirable to focus on creating jobs for everyone who is able and willing to work. As we suggested earlier, there is a lot of work that needs to be done—creating the Green Transition, providing services to the young, the sick, and the elderly, investing in infrastructure—and a lot of people who want to work. Markets often fail to match the need for work and the willingness of people to work. In such circumstances, governments should step in.

Third, one approach to ensuring a modicum of income for all in the long run, with co-benefits of perhaps increasing social cohesion and solidarity, is shared capital ownership: as part of government assistance programs (such as those enacted in the wake of COVID-19 in 2020), firms receiving government help could be asked to contribute shares to a sovereign wealth fund—owned by everyone within the nation. Similarly, firms that build on or employ innovations that are based in part on government-funded research could be required to do the same.⁸

Redistribution vs Pre-Distribution

If the scope for redistribution is limited, especially in developing countries in which the capacity to tax is low, we need to change our emphasis from redistribution to pre-distribution, i.e. we need to focus on how to affect the distribution of market income rather than taking market income as given and imposing taxes and transfers (Stiglitz *et al*, 2015, 2019). As we will explore in the following sections, there are many policies to which developing countries can resort that affect the distribution of market income itself.

5.2 Expenditure and Infrastructure Policy

Expenditure policy can be as important in offsetting the adverse effects of AI as taxation and direct redistribution, and it carries several benefits over transfers that are particularly relevant in developing countries: government expenditures may be easier to target based on need, and they may be more robust to corruption. For instance, expenditures on human well-being, such as on education and health, are naturally targeted to those who receive education and healthcare, rather than being spent on those who already are educated or on those who are healthy. Environmental policies help those who bear the brunt of environmental degradation, including climate change, which are borne disproportionately by the poor.

Expenditure policies that increase the demand for unskilled labor may serve double duty: they raise demand for unskilled labor and increase the equality of market income (what is often now called pre-distribution), and they can be targeted so that the benefits of the expenditure go disproportionately to the less well off, increasing the equality of “well-being.” One important example are infrastructure investments, which are a labor-intensive expenditure that can be designed to be pro-egalitarian.

⁸ Conceptually, government ownership of capital is equivalent to taxes on capital with exemptions on new investment that avoid the negative incentive effects of capital taxation.

Of particular importance are investments in the digital infrastructure of a country, which reduce the “digital divide” and allow citizens to access the vast services provided by the Internet. Recent advances in network technology allow developing countries to leapfrog older technologies in which high-income countries have invested fortunes, for example by using wireless 5G technologies instead of laying vast networks of cables.

Other infrastructure investments include public transportation systems that connect especially lower income workers with jobs and enhance the opportunities available to them. Another example of labor-demand increasing public expenditures is creating service sector jobs, for example in healthcare, caring for the aged, and some aspects of education, which can be designed to serve double duty – disproportionately benefiting the poor and needy as they increase wages by increasing the demand for labor.

5.3 Education

In some quarters, education has come to be seen as the panacea to inequality, since more educated workers receive higher incomes than less educated.

In high-income countries, the potential for education to solve the problem of inequality, let alone the particular problem of labor-saving technological progress that we discuss in this paper, may have been exaggerated. For instance, many countries have seen stagnation of incomes of even those with a college education in recent years. And as we noted, AI has gone some way in replacing even some of the critical tasks performed by highly skilled workers, for example in radiology.

However, a certain level of education is critical for the citizens of developing and high-income countries alike to participate in the modern digital economy and avoid an education-based digital divide whereby some simply do not know how to access and benefit from the resources and opportunities offered by the Internet and related digital technologies.

5.4 New Development Strategies

Developing countries will need a new multi-pronged development strategy to replace the manufacturing-led export-based growth model, which used to be the most successful development model ever. One of the most important aspects of countries’ development strategies have traditionally been industrial policies—interventions that shape the direction in which the economy is moving, with particular emphasis on the secondary sector. However, in an age of increasing automation in manufacturing, development strategies have to broaden their focus beyond manufacturing and beyond the secondary sector to other sectors of the economy, including agriculture and services.⁹

Greenwald and Stiglitz (2014b) point out that every country has, in effect, a sectoral development policy—shaped by infrastructure and education investments and tax and regulatory policy. It is only that some countries don’t know, or won’t admit, that they have such policies. The danger then is that such policies

⁹ Curiously, such policies have continued to be referred to as “industrial policies” even when they move the economy away from the industrial sector. We use the more generic title sectoral policies, but they are broader: they can also be used to change technology within a sector (e.g. towards green or more labor-intensive technologies.)

can be more easily captured by special interests.¹⁰ In developing countries development policies are much more at the center of economic policy. They need to be designed to manage, mitigate the effects of, and adapt to the disruptions of innovations at home and abroad, to ensure that the net societal benefits, broadly defined, are maximized.

A lot of innovation in developing countries focuses on technology adoption rather than developing entirely novel technologies. Whereas high-income countries focus on “steering innovation,” developing countries need to pay attention to “steering the adoption of technologies.” Their development strategy could intentionally focus on steering the adoption of labor-using technologies that have already been developed in high-income countries, adapting them to their own circumstances and needs, redesigning them, and building on them. Decisions on what type of inward FDI to encourage could also be informed by these objectives.

In designing the new development strategies, developing countries will need to think carefully about the rationale for public interventions. Of particular importance is that the direction of technological progress and of technology adoption is endogenous, and there is no presumption that market decisions in this area are socially desirable. Decisions made at one date have effects in later periods, with firms making the decisions appropriating only a fraction of the benefits and bearing only part of the costs of their decisions. For example, this is clearly manifest when there are knowledge spillovers to other firms and when technology evolves over time, e.g. through learning by doing. Firms acting on their own will not fully take into account the dynamic implications of their decisions today on others.

Part of the problem are other market failures beyond the ability to appropriate the returns from current choices—for instance, imperfections of risk and capital markets. The capital market imperfections that impede the reallocation of labor in high-income countries in response to innovation—and that can result in innovations being welfare decreasing—are even more important in developing countries, making it imperative to combine industrial policies with active labor market policies (see e.g. delli Gati et al 2012a, 2012b).

Relatedly, part of the problem is that market prices don’t adequately reflect social shadow values. A well-known example is that, in the absence of appropriate regulation, the price of carbon in the market is zero, but this does not reflect the social cost of carbon.

Similarly, market prices may reflect the scarcity value of resources and guide economic activity towards what enhances efficiency, but they do not reflect the social value of an equitable distribution of resources and do not guide economic activity in that direction. Given the constraints on redistribution, this leaves an important role for the government to steer innovation and foster economic development in a socially desirable direction (Korinek and Stiglitz, 2020b). For example, much could be gained from encouraging innovators to shift their focus from labor-saving towards more labor-using technologies.

Fortunately, while the new technologies necessitate a change away from the old and highly successful development strategies of the past half century, they also open up new opportunities. In agriculture, AI offers the potential for large productivity increases based on algorithms that help farmers fine-tune and optimize a range of decisions that increase their yield. Such algorithms depend on crops, soil and

¹⁰ For example, US bankruptcy provisions favoring derivatives can be thought of a sectoral policy encouraging the growth of derivatives; but until the 2008 financial crisis, few outside of that sector were even aware of the favorable treatment that derivatives have received.

weather conditions and need to be customized to local conditions. Just as agricultural extension services extending general knowledge about agriculture to local farmers played an important role in the development of the US in the second half of the nineteenth century and early part of the twentieth, there is an important role for government agricultural extension services today in developing countries. Digital platforms can also greatly enhance the ability of small farmers to trade their products at fair market prices, reducing the market power of middle men that frequently absorb a significant fraction of the surplus generated in agriculture.

Developing the service sector is crucial for economic development as the role of the primary and secondary sectors is declining. Many developing countries may carve out new areas of comparative advantage in services that will, however, depend on good Internet connections and a certain degree of education of the workforce. For example, call centers and similar business and consumer services rely on requisite language skills. There is also a growing market for simple human services that can be broken down into small components and fed into AI systems (e.g. labeling images). However, what adds to the difficulty is that services that can be outsourced may also be more easily automated. Other services such as tourism have proven a more automation-resistant (although not pandemic-resistant) source of export revenue for countries that have managed to fashion themselves into desirable tourist destinations. Exporting services offers many of the potential growth benefits of the manufacturing-based export-led growth model.

Services that are aimed at a domestic audience, for example healthcare, caring for the aged, as well as education, may not deliver much export revenue but are very important for economic development and welfare. There is much scope for employing AI to improve the delivery and efficiency of these services, and it requires government policy to do so since private service providers are frequently small in size and cannot afford the necessary investments. And even in these areas, there may be significant opportunities for cross border trade, for example via medical tourism and via retirees from advanced countries locating in warmer climates, if adequate health care is available.

Steering Innovation in AI in High-income Countries

The overall direction of innovation in AI will be set to a large extent by high-income countries plus China. This implies that technological progress in those countries – and how labor-saving it is – also matters for developing countries that will be exposed to the new technologies.

Many governmental policies have indirect effects on incentives for innovation. For example, at least in the short run, the cost of capital is influenced by monetary policy, with the goal of stabilizing aggregate demand. In recent years, monetary authorities in many countries have set interest rates such that real returns on safe assets have been very low or even negative, likely below the social shadow price of capital. Stiglitz (2014) shows that this encourages excessive automation in high-income countries. Acemoglu et al (2020) observe that tax policies that favor capital over labor also distort the direction of progress towards saving labor.

And there are immediate implications for developing countries: Once the cost of developing a labor-saving innovation has been incurred in high-income countries, it can frequently be rolled out globally at comparatively low cost, potentially imposing significant welfare costs on workers in developing countries (see e.g. Diao et al, 2021).

Pritchett (2019) observes that migration policies in high-income countries restrict labor supply and lead to comparatively high wages that do not reflect the abundance of labor, and in particular of unskilled labor, at the global level. The high wages then provide innovators in high-income countries with excessive incentives to invest in the automation of tasks that are performed by unskilled labor compared to what is desirable from the perspective of developing countries.

Regional policies

Economists are also becoming increasingly aware of the importance of regional heterogeneity. Unlike in stylized models in which only national borders exist, labor does not move seamlessly across regions within countries. Even in high-income countries, large disparities between regions or between rural and urban areas persist, as illustrated e.g. by the case of northern and southern Italy or by the rural/urban differential in the United States and many other countries. Such disparities call for location-based policies in fostering development.

6 Economic Development and Global Governance

In a globally integrated economy—from which developing countries and emerging markets have benefited enormously in many ways—global rules matter. The global rules have always been set to favor high-income countries; they are, to a large extent, set by the large powerful countries, and frequently by powerful special interests within them, whereas developing countries don't have a seat at the table, or are at least underrepresented.

The global rules have large effects on the ability of these countries to levy taxes in the digital era, on high-income countries' ability to extract rents from the developing countries (say through market power and intellectual property rights), and more broadly on the global terms of trade and distribution of income.¹¹ While developing countries may realize these inequities—and the inefficiencies—of our global economic system, it often seems that there is little they can do.

AI has provided a new arena in which rules need to be set, at the same time that it may exacerbate the imbalances in economic power, as our earlier discussion emphasized.

However, there are reasons for cautious hope when it comes to the rules governing information and AI. First, the rules in this area are still in the process of being set so there is hope that international institutions and civil society may have a positive impact on the shape of these rules. Still, the fact that recent trade agreements between the US and other countries have contained provisions reflecting the interests of big-tech companies—with limited open debate and limiting the scope for these trading partners to design regimes that reflect a broader public interest—is of concern.

Secondly, it should be in the self-interest of high-income countries to avoid the possibility of a strong backlash to globalization in developing countries. The possibility of such a backlash is considerable: The United States and a number of other high-income countries, which have been big beneficiaries of globalization, have experienced such a backlash – in part because they have not ensured that the losers of globalization were compensated. In the past, there was at least some sense that globalization created mutual gains for high-income and developing countries. The backlash in developing countries would be even greater if they come to see globalization as a mechanism of rent extraction from their economies

¹¹ For a discussion of how this plays out in trade rules, see, e.g. Charlton and Stiglitz (2005).

(even if the truth may be that technological change is making them lose some of the earlier gains from globalization).

Moreover, international institutions, some of which are less and less dominated by high-income countries, may play a role in ensuring that the rules are set in a way that more adequately reflects the interests and concerns of all countries, including developing countries.

As the rules for new technologies are being written, there are several areas of particular concern in which reforms in global governance would help developing countries better adapt to advances in AI.

6.1 A Global Tax Regime for the Digital Age

The inadequacies in the global tax regime make it difficult for developing countries to capture much of the rents that the global digital giants earn within their borders, even as their business activities take away from the activities of domestic firms and thereby reduce the domestic tax base. Indeed, even high-income countries have had difficulty with adequately taxing global tech giants. Some of the issues are now being discussed at the OECD in an attempt to establish a global tax regime.

The current global tax regime allows multinational firms to avoid much taxation—often paying taxes at rates markedly lower than local small businesses. It also impairs the ability of developing countries and emerging markets to tax the economic activity which occurs within their territories. This system is both inefficient and inequitable.

The controversy over digital taxation has exposed the deeper problems of multinational corporate taxation based on transfer prices, which are easily manipulated. The issue could be addressed by moving to a formulary apportionment system, whereby the worldwide profits of a corporation are apportioned to different countries according to a formula (see e.g. Clausing and Avi-Yonah, 2007). The exact formula could have large distributive effects across countries. For instance, a simple formula based just on sales, while less manipulatable than other formulae, may disadvantage developing countries. A particular controversy associated with the digital economy is the value assigned to the data that is collected in the process of economic transactions.

The broader debate over international taxation has also led to renewed attention on closing down fiscal paradises, on international initiatives for transparency in capital ownership, which would help developing countries to increase their tax base, and on creating a global minimum multinational corporate tax rate, to prevent a race to the bottom.

6.2 Global Competition Policy

The tendency of digital technologies to give rise to natural monopolies makes competition policy especially important. One challenge is that the countries in which tech giants are based have incentives to protect their own tech firms since they share in the rents that these firms earn globally. For example, when the European Union investigated Google for anti-competitive practices or when Germany investigated the privacy practices of Facebook, the US treated it as a political question rather than a matter of economic policy and responded by accusing Europe of being anti-American. Such accusations were misplaced: although it was true that the policy remedies suggested by the Europeans may have reduced the rents the companies could earn in Europe, they simply wanted to ascertain that these firms' practices did not violate the norms on competition and privacy established in Europe. The tendency for

matters of competition policy to turn into arguments over rents may get worse, given the global concentration of market power in AI in two countries, China and the United States.

Individual developing countries and emerging market economies stand little chance in reining in the behaviors of powerful global corporations on their own – in many instances, the corporations have a higher market capitalization than the GDP of the countries in question. This makes it important for developing countries to coordinate and develop competition policy together, for example via a common competition authority for developing and emerging economies that can exert sufficient power over large global corporations, just as the countries of Europe would not be able to police the competitive behavior of American corporations on their own but are able to do so through the European Union.

Given the breadth and reach of the new digital giants, it may be desirable for consumers to enact stronger rules preventing conflicts of interest for companies that e.g. simultaneously own a marketplace and participate in it, and stronger rules preventing pre-emptive mergers, i.e. mergers and acquisitions designed to stifle the threat of a competitive marketplace in the future. It may also be desirable for consumers to engage in ex-post remediation by breaking up mergers when they prove to be anti-competitive.¹² As the experiences cited above have shown, the countries in which digital giants are based may not have the correct incentives to police these companies' competitive practices, given the large global rents that are at stake.

6.3 Intellectual Property Rights

Our system of intellectual property (IP) rights is designed to give monopoly rents to innovators to compensate and reward them for their innovative activities. There has been much concern in recent years that the prevailing IP system gives excessive protection to innovators, with particularly adverse effects on developing countries. As the World Commission on the Social Dimensions of Globalization (2004) emphasized, there is a need to rebalance the international intellectual property regime to ensure an equitable distribution of the gains from technological progress. We demonstrated in Korinek and Stiglitz (2019) that reducing the length of patent protection can ensure that the gains from AI-based innovations are shared among society and lead to a welfare improvement.

The most efficient way of distributing technological advances is to keep them in the public domain, financed via governments, international organizations, donors or charities. This avoids restrictions in access to new technologies and the creation of monopolies that concentrate rents and power. There is much scope for publicly financed research and development to benefit developing countries, for example in the areas of agriculture where new technologies increase the productivity of crops, or in healthcare where developing countries face unique challenges that do not attract sufficient research by private corporations in high-income countries.

When research and development is financed privately, there is a strong case for granting different patent protection in developing countries than in high-income economies. The length of patent protection trades off how much surplus to allocate to innovators to compensate them for their efforts versus how much to let the broader public benefit from an innovation. Most patents are developed in high-income countries and are financed by the surplus that innovators extract from the patent protection there; innovators would not incur significant losses if developing countries could use their technology for free before their patents

¹² There is by now a large literature describing the new competition policies that may be required. See Stiglitz (2019) and Wu (2018).

expire in high-income countries. Indeed, in the drug sector, this realization has led drug companies to offer life-saving drugs to some of the poorest countries at steeply discounted prices. Compulsory licenses (part of TRIPS and other international agreements) give the right to developing countries to access such life-saving IP at appropriate royalties, but developing countries often do not have the capacity to exercise those rights and are intimidated from doing so by threats from developed countries. Worse, trade agreements have impeded access to generic medicines, forcing developing countries to pay high prices for drugs.

Before the advent of AI, it was clear that there was a need for a developmentally oriented intellectual property regime—in some ways markedly different from that currently prevailing (Cimoli *et al* 2014).

But AI has made the challenge of access to knowledge even greater. Part of the nature of AI is it may not even need much protection by the patent system. Algorithms can be kept proprietary, and they are always evolving. Requiring disclosure of certain key algorithms would be imperative to ascertain whether algorithms are discriminatory, for example by engaging in price discrimination.¹³

6.4 Data and Information Policy

Data is a critical input underlying the new AI economy. That is why information policy – the rules governing the control over and use of data – has moved to the top of the policy agenda. There appears to be an on-going attempt of global AI firms to set the data regulatory agenda while nobody is looking. This is already happening in recent trade agreements. For instance, while the new trade agreement between Canada, Mexico and the US had stronger provisions protecting labor and access to healthcare as well as better investor-state dispute settlement provisions, rules on the digital economy moved in the opposite direction, providing better protection for the tech giants. Being part of an international agreement, it may be difficult to change the data regulation regime in the future.

This is particularly important for developing countries: the rules are now being set by the tech giants and for the tech giants, with little concern for the views of citizens in the high-income countries, let alone those in the rest of the world.

Moreover, the monopolization of data by global AI firms also makes it more difficult for developing countries to catch up and develop their own AI-based companies. Global firms can use their access to vast troves of data from across the world to refine their products and offerings to consumers ever further. This makes it more and more difficult for newcomers in developing countries to catch up.

Europe has actively worked on rules to ensure that the benefits of new digital technologies are shared and the harms are minimized. For instance, the EU has put forward proposals to require data sharing, with the goal of preventing accretion of monopoly power by monopolizing data. But giving control rights over data to individuals will not suffice; there is evidence that, without regulation, individuals turn their data over to the digital giants and internet providers, receiving but a pittance: asymmetries in information and power are just too great to ensure an equitable outcome.

New transparency regulations, e.g. on the algorithms and targeting of advertising, would be desirable to ensure that harmful practices cannot go undetected. This would also allow us to address discriminatory

¹³ It is sometimes argued that such disclosure is not possible because algorithms are always evolving. While true, they could still be disclosed as of a particular moment in time.

impacts of pricing and advertising. It may be desirable to enact stronger rules protecting privacy and stemming the rapid spread of misinformation and messages that promote violence and hate as well as other harmful messaging, even when conducted as part of a political campaign. There is no basic right to virality. In the US, it may be desirable to reconsider the Section 230 provision, which reduces the accountability of internet companies compared to other publishers.

As in the case of competition policy, the countries in which tech giants are based may not face the correct incentives to police the worldwide behavior of their companies since they share in the rents that these companies earn around the world. Developing countries would need to cooperate and band together to have sufficient clout to impose regulation on global giants that reflects their developmental interests.

7 Conclusion

Advances in AI and related technologies may, like the Industrial Revolution, represent a critical turning point in history. Increasing automation in manufacturing may lead to the demise of the manufacturing-export-led developmental model, which has had such profound positive effects on many emerging market economies. The worst-case scenario is the unravelling of much of the gains in development and poverty reduction that we have seen over the last half century.

The concern is, while earlier technological advances were associated with shared increases in prosperity, increasing equality between countries, as suggested by the convergence hypothesis in standard neoclassical theory, the new advances may arrest the convergence in standards of living between rich countries and developing countries. They may instead result in increasing inequality both within and across countries, unless we design policies that counterbalance them.

The new era will be governed by different rules and will require a different kind of economic analysis. Just like the production functions that Ricardo used to analyze agrarian and rural economies were very different from those in the models of manufacturing that dominated the mid-20th century, we need to adjust and update our economic frameworks to think about the models that will describe the next 50 years. For instance, the competitive equilibrium model may be even less relevant to the 21st century AI economy than it was to the 20th century manufacturing economy.

It is obvious that the future is uncertain—and that there are large downside risks that should not be ignored. Economic analysis, based on models appropriate to this new era, has the potential to help in the development of policies—both at the global and national level—that can mitigate these adverse effects, to ensure that this new era of innovation will lead to increased standards of living for all, including the billions living in developing countries.

References

- Acemoglu, Daron (1998), "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality," *Quarterly Journal of Economics*, 113(4), pp. 1055-1089.
- Acemoglu, Daron (2002), "Directed Technical Change," *Review of Economic Studies*, 69(4), pp. 781-809.
- Acemoglu, Daron and David Autor (2011), "Skills, Tasks and Technologies: Implications for Employment and Earnings," *Handbook of Labor Economics* 4b, pp. 1043-1171.
- Acemoglu, Daron, Andrea Manera, and Pascual Restrepo (2020), "Does the US tax code favor automation?" *Brookings Papers on Economic Activity* 2020(1).
- Aghion, Philippe, Ufuk Akcigit, Antonin Bergeaud, Richard Blundell, and David Hemous (2019), "Innovation and Top Income Inequality," *Review of Economic Studies* 86(1), pp. 1-45.
- Alonso, Cristian, Andrew Berg, Siddharth Kothari, Chris Papageorgiou and Sidra Rehman (2020), "Will the AI Revolution Cause a Great Divergence?" *IMF Working Paper* 20/184.
- Arntz, Melanie, Terry Gregory and Uleich Zierahn (2016), "The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis", *OECD Social, Employment and Migration Working Papers* 189.
- Atkinson, Anthony and Joseph Stiglitz (1969), "A New View of Technological Change," *Economic Journal* 79 (315), pp. 573-578
- Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen (2020), "The Fall of the Labor Share and the Rise of Superstar Firms," *Quarterly Journal of Economics* 135(2), pp. 645-709.
- Autor, David, Frank Levy, and Richard J. Murnane, (2003), "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics* 118(4), pp. 1279-1333.
- Barro, Robert J., and Xavier Sala-i-Martin (1992), "Convergence," *Journal of Political Economy* 100(2), pp. 223-251.
- Berg, Andrew, Edward F. Buffie and Luis-Felipe Zanna (2018), "Should we fear the robot revolution? (The correct answer is yes)," *Journal of Monetary Economics* 97, pp. 117-148.
- Boix, Carles (2019), *Democratic Capitalism at the Crossroads: Technological Change and the Future of Politics*, Princeton University Press.
- Brynjolfsson, Erik, Avinash Collis, Erwin Diewert, Felix Eggers, and Kevin Fox (2020), "Measuring the Impact of Free Goods on Real Household Consumption," *AEA Papers and Proceedings* 110, pp. 25-30.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson (2019), "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics," *The Economics of Artificial Intelligence: An Agenda*, Agrawal, Gans, and Goldfarb.
- Caselli, Francesco, and Alan Manning (2019), "Robot Arithmetic: New Technology and Wages," *American Economic Review: Insights* 2019 1(1), pp. 1-12.
- Charlton, Andrew, and Joseph E. Stiglitz (2005), *Fair Trade for All*, Oxford University Press.

Cimoli, Mario, Giovanni Dosi, Keith E. Maskus, Ruth L. Okediji, Jerome H. Reichman, and Joseph Stiglitz (eds.) (2014), *Intellectual Property Rights: Legal and Economic Challenges for Development*, Oxford University Press

Clausing, Kimberley and Reuben Avi-Yonah (2007), "Reforming Corporate Taxation in a Global Economy: A Proposal to Adopt Formulary Apportionment," *The Hamilton Project, Brookings Institution*.

Diao, Xinshen, Mia Ellis, Margaret McMillan, and Dani Rodrik (2021), "Africa's Manufacturing Puzzle: Evidence from Tanzanian and Ethiopian Firms," working paper, International Food Policy Research Institute.

Delli Gatti, Domenico, Mauro Gallegati, Bruce Greenwald, Alberto Russo, and Joseph E. Stiglitz (2012a), "Sectoral Imbalances and Long Run Crises," *The Global Macro Economy and Finance*, Franklin Allen, Masahiko Aoki, Jean-Paul Fitoussi, Nobuhiro Kiyotaki, Roger Gordon, and Joseph Stiglitz, eds., *IEA Conference Volume No. 150-III*, Palgrave, pp. 61-97.

Delli Gatti, Domenico, Mauro Gallegati, Bruce Greenwald, Alberto Russo, and Joseph E. Stiglitz (2012b), "Mobility Constraints, Productivity Trends, and Extended Crises," *Journal of Economic Behavior & Organization* 83(3), pp. 375-393.

Faber, Marius (2021), Robots and Reshoring: Evidence from Mexican Labor Markets, forthcoming, *Journal of International Economics*.

Frey, Carl Benedikt and Michael A. Osborne (2017), "The future of employment: How susceptible are jobs to computerisation?" *Technological Forecasting and Social Change* 114, pp. 254-280.

George, Henry (1879), *Progress and poverty: An inquiry into the cause of industrial depressions, and of increase of want with increase of wealth, the remedy*, D Appleton and Company.

Good, Irving John (1965), "Speculations Concerning the First Ultraintelligent Machine," *Advances in Computers*, vol. 6.

Gordon, Robert (2016), *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War*, Princeton University Press.

Greenwald, Bruce and Joseph Stiglitz (2014a), *Creating a Learning Society: A New Approach to Growth, Development, and Social Progress*, Columbia University Press.

Greenwald, Bruce and Joseph E. Stiglitz (2014b), "Industrial Policies, the Creation of a Learning Society, and Economic Development," *The Industrial Policy Revolution I: The Role of Government Beyond Ideology*, Joseph E. Stiglitz and Justin Yifu Lin (eds.), Palgrave Macmillan, pp. 43-71.

Hacker and Pierson (2020), *Let them Eat Tweets: How the Right Rules in an Age of Extreme Inequality*, Norton.

Harris, Karen, Austin Kimson, and Andrew Schwedel (2018), "Labor 2030: The Collision of Demographics, Automation and Inequality," Bain and Company Reports.

Karabarbounis, Loukas and Brent Neiman (2013), "The Global Decline of the Labor Share," *Quarterly Journal of Economics* 129(1), pp. 61-103.

Kennedy, Charles (1964), "Induced Bias in Innovation and the Theory of Distribution," *Economic Journal* LXXIV, pp. 541-47.

Koopmans, Tjalling C. (1960), "Stationary ordinal utility and impatience," *Econometrica* 28, pp. 287-309.

Korinek, Anton (2020), "Taxation and the Vanishing Labor Market in the Age of AI," *Ohio State Technology Law Journal* 16(1), pp. 244-257.

Korinek, Anton and Megan Juelfs (2020), "Preparing for the (Non-Existent?) Future of Work," technical report, Darden School of Business.

Korinek, Anton and Ding Xuan Ng (2019), "Digitization and the Macro-Economics of Superstars," working paper, University of Virginia.

Korinek, Anton and Joseph E. Stiglitz (2019), "Artificial Intelligence and Its Implications for Income Distribution and Unemployment," in Agrawal et al.: *The Economics of Artificial Intelligence*, NBER and University of Chicago Press.

Korinek, Anton and Joseph E. Stiglitz (2020a), "Will COVID-19 Drive Advances in Automation and AI that Exacerbate Economic Inequality?" accepted, *BMJ*.

Korinek, Anton and Joseph E. Stiglitz (2020b), "Steering Technological Progress," working paper.

Korinek, Anton and Joseph E. Stiglitz (2021), "Artificial Intelligence, Worker-Replacing Technological Progress, and Income Distribution," with Joseph Stiglitz, working paper.

Kurzweil, Ray (2005) *The Singularity is Near*, Penguin Group, pp. 135–136.

Maddison, Angus (2003), *Contours of the World Economy 1-2030 AD: Essays in Macro-Economic History*, Oxford University Press.

Malthus, Thomas R. (1798), *An Essay on the Principle of Population*, J. Johnson.

Mankiw, N. Gregory, David Romer and David N. Weil (1992), "A Contribution to the Empirics of Economic Growth," *Quarterly Journal of Economics* 107(2), pp. 407-437.

MacLean, Nancy (2017), *Democracy in Chains: The Deep History of the Radical Right's Stealth Plan for America*, Penguin Random House.

McKinsey Global Institute (2017), "Jobs lost, jobs gained: What the future of work will mean for jobs, skills, and wages," McKinsey Reports.

Myrdal, Gunnar (1968), *Asian Drama: An Inquiry into the Poverty of Nations*, Pantheon Books.

Nordhaus, William (2015), "Are We Approaching an Economic Singularity? Information Technology and the Future of Economic Growth," *NBER Working Paper* No. 21547.

Ostry, Jonathan, Prakash Loungani, and Andrew Berg (2019), *Confronting Inequality: How Societies Can Choose Inclusive Growth*, Columbia University Press.

Pritchett, Lant (2019), "The Future of Jobs is Facing One, Maybe Two, of the Biggest Price Distortions Ever," *Economic Research Forum Working Paper* 1370. <https://erf.org.eg/app/uploads/2019/12/1370.pdf>

Saez, Emmanuel and Gabriel Zucman (2019), *The Triumph of Injustice: How the Rich Dodge Taxes and How to Make Them Pay*, Norton.

Samuelson, Paul (1965), "A Theory of Induced Innovations along Kennedy-Weisäcker Lines," *Review of Economics and Statistics* XLVII:444-64.

Solow, Robert (1987), "We'd Better Watch Out." *New York Times Book Review*, July 12, pp. 36.

Stiglitz, Joseph E. (2014), "Unemployment and Innovation," *NBER Working Paper* 20670.

Stiglitz, Joseph E. (2015), "Leaders and followers: Perspectives on the Nordic model and the economics of innovation," *Journal of Public Economics, Elsevier*, 127(C), pp 3-16.

Stiglitz, Joseph E. (2018a), "From Manufacturing-Led Export Growth To A Twenty-First-Century Inclusive Growth Strategy: Explaining The Demise Of A Successful Growth Model And What To Do About It," *WIDER Working Paper 2018/176 Helsinki: UNU-WIDER*.

<https://www.wider.unu.edu/publication/manufacturing-led-export-growth-twenty-first-century-inclusive-growth-strategy>

Stiglitz, Joseph E. (2018b), "Pareto Efficient Taxation and Expenditures: Pre- and Re-distribution," *Journal of Public Economics Special Issue in Honor of Sir Tony Atkinson (1944-2017)*, 162, pp. 101-119.

Stiglitz, Joseph E. (2019), *People, Power, and Profits: Progressive Capitalism for an Age of Discontent*, Norton.

Stiglitz, Joseph E., Nell Abernathy, Adam Hersh, Susan Holmberg and Mike Konczal (2015), *Rewriting the Rules of the American Economy: An Agenda for Growth and Shared Prosperity*, A Roosevelt Institute Book, Norton.

Stiglitz, Joseph E., Carter Daugherty, and the Foundation for European Progressive Studies (2019), *Rewriting the Rules of the European Economy: An Agenda for Growth and Shared Prosperity*, Norton.

Straub, Ludwig (2020), "Consumption, Savings, and the Distribution of Permanent Income," working paper, Harvard.

Trajtenberg, Manuel (2019), "Artificial Intelligence as the Next GPT: A Political-Economy Perspective," pp. 175-186, in Agrawal, Ajay, Joshua Gans and Avi Goldfarb (2019), *The Economics of Artificial Intelligence: An Agenda*, National Bureau of Economic Research and University of Chicago Press.

Trammell, Philip and Anton Korinek (2021), "Economic Growth under Transformative AI," working paper, Oxford.

Ulam, Stanislaw (1958), "Tribute to John von Neumann," *Bulletin of the American Mathematical Society* 64, #3, part 2, pp. 5.

Varian, Hal (2020), "Automation versus procreation (aka bots versus tots)," VoxEU, <https://voxeu.org/article/automation-versus-procreation-aka-bots-versus-tots>

Von Weizsäcker, Carl (1966), "Tentative Notes on a Two-Sector Model with Induced Technical Progress," *Review of Economic Studies* 33, pp. 245-51.

Vinge, Vernor (1993), "The Coming Technological Singularity: How to Survive in the Post-Human Era," *Vision-21: Interdisciplinary Science and Engineering in the Era of Cyberspace*, Geoffrey Landis, ed., NASA Publication CP-10129, pp. 11–22.

World Bank (1998), *World Development Report 1998/1999: Knowledge for Development*, Oxford University Press. <https://openknowledge.worldbank.org/handle/10986/5981>

World Bank (2011), *The Changing Wealth of Nations: Measuring Sustainable Development in the New Millennium*. <https://openknowledge.worldbank.org/handle/10986/2252>

World Commission on the Social Dimension of Globalization (2004), *A Fair Globalization: Creating Opportunities for All*. <https://www.ilo.org/public/english/wcsdg/docs/report.pdf>

Wu, Tim (2018), *The Curse of Bigness: Antitrust in the New Gilded Age*, Columbia Global Reports.

Appendix: Isoquants and Factor Price Frontier for Cobb-Douglas Production Functions

Consider the Cobb-Douglas production function

$$F(K, L) = AK^\alpha L^{1-\alpha}$$

Isoquants

An isoquant is the locus of (K, L) that satisfies $F(K, L) = AK^\alpha L^{1-\alpha} = \bar{Y}$. Solving for L as a function of K , we obtain

$$L(K) = \left(\frac{\bar{Y}}{AK^\alpha} \right)^{\frac{1}{1-\alpha}}$$

Increases in technology A move the isoquants inwards.

In the space of log factor inputs, isoquants are linear,

$$\ln L(K) = \frac{1}{1-\alpha} \ln \bar{Y} - \frac{1}{1-\alpha} \ln A - \frac{\alpha}{1-\alpha} \ln K$$

Factor Price Frontier

The factor price frontier represents the combinations of factor prices (w, r) that would prevail in competitive markets as we move along a given isoquant. Equivalently, it can be expressed as the maximum wage w for a given return on capital r that is consistent with profit-maximizing behavior under the given technology. For constant-returns-to-scale production functions, this is equivalent to the locus of factor prices (w, r) for different capital/labor ratios. In our Cobb-Douglas example,

$$r = \alpha A(L/K)^{1-\alpha}$$

$$w = (1 - \alpha)A(K/L)^\alpha$$

Combining the two equations to eliminate $(K/L) = \left(\frac{r}{\alpha A} \right)^{-\frac{1}{1-\alpha}}$, we can express the factor price frontier as a function

$$w(r) = (1 - \alpha)A \left(\frac{r}{\alpha A} \right)^{-\frac{\alpha}{1-\alpha}} = (1 - \alpha)\alpha^{\frac{\alpha}{1-\alpha}} A^{\frac{1}{1-\alpha}} r^{-\frac{\alpha}{1-\alpha}}$$

Increases in technology A shift out the factor price frontier, reflecting that it is technologically feasible for both factors to be compensated at a higher rate. In log space, the factor price frontier for Cobb-Douglas production functions is linear.