



WP/21/166

IMF Working Paper

Technological Progress, Artificial
Intelligence, and Inclusive Growth

by Anton Korinek, Martin Schindler, and Joseph E. Stiglitz

***IMF Working Papers* describe research in progress by the author(s) and are published to elicit comments and to encourage debate.** The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

I N T E R N A T I O N A L M O N E T A R Y F U N D

IMF Working Paper

Institute for Capacity Development

Technological Progress, Artificial Intelligence, and Inclusive Growth**Prepared by Anton Korinek (University of Virginia), Martin Schindler (IMF), and Joseph E. Stiglitz (Columbia University)****Authorized for distribution by Valerie Cerra¹**

June 2021

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate. The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

Abstract

Advances in artificial intelligence and automation have the potential to be labor-saving and to increase inequality and poverty around the globe. They also give rise to winner-takes-all dynamics that advantage highly skilled individuals and countries that are at the forefront of technological progress. We analyze the economic forces behind these developments and delineate domestic economic policies to mitigate the adverse effects while leveraging the potential gains from technological advances. We also propose reforms to the global system of governance that make the benefits of advances in artificial intelligence more inclusive.

JEL Classification Numbers: E0, O1, O3**Keywords:** Artificial Intelligence; Inclusive Growth; Economic Development**Authors' E-Mail Addresses:** anton@korinek.com, mschindler@imf.org, jes322@gsb.columbia.edu

¹This paper is an extended version of a chapter that has been accepted for publication by Oxford University Press in the forthcoming book on “How to Achieve Inclusive Growth,” edited by V. Cerra, B. Eichengreen, A. El-Ganainy, and M. Schindler due for publication in 2021. Parts of this paper subsume ideas presented in previous work by Korinek and Stiglitz (2021a) under the title “Artificial Intelligence, Globalization, and Strategies for Economic Development.” We thank Avital Balwit, Andy Berg, Valerie Cerra, Sharmini Coorey, Romain Duval, Barry Eichengreen, Katya Klinova, Nikola Spatafora, Jeromin Zettelmeyer, participants in the IMF IG seminar series, and numerous IMF colleagues for insightful comments and suggestions; David Autor, Adrian Peralta-Alva and Agustin Roitman for helpful data and charts; and Jaime Sarmiento for excellent research assistance. Financial support from the Institute for New Economic Thinking is gratefully acknowledged.

Contents

Abstract	2
I. Introduction.....	4
II. Downside Risks of Technological Progress	6
A. Labor-Saving Technological Progress.....	9
B. Resource-Saving Technological Progress.....	12
C. Information, Digital Monopolies and Superstars	13
D. Misguided Technological Progress.....	14
E. Broader Harms Associated with AI.....	16
III. Evaluating the Uncertainties and Opportunities.....	17
A. Uncertainty about the Pace and Scale of Progress.....	17
B. The Productivity Puzzle: Are We Really in an Era of Unprecedented Innovation?	18
C. Putting AI in the Broader Context of Development.....	19
IV. Lessons from Past Technological Transformations	21
A. Pre-Industrial Revolution.....	22
B. Industrial Revolution.....	22
C. Manufacturing-Based Export-Led Growth	24
D. What is Different This Time	26
V. Domestic Policy Responses.....	27
A. Taxation, Redistribution, and Government Expenditures.....	27
B. Pre-Distribution.....	31
VI. Global Governance.....	35
A. A Global Tax Regime for the Digital Age.....	36
B. Global Competition Policy.....	36
C. Intellectual Property Rights.....	37
D. Data and Information Policy.....	38
VII. Conclusion	39
References	41

Figures

Figure 1. Productivity and Earnings Growth.....	7
Figure 2. Income Inequality over Time, 1960-2015.....	8
Figure 3. Rising Wage-Skill Premia in the US	10
Figure 4. Employment Trends in Routine vs. Non-Routine Jobs	10
Figure 5. Population Aging and Automation.....	20
Figure 6. Federal Minimum Wage in the US, 1938-2020.....	23
Figure 7. Technology Adoption Lags	26

I. INTRODUCTION

Advances in artificial intelligence (AI) and related forms of automation technologies have led to growing fears about job losses and increasing inequality. This concern is widespread in high-income countries. Developing countries and emerging market economies should be even more concerned than high-income countries, as their comparative advantage in the world economy relies on abundant labor and natural resources. 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 could 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 many decades, there was a presumption that advances in technology would benefit all—embodied by the trickle-down dogma that characterized neoliberalism. And for some time, this presumption was in fact justified. For example, for the three decades following World War II, the US economy and many other high-income and developing countries experienced broadly shared increases in living standards. However, over the past half-century, output growth and median worker incomes started to decouple.

Moreover, economic theory cautions that technological progress is likely to 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, then entire nations may be worse off except if the winners within one country compensate the losers in other countries, which seems politically very difficult.

This paper argues that concerns about whether technological progress leads to inclusive growth are indeed justified – and that especially developing countries may face a stark new set of challenges going forward. However, we propose policies that can mitigate the adverse effects so that advances in technology lead to a world with greater *shared* prosperity. This will require new domestic policies and development strategies as well as strong international cooperation and a rewriting of the global rules governing the information economy.

We start by laying out the key properties of AI and related automation technologies that underlie the concerns about recent technological progress.² AI is likely to be labor-saving and resource-saving, devaluing the sources of comparative advantage of many developing countries and deteriorating their terms of trade. Being an information technology, AI also tends 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. Moreover, under reasonable assumptions, the rate and direction of technological progress chosen by the market are generally suboptimal (Korinek and Stiglitz, 2019). This creates the possibility of

² AI is frequently defined as the ability of machines to perform tasks that previously required human intelligence. This implies that AI carries by its nature the potential to replace human labor. A number of arguments put forward in this paper are applicable not only to AI but also to other forms of technological progress.

steering innovation in AI and other technologies in directions that are more beneficial to humanity at large, for example, preserving the planet or creating satisfying employment opportunities, rather than substituting for labor and creating more unemployment and inequality.

Taking a step back, we evaluate to what extent the discussed concerns about technological progress are justified, given what we know at present. There is vast uncertainty about the impact of AI, 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.

In this context, we 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. We also analyze how the forces generated by progress in AI interact with other recent developments, in particular with the recovery from COVID-19, with secular population dynamics, and with the need for a Green Transition.

Despite the uncertainties surrounding AI, its potentially dramatic consequences suggest that we should steer our own research in directions where the expected social value added of economic analysis is greatest: we need 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, making it a mere blip in the history of human civilization. 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 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 erase much of the gains that the developing countries have made in recent decades.

However, 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 will require a new development model that follows a more multi-pronged strategy to replace the manufacturing-based export-led growth model.

The key policy question is how countries can improve the likelihood of benign outcomes from technological progress. This is especially pertinent for developing countries, but it is also a challenge for advanced economies to develop policies that ensure that technological

advances lead to broadly shared prosperity and that their adverse effects are mitigated. We delineate here a number of such policies. Taxation and redistribution are a first line of defense 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 both income to workers and a valuable social return – for example, investments in education or infrastructure are labor intensive and enhance human capital and the physical infrastructure of countries, both of which are important in bridging the digital divide and ensuring that all citizens can participate in the opportunities afforded by digital technologies.

To replace the manufacturing-based export-led growth model, developing countries will need to steer technological progress and technology adoption in new directions, in part by leveraging the opportunities that 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 developing countries can benefit from advances in AI and other information technologies while addressing the downsides of these new technologies. We need to design a global tax regime for the digital age that enables countries to raise taxes on transactions that occur within their borders. Competition policy is also increasingly a question that transcends national borders as the footprint of the digital giants is global and authorities in their countries of origin do not face the correct incentives to ensure a competitive marketplace. Intellectual property regimes need to be adapted so they are attuned to the needs and circumstances of developing countries. Moreover, information policy including the regulation of data needs to be discussed at the supra-national level to provide a voice to developing countries that could otherwise not influence the design of such policies.


The remainder of this paper is organized as follows. In the second section, we provide an overview of the downside risks of technological progress, with special emphasis on potential AI-induced economic disruptions; in the third, 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 related technologies.

II. DOWNSIDE RISKS OF TECHNOLOGICAL PROGRESS

Many technology optimists suggest that productivity gains go hand in hand with real wage gains. This presumption that technological progress would benefit all was also embodied by the trickle-down dogma that has characterized neoliberalism. However, the presumption was supported neither by theory nor evidence; indeed, economic theory has always held that advances in technology do not necessarily benefit all and may create winners and losers. The data (Figure 1) show that in recent decades, many countries have experienced episodes


during which wages lagged productivity growth. Moreover, as we argue below, even where average wages did keep up with productivity, median wages may not have, and there is a risk that any positive gains seen in the past may not continue.

Figure 1. Productivity and Earnings Growth



Source: OECD, Bureau of Labor Statistics, IMF staff calculations.

Figure 2 illustrates that the income gains associated with technological progress have been highly unevenly distributed. In the US and other high-income countries, most of the benefits of growth have gone to those at the top, resulting in widening income inequality in most advanced economies since the early 1980s, reversing an earlier downward trend in many countries.



Source: World in Data; data depict the cumulative share of national (pre-tax) income accruing to the top 5 percent of income earners (<https://wid.world/>).

How can we reconcile this with economic theory? In the context of a competitive economy, we can think of technological progress as moving out the production possibility frontier: one can get more of any output for a given amount of inputs. But this increase in production possibilities does not tell us how the gains from progress will be distributed. In our simplest economic models, for example, if we assume a competitive economy with a Cobb-Douglas production function, relative shares are fixed.

However, in the more general case, technical change may change the distribution of income, so that, for instance, labor gets a smaller share of a larger pie. If its share decreases enough, workers could even be worse off. Whether wages increase or decrease depends on what happens to the demand for labor at existing wages. Using the terminology first introduced by Hicks, technical change that leads to a decrease in the relative share of labor is called capital-biased; if it leads to a decrease in the share of unskilled labor, it is called skill-biased; if it leads to an outright reduction in wages, it is called labor-saving. The US, for example, has experienced routine-biased technological change that has replaced workers engaged in both manual and cognitive routine activities since the 1980s and that has contributed to the hollowing out of the middle class (Autor et al., 2003).

Korinek and Stiglitz (2019) show that the distributive effects of innovations can be seen as generating quasi-rents –aside from delivering direct gains to innovators, innovations lead to changes in factor demands, for example lowering demand for unskilled labor and raising demand for skilled labor, and the affected workers experience gains or losses. The winners of progress (e.g. the skilled workers in our example) experience these gains without having contributed to the innovation, obtaining quasi-rents, whereas the losers experience losses without any fault of their own. That, in turn, has an important implication: governments can capture some of the quasi-rents by taxing the winners and redistributing it; and given the nature of the gains, governments may even be able to raise taxes in ways that have no or limited distortionary effects, for example if the winners include owners of fixed factors such

as land. Thus, “managed” technological progress could allow for Pareto-improving outcomes.

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 cost is borne in another, a Pareto improvement would require that the winners compensate the losers across national boundaries. Today, such cross-border transfers are voluntary and limited.


As a result, the fruits of technological progress will be unequally shared; but more troublesome is that while some countries may gain a great deal, others will lose. These differences will be reflected, respectively, in improvements and deteriorations of countries’ terms of trade. 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.

A. Labor-Saving Technological Progress

Many observers are concerned that AI may be labor-saving, that is, cause a decline in the demand for labor at existing factor prices. If this occurs, equilibrium wages will decrease and workers will be worse off.

As we have noted, over the past half-century, the US and many other countries seem to have experienced technological progress that was biased against workers with lower levels of education performing routine tasks, sufficiently biased that it may even have been labor-saving in that segment, reducing such workers’ real incomes. For example, Autor et al. (2003) observe that from the 1970s to the 1990s, while computerization was a substitute for an increasing number of routine tasks, technological change increased the productivity of workers in non-routine jobs that involved problem-solving and complex communications tasks. These changes in technology may have explained nearly 2/3 of the relative demand shift toward college-educated labor over that period. Similarly, more recently, Acemoglu and Restrepo (2020) estimated significant adverse employment and wage effects from the introduction of industrial robots in the US, concentrated in manufacturing and among routine manual, blue-collar, assembly, and related occupations, helping to explain the dramatic increase in wage dispersion across skill groups over the past five decades (Figure 3).


Figure 3. Rising Wage-Skill Premia in the US
(Real Wages of Full Time U.S. Male Workers, 1963 = 100)



Source: Autor (2019).

This job polarization in terms of wages has also been reflected in relative employment dynamics. Employment in nonroutine jobs has continued to grow steadily in the US, while that in routine jobs has stagnated, or in some periods declined, since around 1990, contributing, as we have noted, to a “hollowing out of the middle” (Figure 4).³ OECD (2019) note that middle-skilled jobs may be the most prone to both automation and offshoring, as they most encompass routine tasks that are relatively easy to automate (or offshore).⁴

Figure 4. Employment in Routine vs. Non-Routine Jobs
(Persons, millions)



Source: Current Population Survey and Federal Reserve Bank of St. Louis.

³ As can be seen in the figure, the COVID shock in 2020 has clearly accelerated the trend, at least temporarily, giving rise to a large decline in employment in routine manual jobs but only a modest dip in nonroutine cognitive jobs.

⁴ See also WEO (2018) for a broader review of employment and earnings dynamics across countries and sectors.

Standard models of aggregate production functions with skilled and unskilled labor-augmenting progress and capital-augmenting progress can generate the observed patterns of movements in factor prices and shares, depending on patterns of progress as well as elasticities and cross-elasticities of substitution. Acemoglu and Restrepo (2019a) formulate a particular model in which the displacement of workers by robots will reduce the labor share of income and may be labor-saving if the productivity gains from the robots are modest. Berg et al. (2018) focus on the differential effects of technological progress across worker groups and shows that technological progress may be unskilled-labor-saving because that type of labor is easily substituted for by robots; by contrast, high-skilled labor is likely complementary to robots and will benefit from technological progress; as a result, technological advances risk bringing about large increases in inequality. Automation may also worsen inequality along other dimensions—for example, in sectors where women occupy more routine jobs (Brussevich et al., 2018).

Even if technological progress is labor-saving in the short run, it may also trigger additional accumulation of capital that is complementary to labor, benefiting labor in the long run. For example, Stiglitz (2015) and Caselli and Manning (2019) show that in an economy with capital and labor only, in which long-run capital accumulation is determined by an exogenous interest rate, labor will always gain.⁵ Ultimately, however, impacts on inequality depend on whether there are other scarce limiting factors in the economy, for example, natural resources or land, which would benefit from technological progress and ultimately become more scarce as the factors “capital” and “machine-replacing labor” become more abundant and cheaper. Indeed, Korinek and Stiglitz (2021a) show that if this is the case, then, without government intervention, labor may lose out from technological progress even in the long run.

At a global level, similar dynamics may play out. 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 particularly 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. Labor-saving progress may not only create winners and losers within the affected developing countries, but it may make entire countries on net worse off. Alonso et al. (2020) find that improvements in the productivity of “robots” could drive divergence, as advanced countries benefit from computerization more given their higher initial capital stock.

However, it is also conceivable that other forms of advances in technology could benefit workers: intelligence-assisting devices and algorithms (IA) may be complementary to labor rather than substituting for it, thus enhancing the prospects of labor. Innovations that fall into this category may include augmented reality (AR), machine learning (ML) algorithms that help analyze complex data, and other forms of integration of AI with humans.⁶ Automation technologies frequently affect particular *tasks* but not (entire) jobs, which consist of multiple

⁵ The result is intuitive: the dual to the production function is the factor price frontier. Technological change shifts out the factor price frontier, implying that if the interest rate is unchanged, wages must increase.

⁶ One extreme example is Elon Musk’s *Neuralink* which aims to achieve a symbiosis of humans and AI by surgically implanting technology into the brain.

tasks (see, e.g., Acemoglu and Autor, 2011)—IA innovations may help workers be more productive in their jobs by taking over, or improving, certain tasks. 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, at least not yet.

Driverless trucks provide another example: truck driving provides significant employment opportunities for men with only a high school education so there is understandably concern for the disruption that self-driving trucks might bring about. But truck drivers also perform a number of related tasks – they fill orders, load and unload, monitor the truck, and more – not all of which may be easily automated. More generally, most jobs have multiple dimensions and consist of multiple tasks. With some tasks automated, workers will be able to devote more attention to, and perform better at, those tasks that are not. Importantly, both AI and IA imply extensive restructuring of the economy.

The central concern of this paper remains: there may be a reduction in the demand for labor, especially for unskilled labor. We will further evaluate whether or not these fears are justified below in Section 3. If, however, it turns out that AI is labor saving, and especially if it is unskilled labor saving, the consequences for developing countries would be severe. This is the “resource” which constitutes their comparative advantage and in which they are relatively rich. The convergence in standards of living between developing countries and developed that has marked the past half century would be arrested, even reversed. It would also 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 would exacerbate those inequalities – and developing countries often lack the institutional capacities to counteract them.

B. Resource-Saving Technological Progress

Another type of progress that is of great concern to some developing countries is resource-saving technological progress. This has gotten less attention than labor-saving progress so far (e.g., Solow, 2009), but 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 reduce the demand for depletable natural resources and lower carbon emissions. Examples include algorithms that optimize efficiency in data centers or that make transportation networks more efficient. Technologies that enable telework may also reduce the carbon footprint of workers.⁷ Thus, such resource-saving innovations may have adverse distributional effects on developing countries that have a comparative advantage in natural resources, and that have specialized in exporting them. The impact on exporters of different

⁷ As always, calculating the full consequences of a new technology on the demand for any natural resource, or carbon emissions, is complex. It must be done on a full life-cycle basis, incorporating initial investment, maintenance, as well as day-to-day operations. That said, for instance, data centers running cutting-edge AI applications are typically energy-intensive and may lead to increases in demand for electricity and depletable natural resources. Still, on net, it is likely that the demand for carbon-based energy sources will decrease. Some natural-resource-rich economies may benefit, such as those rich in rare earths or other metals that are inputs in the production of batteries, microchips, solar panels, wind turbines etc.

types of natural resources may be quite different – for example, exporters of carbon-based energy will fare differently from exporters of rare earth metals.

Consider oil-exporting countries, which have already experienced many 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. The challenges of addressing global inequality under such a scenario would be an order of magnitude larger than they are even today, posing a test for the global community. A number of oil-exporting countries rely on their export revenue to buy food and other basic essentials—if they lose their ability to export oil, the consequences would be dire. Thus, as in the case of labor-saving technological progress, the world as a whole may be better off—in this case by undoing resource scarcity and reducing climate change—but not all countries would benefit.⁸

C. Information, Digital Monopolies and Superstars

So far we have considered the effects of technological change in a competitive environment. However, the rise of AI and other information technologies may also lead to greater concentrations of market power. As a result, the economy may move to an equilibrium that is less competitive and more distorted by market power, with greater rents for dominant firms. Actors with market power will use that power to advantage themselves. The resulting distortions may offset part of the benefits of innovation, exacerbating the adverse distributive effects of labor-saving or resource-saving innovation. With any inequality-averse social welfare function, societal welfare could decrease.

While the assumption of competitive markets often provides a useful benchmark, that model becomes less appropriate as one considers an economy that is dominated by AI. It is hard to conceive of an AI economy being competitive, or at least well-described by the standard competitive equilibrium model.

There are several reasons why advances in AI may 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 codes or ML algorithms 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 creating large monopolies, sometimes also called “superstar” effects (see, e.g., Korinek and Ng, 2019, and Stiglitz and Greenwald, 2014a).

⁸ In addition, many fossil-fuel-dependent countries have not yet diversified their export base, and may face limited options to diversify into job-rich manufacturing growth given this sector’s vulnerability to automation. See, e.g., Peszko et al. (2020).

Some authors have identified 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.

Such superstar and monopoly effects are likely to play out not only at a company level but also at a country level, and they are likely to be particularly severe in the context of 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, with the notable exception of China – one of the leaders in AI. Moreover, to the extent that firms or countries can protect their knowledge, the resulting monopolization of knowledge may also impede the catching-up process. Importantly, even if competitors could “steal” a superstar’s knowledge, this may not necessarily be sufficient as the superstars can continuously improve their algorithms based on their users’ data, thus remaining, perhaps permanently, ahead. 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 too may change with AI.

Some observers suggest as a silver lining for developing countries in that ML technologies are reliant on data and that more diverse data contain more information. Thus, selling data might generate some income for developing countries. However, this is unlikely to make up for their lost income as the marginal return to more diverse data may be limited. Moreover, future advances in ML algorithms may make them less reliant on large quantities of data and instead require more specific, tailored data.

D. Misguided Technological Progress

Economic theory has illuminated why the nature of innovation (e.g., the factor bias) 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 the literature on induced innovation (e.g. Kennedy, 1964; von Weizsäcker, 1966; Samuelson, 1965; Atkinson and Stiglitz, 1969; Acemoglu, 1998, 2002; Stiglitz, 2006). 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. The market may even provide incentives for innovations that reduce efficiency by absorbing more resources than they create for society, as may be the case, for example, for high-frequency trading. This calls for policy to actively steer technological progress, as we will discuss further below.

The fundamental problem is that knowledge is a public good, in the Samuelsonian sense. If it is to be privately financed and produced, there must be inefficient restraints on the use of knowledge, and those restraints typically also give rise to market power. If there are no restraints on the use of knowledge, then innovators cannot appropriate the returns to their production of knowledge, and so they will have little incentive to innovate.⁹ When knowledge is produced as a by-product of learning or investing, the inability to fully appropriate all the learning benefits will lead to under production or underinvestment in sectors of the economy associated with high learning and learning spillovers. As Greenwald and Stiglitz (2006) and Stiglitz and Greenwald (2014a) point out, this has important implications for developmental policy, providing a rationale for industrial and trade policies.¹⁰

More recent literature has drawn attention not only to biases in the level and pace of innovation but also to the direction. In economies with incomplete risk markets and imperfect and/or asymmetric information (i.e., in all real-world economies), the equilibrium is not constrained Pareto efficient, and prices do not necessarily give the “correct” signal to innovators on the direction of innovation. There are pecuniary externalities that matter.¹¹ For instance, in the Shapiro-Stiglitz (1984) efficiency wage model, where unemployment acts as a disciplining device to discourage shirking in the context of a labor market with imperfect and costly monitoring, there will be too much labor-augmenting technological progress, resulting in too high a level of unemployment (Stiglitz, 2006). There are multiple other biases, for example, towards innovative activities in which intellectual property rights are more easily secured, and in drugs, to me-too innovations, where private returns can markedly exceed social returns.

Markets do not care about income distribution. Market forces may drive economic decisions towards efficiency—in the narrow, microeconomic sense—but will not give any consideration to the distributive consequences. Recent contributions, however, have emphasized that overall economic performance can be affected by inequality (Ostry et al., 2019; Stiglitz, 2013); obviously, individual entrepreneurs will not take into account this macroeconomic externality, and accordingly the market will be biased towards producing too much labor-saving innovation, creating a role for redistributive policies. In addition, Korinek and Stiglitz (2020) show that in the presence of constraints on redistribution, policy can improve welfare by steering innovation to take into account its distributive 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

⁹ There is a large literature on the welfare economics of innovation, dating back to Arrow (1962a). Stiglitz (1975a, 1987) drew attention explicitly to the public good aspects of knowledge, and the similarity between the economics of information and the economics of knowledge. See also Romer (1986).

¹⁰ The inefficiencies in economies with learning by doing was first noted by Arrow (1962b).

¹¹ See Greenwald and Stiglitz (1986). These, in turn, give rise to macroeconomic externalities, the consequences in the context of innovation have been studied by Korinek and Stiglitz (2019).

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.¹²

What is most relevant for developing countries is that these distributive implications extend 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 consider. 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 other countries. In developing countries, a key question is about *adopting* appropriate technologies rather than innovating, but the same kind of analysis that 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 further adaptation to local circumstances, and the benefits and costs of the technological evolution are not fully appropriated, for example, in the process of learning by doing. These concerns have long been at the center of concern of industrial policy.

E. 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 economies, an important concern is that the international community may address them in a way that does not reflect the priorities and needs of developing countries. Policymakers in many countries are beginning to discuss appropriate regulatory regimes and a 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 (e.g., Ding et al., 2018; Sacks, 2018), 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-vaccine movement.

¹² More generally, the direction of innovation is affected by the *share* of the factor. If the elasticity of substitution is high, a lower factor price will be associated with an increased factor share, and this can induce greater efforts at increasing the productivity of that factor. In that case, the equilibrating force just described does not arise, and the opposite occurs (Stiglitz, 2014).

III. EVALUATING THE UNCERTAINTIES AND OPPORTUNITIES

A. 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 great degree of uncertainty about the pace of change and the magnitude of the coming disruption, even among experts in this area. Some economists (e.g., Gordon, 2016) 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. In fact, Gordon (2016) 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 (GPT)—that has the potential to revolutionize every sector of the economy (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 AI may eventually advance to a point where AI systems reach human levels of general intelligence. This may imply that they can also do research, design better versions of themselves and thereby recursively self-improve, giving 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).¹³ The concept of such a singularity has been popularized by Good (1965), Vinge (1993) and Kurzweil (2005), and is being increasingly discussed among economists (e.g. Nordhaus, 2015; Aghion et al., 2017). Predictions of when such a chain of events might occur, however, continue to be perpetually revised—Armstrong et al. (2014) note that over the past six decades or so analysts have continued to expect “the development of [general] AI [to occur] within 15-25 years from whenever the prediction is made.”¹⁴

This last perspective emphasizes that AI-driven machines may not only be physically stronger than humans and better and faster at processing information, but in an increasing number of domains, they may also learn better and faster than humans.¹⁵ Thus, AI may be much more disruptive than a “mere” GPT; AI programs are increasingly replacing tasks previously performed by humans. If machines can engage in all tasks that have traditionally been performed by labor, and if they can do so at ever lower cost, then traditional labor would eventually become redundant, with the marginal product of human labor possibly falling so low that it no longer covers the subsistence cost necessary to keep a human alive (Korinek and Stiglitz, 2019). This would represent the extreme case of labor-saving

¹³ As Vinge (1993) noted: “Within thirty years, we will have the technological means to create superhuman intelligence. Shortly after, the human era will be ended.” However, it should be noted that general AI does not in itself imply the singularity (e.g., Walsh, 2016).

¹⁴ Responding to Kurzweil's (2005) thesis that “The Singularity is Near,” Walsh (2016) provides arguments for why “The Singularity May Never Be Near.”

¹⁵ There is even a perspective that holds that AI-powered machines could become agents of their own (Korinek, 2019).

innovation: it is in fact labor-replacing innovation – employing labor would become a strictly dominated technology.¹⁶

We discussed earlier some studies examining which 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 14% of all jobs (OECD, 2019) to an estimate of 20–25% (Bain and Company, 2018) and almost 50% 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 in certain industries and among certain groups of workers, specifically among unskilled and routine jobs. Knowing what fraction of all jobs will be lost to AI therefore does not necessarily provide a good metric of the impact on income distribution, and especially so in the short run.

Applying our earlier insights on steering innovation to economic research, economists should steer their research in directions where the expected social value added of economic analysis is greatest, that is, where it has the highest welfare impact.

Even if some of the described scenarios have a relatively low probability, it is important to think particularly hard about events that will be highly disruptive to society, to think through the consequences, and to prepare for how we might ameliorate some of the more adverse effects. Extensive labor replacing innovation would be such an event. Even if one places a relatively low probability on such an event—and one may argue that it is not actually a low-probability event—the associated social repercussions would be sufficiently large that it makes sense to focus attention on such an event. Studying scenarios that pose the most adverse social impacts would better prepare economies to deal with them when they occur—and they also provide valuable lessons for scenarios in which the impact is less stark.

B. 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 are the increases in innovation not showing up in GDP data? This is analogous to the puzzle of missing productivity growth from computerization that Bob Solow described in the 1980s when the GPT of the time – computers – spread throughout the economy (Solow, 1987). 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).

¹⁶ Note that this is in contrast to a long tradition in the traditional economics literature that viewed labor as an essential input for any production process.

Another part of the explanation of the productivity puzzle is that there are difficult measurement problems. Many recent technologies may have led to increases in societal welfare that are not captured by GDP (see e.g. Brynjolfsson, 2020). For example, when online services are exchanged against “eyeballs,” i.e., when users are exposed to advertisements instead of paying for services, the benefits to consumers are not included.¹⁷

C. Putting AI in the Broader Context of Development

There are several 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, which is likely to accelerate the automation of jobs that require physical interaction (see e.g. Korinek and Stiglitz, 2021b). 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.¹⁸

Population Dynamics

Population dynamics will interact in important ways with labor-saving or -replacing technologies (see e.g. Varian, 2020). In countries with rapidly growing working-age populations, such as in many African countries, lots of new jobs will have to be created to maintain a given employment rate. Advances in automation that are developed in high-income countries and easily deployed around the world will make this more difficult. However, the large supply of labor may slow down the development and adoption of automation technologies within such countries (although the evidence in several countries suggests that at least in large export-oriented manufacturing, the technologies employed are remarkably similar to those in advanced countries; see Rodrik, 2011).¹⁹ Moreover, young populations also generate significant demand for education, which in turn creates jobs. Overall, even countries like India face difficulties in creating enough formal sector jobs to keep pace with the growing working age population. The faster growth of population makes capital deepening more difficult, slowing the pace of growth in income per capita.


¹⁷ The measurement problems are still more complicated: advertising is an “intermediate” product and does not directly enter into the value of the final goods and service that constitutes GDP. If advertising were a normal input, and markets were competitive, an increase in the efficiency of production of an intermediate good would be reflected in a lowering of the final goods price, and that in turn would be associated with an increase in GDP. Better advertising engines may, as we noted earlier, actually increase market power and decrease overall economic efficiency. Moreover, they may induce an adverse redistribution, lowering welfare still more.

¹⁸ Any innovations to deal with Covid-19 will still be available in the post-Covid-19 world. Moreover, the development of research strategies in response to Covid-19 may set in motion a process of “learning to learn,” learning better how to innovate in human-replacing dimensions. See Atkinson and Stiglitz (1969) and Stiglitz (1987b).

¹⁹ This would not, of course, be true if the factor price equalization theorem held. More generally, differences in domestic factor ratios do not necessarily align well with differences in factor prices.

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. Overall, the evidence suggests that aging societies adopt new technologies and automate (Acemoglu and Restrepo, 2019b and Figure 5).²⁰

Figure 5. Population Aging and Automation



Source: Reproduced from Acemoglu and Restrepo (2019b).

Note: Aging is measured as the change in the ratio of workers above 56 to workers aged 21-55; robots are measured as the number of industrial robots per thousand workers. Based on regression analysis in Acemoglu and Restrepo (2019b) controlling for a number of covariates.

The Green Transition

A third 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. Without global policies to save our planet, developing countries will experience some of the largest losses from global warming.

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. The Green Transition is similar to resource-saving innovation and risks undermining the standard of living of oil-exporting countries, among which there are a number of low-income countries.

²⁰ There are countervailing forces to the scarcity of labor associated with a declining working age population. A younger population may be more tech savvy, better able to pick up, adopt and adapt to new technologies. The figure suggests that the scarcity effect dominates. There are other factors too that play a role in robotization.

There is also an important complementarity between the Green Transition and AI: the Green Transition is likely to increase the demand for labor which could offset some of the negative effects on labor demand of automation and AI. Indeed, given the labor needed for the Green Transition, the labor replacement due to automation and AI in many activities, including manufacturing, could be considered a fortunate development enabling countries to better address the challenges of climate change. There is thus an inherent tension in frequent claims that on the one hand economies cannot afford to mitigate climate change (i.e., that there are insufficient resources), and on the other hand concerns over a potential crisis with a surplus of labor arising from labor-saving AI.^{21,22}

However, we do face challenges in how to channel surplus resources into what is required for the Green Transition. Some of the skill sets of those labor resources freed up by technological progress will differ from those needed in the Green Transition, although Louie and Pearce (2016) argue that the retraining costs would be moderate, and many of the investments (such as installing solar panels) require only limited skills.

There may be institutional constraints that make it difficult to reallocate capital towards green investment. 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 in between are short-term financial markets. Local, national, and multilateral Green Development Banks may be helpful in financing the private green transition. Better disclosure to investors of risks associated with “brown” investments (i.e., ones that contribute to pollution) and changes in fiduciary standards for asset managers towards their investors, would help move resources into green investments. Of course, without strong incentives, provided by price signals and environmental regulatory constraints, incentives for green investments and innovation will be greatly attenuated.

IV. LESSONS FROM PAST TECHNOLOGICAL TRANSFORMATIONS

To grasp the historical nature of what is going on, it is necessary to put the advent of AI and related technologies in the context of the broader history of technological progress. Humanity spent much of its history at a Malthusian stage. The Industrial Revolution started a little 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 – 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

²¹ There is a similar dissonance between those who argue that the economy faces secular stagnation and those who say there are not the resources required for a rapid green transition.

²² Over the long run, the effect of the green transition on the demand for labor is more problematic. While many of the green technologies have higher upfront costs, maintenance costs are markedly lower, and not only are life-cycle carbon emissions lower, but so is labor usage.

before the Industrial Revolution and the early decades of the revolution itself serves as a reminder.

A. 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, several 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 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? Thus, populations may decline not as a result of choice, as in many developed countries, but from Malthusian dynamics. In today's globally connected world, that presents ugly alternatives: Will the rich countries simply look away, as they see this suffering and near-starvation in poorer countries? Will they create ever-increasing barriers to stave off the inevitable pressures of migration?

B. Industrial Revolution

The Industrial Revolution marked the beginning of rapid growth in high-income countries. After centuries in which standards of living had been stagnant, growth started to increase markedly. It transformed the world. 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, not everyone benefits. It can give rise to very large disruptions during the transition. In the short run, there was significant social upheaval from the industrial revolution—Charles Dickens' novels make it clear that not everyone prospered. In the UK, some people 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. Even indicators such as life expectancy initially went down. Looking at those who suffered, the Industrial Revolution was clearly 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. Eventually, governments played an important role in mitigating the adverse effects, including the problems posed by urbanization, such as challenges in sanitation, environmental degradation, public health, infrastructure, and congestion.

Government took a strong role too in advancing the positive effects of the new economy. Education was an important element in creating a productive workforce – it was therefore also in the interests of capitalists, and public education received broad public support.

In high-income countries, institutions related to labor legislation, unionization, and social safety nets 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 today is taken for granted, but in 1900, it was not obvious if meaningful labor legislation would ever be enacted. Strikingly, some of the tough political battles that made the adoption of such legislation problematic a century ago are playing out once again in the United States, where there has been an erosion of protections, for example, those associated with minimum wages, health and safety standards, or overtime pay, among others.

Figure 6. Federal Minimum Wage (adjusted for inflation) in the US, 1938-2020



Source: Federal Reserve Economic Data (FRED).

These labor market reforms helped support the structural transformation that occurred with the rise of manufacturing, and they showed that equality and growth are complementary (e.g., 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.

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 fifty years ago (Figure 6), 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 in the form of rising wages, but also, in parallel, unprecedented political power. However, this power has been eroded more recently (see e.g. Boix, 2019). In simple models of democracy, the median voter (or more broadly, the “majority”) determines political outcomes. But the evidence is that that model provides a poor description of the outcomes of the political process. For instance, the majority of voters want a more egalitarian society (see, e.g., chapter 1 in Cerra et al., forthcoming). But in recent decades, in many countries, the political and economic rules have evolved in the opposite direction, giving more influence to the power of “money”.^{23,24}

C. 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 multifold.

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 seen by many to have a comparative advantage in agriculture. It instead pursued a strategy of creating its own dynamic comparative advantage via an industrial policy that led it towards industrialization. That model served most of East Asia remarkably well, in a way few had anticipated (e.g., Myrdal (1968) who predicted that Asia would never develop). See also Aghion et al. (2021).

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

²³ For example, based on data for 1981-2002, Gilens (2005) finds that in the US actual policy outcomes strongly reflected the preferences of higher-income groups, with little relationship to the preferences of the poor or middle-income citizens. For a broader discussion of the interplay of economic and political inequality, see Stiglitz (2013, 2019).

²⁴ Harari (2017) also explores the implications of super-human artificial intelligence on society and politics.

employment opportunities, foreign exchange, tax revenue—everything that was needed for a quick developmental transition.

While their development trajectory began with taking advantage of their static comparative advantage in cheap labor, and especially cheap unskilled labor, over time, many East Asian countries moved 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 may have been 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.²⁵ 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). A quarter century ago, the World Bank began thinking 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 increase the gap in knowledge and make it more difficult to catch up. While technology adoption lags have declined over the past centuries (Comin and Hobijn, 2010 and Figure 7), the specific nature of AI may reverse that. Cutting-edge AI technology is highly specialized, and improvements are driven to a large degree by learning from large datasets, creating a winner-takes-all dynamic, as we noted earlier. In addition, 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 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. In fact, the exponential nature of growth in AI technology may imply that laggards not only cannot catch up, but that the gap between them and the front runners may grow, compounding the potential adverse

²⁵ The emphasis here is on (traded) goods rather than (non-traded) services—while learning by doing could occur in both, it is the former that drives export-led development. See, e.g., McMillan and Rodrik (2011) who note that non-traded service sector development on its own typically has not had a substantial impact on overall productivity.

effects that developing countries may suffer from labor-saving or resource-saving technological progress.²⁶

Figure 7. Technology Adoption Lags



Source: Comin and Hobijn (2010).

Notes: Technology adoption lag is a mean estimated lag in cross-country technology diffusion (Comin and Hobijn, 2010, Table 2).

D. What is Different This Time

Not only may the AI revolution make it more difficult for developing countries to catch up, the AI revolution may also 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, what risks becoming our “destination”—a service sector economy, marked by greater inequality, with less support for public education and more concentrations of market power—may be less attractive in many ways than the current situation, and the process of getting there may be more disruptive; that is, 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.²⁷ This may reduce support for

²⁶ Stiglitz (2015) models the relationship between technological leaders and followers.

²⁷ We emphasize that the focus here is on the more adverse scenarios, to help prepare policies; should they not materialize, so much the better. We noted countervailing forces—the need for labor for the green transition, that
(continued...)

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. In addition, worker tenure has declined, making it harder to unionize the workforce (Choi and Spletzer, 2012). 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. Unskilled workers in these countries may suffer the most.

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. (See, e.g., Gilens, 2005.)

V. DOMESTIC 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 can be employed to address the effects of technological disruption, both in developing countries and to protect vulnerable segments in advanced economies. 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.

A. Taxation, Redistribution, and Government Expenditures

Among the critical policies to combat rising inequality are those of taxation and redistribution, with a particularly important role for progressive taxation. However, 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. However, this also means that there

even within advanced economies, people may still be needed for service jobs requiring physical proximity and/or the “human touch” (such as elderly care, housekeeping, etc.). Most important, these outcomes are not inevitable: we can steer innovation in a different direction and, as the discussion below will hopefully make clear, there are multiple actions that can be taken to mitigate some of the adverse effects.

is significant scope for developing economies to enhance their tax structures and expenditure systems (e.g., scrapping harmful subsidies and tax exemptions) to build fiscal space for public spending, and improvements in tax capacity. Also, 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 governments 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.²⁸

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, such as land, acts like a lump sum tax, and taxing rents generated by market power and political activity may discourage such rent-seeking, enhancing efficiency.

We have argued earlier that technological progress 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). High on the list of what is desirable to tax are “bads” rather than goods, i.e., Pigouvian taxes on activities and goods that create negative externalities, for example, polluting or carbon-emitting goods. This would contribute to the Green Transition in a dual way, not only by providing tax revenue for public investments but also by correcting market prices to reflect the negative externalities.²⁹

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). But such insurance is not available. One of the functions of social insurance is to socialize these risks that otherwise would have been borne by individuals. But

²⁸ Some are justifiably concerned that digital platforms are in fact very efficient at exploiting workers. 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.

²⁹ See also Pouokam (2021).

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.

Universal Basic Income

Many commentators have responded to concerns about the impact of technological progress on employment by advocating a universal basic income (UBI). While proposals differ in their detail, they typically entail that all individuals are paid a UBI independent of their employment or wealth status, and with a level of UBI payments geared above the poverty line. While such programs would imply formidable fiscal costs, and with it, possibly large distortionary taxes, those could be contained if a UBI replaced other social safety programs (such as social security, welfare, or unemployment insurance systems). By doing so, it would also reduce the overall administration costs.

From a global welfare perspective, a global UBI that was truly “universal” as the name suggests, i.e., that covers all citizens of the world equally, would be most desirable, given the potentially large global implications of AI. Currently, access to prototypes of a UBI is exclusive to 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). But given the limitations on cross-border transfers that have been the center of attention of this paper, a global UBI is clearly still in the realm of fantasy.^{30,31}

However, in the short- to medium-run, the focus should be on creating jobs for everyone who is able and willing to work, especially in light of the earlier discussion of how much labor will be needed for the Green Transition, to provide services to the young, the sick, and the elderly, and to invest in infrastructure. Governments may have a role to play in helping match the need for work and people willing and able to work. However, while a clear need for a UBI may be in a more distant future, there are other policies that may achieve similar objectives to a UBI. For example, one approach to ensuring a modicum of income for all over the long run, with co-benefits of perhaps increasing social cohesion and solidarity, is shared capital ownership (e.g., Solow, 2009): as part of government assistance programs (such as those enacted in the wake of COVID-19 in 2020), firms receiving government help should contribute shares to a sovereign wealth fund—owned by everyone within the nation.

³⁰ Some countries have started to experiment with schemes that have some characteristics of a UBI. E.g., Spain introduced in early 2020 a “minimum vital income” to ensure a guaranteed minimum income for the poorest. However, it is not unconditional, but instead tops up incomes below the minimum income, which may create disincentive effects to continue work in jobs that pay below that threshold. Several other countries have run pilot programs, often on a small scale and/or for a limited time. Overall, these programs appear to indicate that such schemes tend to have little impact on labor supply (see, e.g., <https://www.vox.com/future-perfect/2020/2/19/21112570/universal-basic-income-ubi-map>). Earlier research on a negative income tax in the US suggested that by enabling individuals to search more for a better matching job may actually enhance productivity.

³¹ UBI programs may turn out to be important policies in a future in which labor truly becomes redundant (Korinek and Juelfs, 2021). There is uncertainty over when that future may arise, as the earlier discussions indicated—but given the complexities of transitioning to such a new regime, there may be a rationale for countries to start experimenting with UBI systems.

Similarly, firms that build on or employ innovations that are based in part on government-funded research should be required to do the same.³²

Starting with Keynes (1931), economists have argued that technological progress and automation would in principle enable people to work less and spend more time on more meaningful activities rather than tedious and repetitive tasks – a point also emphasized, for example, by Varian (2020). However, this requires either that wages go up in tandem with productivity growth, unlike in recent decades, or that the fruits of progress are shared more widely using transfers. If these questions of distribution can be solved satisfactorily, then individuals could indeed respond to productivity growth by working less without experiencing material losses. There is considerable evidence that many workers would prefer to work less and with more flexible work sessions. The Dutch model, which provides all workers with a right to part-time work (at pro-rated wages) could serve as an example, assuming that wages are sufficiently high.^{33,34}

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 for whom the social returns of those expenditures may be high. For instance, expenditures on human well-being, such as on education and health, are naturally targeted to those who need education and healthcare, rather than being spent on those who already are educated or on those who are healthy. Expenditures to protect the environment help those who bear the brunt of environmental degradation, including climate change, which disproportionately affect the poor.³⁵

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

³² Notably, some have discussed a “robot tax” that could help finance redistributive fiscal measures (e.g., Rubin, 2020). However, such a robot tax may be difficult to implement (e.g., what distinguishes a “robot” from traditional capital?) and may discourage innovation (e.g., Summers, 2017). Conceptually, government ownership of capital is equivalent to taxes on capital with exemptions on new investment that avoid any negative incentive effects of capital taxation, although it may be insufficient to provide funding for large-scale redistributive programs that may be needed in a long-term equilibrium with low employment levels. See also Korinek (2020).

³³ The reduced labor supply may itself help sustain higher wages.

³⁴ For a broader view of how to achieve inclusivity in the labor market beyond the challenges posed by technological advances, see El-Ganainy et al. (2021).

³⁵ For example, Colmer et al. (2020) find that while fine-particle air pollution has decreased overall in the US over the past four decades, whiter and richer neighborhoods have become relatively less polluted, while poor and minority communities are (still) the most polluted.

Of particular importance are investments in digital infrastructure that 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 elderly, and some aspects of education, which can again be designed to serve double duty – disproportionately benefiting the poor and needy as they increase wages by increasing the demand for labor.

B. Pre-Distribution

Our concern here is the distribution of consumption (or more broadly, of well-being) among the citizens of a country. That is affected by inequalities in market incomes and the extent of redistribution. The previous subsection discussed redistribution through tax and expenditure policies. But a society with a more equalitarian market distribution needs to place less burden on redistribution. Good policy entails an optimal mix of “pre-distribution”—actions to increase the equality of market income—and redistribution. This is especially so because some of the actions to increase the equality of market distribution are actually efficiency-enhancing, i.e., have a negative cost. For instance, actions which reduce market power, the ability of firms to exploit information asymmetries, or to engage in a variety of other exploitive practices.

There are two categories of policies which affect the distribution of market incomes: (1) Policies that affect individuals endowments of assets—human capital (education) and financial assets. These are affected by the public provision of education and more broadly, policies which affect the intergenerational transmission of advantage and disadvantage (such as inheritance taxes.) And (2) policies that affect the returns on factors, which include the laws and regulations that determine the “rules of the game.” These include competition laws, labor legislation, and rules governing globalization, the financial sector, and corporate governance. These rules affect simultaneously efficiency and distribution.³⁶

Education Regarding the first set of such policies, the fact that more educated workers receive higher incomes than less educated ones may invite the conclusion that education is the solution to inequality. While providing more equal access to high-quality education especially for the poor may reduce inequality—and is absolutely essential to 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—education is far from a panacea. Indeed, if there are large innate differences in ability, education can identify and amplify these differences, actually increasing inequalities

³⁶ For an extensive discussion of some of the critical “rules,” see Stiglitz et al (2015, 2019). Later, we discuss a particularly important set of policies that can affect the returns to factors—those associated with steering the development and adoption of technologies.

in market income. (Stiglitz, 1975b). Moreover, education cannot address the problems arising from the declining share of labor income overall.

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 the direction of technological progress in those countries – how labor-saving it is – also matters for developing countries that will be exposed to the new technologies.

Korinek and Stiglitz (2020) make the case for actively steering technological progress so that it is more labor-using. They show that whenever lump-sum transfers are not available, it is desirable to encourage technological progress that leads to higher demand for those types of workers with the lowest incomes. This can be done by nudging entrepreneurs, by considering the labor market implications of government-sponsored research, or by explicit incentives provided to the private sector. Klinova and Korinek (2021) and Partnership on AI (2021) describe how to develop and how to operationalize frameworks for steering advances in AI towards greater shared prosperity.

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. Examples include self-checkout kiosks that harm workers, whatever their benefits or costs may be for consumers and global corporates.

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 (or from the perspective of global efficiency).³⁷

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, for

³⁷ The alternative – to allow for greater migration – would of course also put downward pressure on wages in advanced countries and might increase inequality within those countries.

example, by the case of northern and southern Italy or by the rural/urban differential in the United States and many other countries. Looking ahead, the “good” jobs of the future (e.g., those hard to automate and/or complementary to AI) might be located in major urban centers (such as non-routine manual occupations and personal services). Such geographical disparities call for location-based policies in fostering development, although the details of such policies can be complex. For example, they may entail a trade-off between income growth (which benefits from geographical concentration) and geographical inequality (which does not).

New Development Strategies

Developing countries will need a new multi-pronged development strategy to replace the manufacturing-led export-based growth model. Industrial policies have traditionally been among the most important aspects of countries’ development strategies—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 the secondary sector to other sectors of the economy, including agriculture and services.³⁸

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 do not know (or admit) that they have such policies. The danger then is that such policies can be more easily captured by special interests.³⁹ In developing countries development policies are much more at the center of economic policy. They need to be designed to manage innovations and mitigate the effects of and adapt to the disruptions that innovations may engender, to ensure that the net societal benefits, broadly defined, are maximized.

A lot of innovation in developing countries focuses on technology adoption and adaptation 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 should 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 should 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: how can government improve upon the decisions made by decentralized agents? Of particular importance is that the direction of

³⁸ 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 term *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).

³⁹ 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.

technological progress and 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 consider the dynamic implications of their decisions today on others.

There are also 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 which decrease welfare—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 do not 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 do not reflect the social value of an equitable distribution of resources and do not guide innovation 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, 2020). 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 weather conditions and need to be customized to local conditions. Just as agricultural extension services, which extended general knowledge about agriculture to local farmers, played a critical role in the development of the US, there is an important role for government agricultural extension services today in developing countries.⁴⁰ Digital platforms can also 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

⁴⁰ In 1914, the U.S. Department of Agriculture created a system of “extension” services, with the aim of providing farmers with expert advice on agriculture and farming. See, e.g., <https://www.almanac.com/cooperative-extension-services>.

market for simple human services that can be broken down into small components and fed into AI systems (e.g., labeling images). However, as we noted earlier, services that can be outsourced are often also 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 elderly, as well as education, may not deliver much export revenue but are 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 relocating to warmer climates, if adequate health care is available.

VI. 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 do not 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.⁴¹ 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

⁴¹ For a discussion of how this plays out in trade rules, see, e.g., Charlton and Stiglitz (2005).

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 (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.

A. 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 activities take away business from 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 are collected in the process of economic transactions and how and whether that value should be taxed.

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.

B. 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. While the policy remedies suggested by the Europeans may have reduced the rents the companies could earn in Europe, their purported aim was 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, there is a need for stronger rules preventing conflicts of interest for companies that 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. There will also be a need for more ex-post remediation: breaking up mergers when they prove to be anti-competitive.⁴² 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.

C. Intellectual Property Rights

The current system of intellectual property (IP) rights is designed to give (temporary) 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 IP regime to ensure an equitable distribution of the gains from technological progress. Korinek and Stiglitz (2019) demonstrated that reducing the length of patent protection can ensure that the gains from AI-based innovations are better shared among society and can thus 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

⁴² There is by now a large literature describing the new competition policies that may be required. See Stiglitz (2019) and Wu (2018), as well as Akcigit et al. (2021) for an overview of emerging issues and complexities in competition policy.

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 expire in high-income countries. Indeed, in many sectors, including pharmaceuticals, there is extensive cross-border price discrimination; drug companies could 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 access such life-saving IP at appropriate royalties, but many developing countries do not have the capacity to exercise those rights; and those that do have the capacity are intimidated from doing so by threats from developed countries. Trade agreements have done everything they can to impede 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 IP 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 that 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 is imperative to ascertain whether algorithms are discriminatory, for example, by engaging in price discrimination.⁴³

D. 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. Global tech firms are setting the data regulatory agenda in their interest without sufficient public oversight. This has already happened 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

⁴³ It is sometimes argued that such disclosure is not possible because algorithms are always evolving. While they are always changing, they could still be disclosed as of a particular moment in time. There are other (often costly) ways of monitoring the behavior of algorithms at any point in time.

countries: the rules are currently being set 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 close the gap between themselves and the leading firms.

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; without proper 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, for example, regarding the algorithms and targeting of advertising, are necessary, but again not sufficient. Policymakers must be able to address the discriminatory impacts of pricing and advertising.

There is also a need for stronger rules protecting privacy and 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. In the US, the Section 230 provision which reduces the accountability of internet companies—unlike other publishers—is an example of a regulation that should be reconsidered.

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 need to cooperate and band together to have sufficient clout to impose regulation on global giants that reflects their developmental interests.

VII. 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 increases in wage inequality, declining labor demand, and increased skill premia in most countries; as well as to the demise of the manufacturing-export-led developmental model, which has historically had 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 could be observed over the last half century.

While earlier technological advances were associated with more shared prosperity and increasing equality between and within countries, the new advances may result in increasing inequality along both dimensions unless policies are designed to counterbalance them.

The new era will be governed by different rules and will require a different kind of economic analysis. Just as the production functions that Ricardo used to analyze agrarian and rural economies are very different from those in the models of a manufacturing economy that dominated the mid-20th century, current economic frameworks must be adjusted and updated 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.

There is a particular high degree of uncertainty across the possible scenarios of technological development and their impact, but what we do know is that there are large potential 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).
- Acemoglu, Daron, and Pascual Restrepo (2019a), “Automation and New Tasks: How Technology Displaces and Reinstates Labor,” *Journal of Economic Perspectives* 33(2), pp. 3–30.
- Acemoglu, Daron, and Pascual Restrepo (2019b), “Demographics and Automation,” Unpublished Working Paper, MIT. <https://economics.mit.edu/files/16788>
- Acemoglu, Daron, and Pascual Restrepo (2020), “Robots and Jobs: Evidence from US Labor Markets,” *Journal of Political Economy* 128(6), pp. 2188–2244.
- 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.
- Aghion, Philippe, Reda Cherif, and Fuad Hasanov (2021), “Competition, Innovation, and Inclusive Growth,” IMF Working Paper No. 2021/080. Also forthcoming as Chapter 6 in Cerra et al. (forthcoming).
- Aghion, Philippe, Benjamin F. Jones & Charles I. Jones (2017), “Artificial Intelligence and Economic Growth,” NBER WP 23928.
- Akcigit, Ufuk, Wenjie Chen, Federico J. Díez, Romain Duval, Philipp Engler, Jiayue Fan, Chiara Maggi, Marina M. Tavares, Daniel Schwarz, Ipei Shibata, and Carolina Villegas-Sánchez (2021), “Rising Corporate Market Power: Emerging Policy Issues,” *IMF Staff Discussion Note* SDN/21/01.
- 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.
- Armstrong, Stuart, Kaj Sotala and Seán S. ÓhÉigeartaigh (2014), “The errors, insights and lessons of famous AI predictions – and what they mean for the future,” *Journal of Experimental & Theoretical Artificial Intelligence*, 26(3), pp. 317-342. <http://www.fhi.ox.ac.uk/wp-content/uploads/FAIC.pdf>
- Arrow, Kenneth (1962a), “Economic Welfare and the Allocation of Resources for Invention,” in: R. Nelson, ed., *The Rate and Direction of Inventive Activity: Economic and Social Factors*, Princeton University Press.
- Arrow, Kenneth (1962b), “The Economic Implications of Learning by Doing,” *The Review of Economic Studies* 29(3), pp. 155-173.

- Atkinson, Anthony, and Joseph Stiglitz (1969), “A New View of Technological Change,” *Economic Journal* 79 (315), pp. 573-578
- Autor, David H., (2014), “Skills, Education, and the Rise of Earnings Inequality Among the ‘Other 99 Percent’,” *Science*, 23 May 2014: 344 (6186), pp. 843-851.
- Autor, David (2015), “Why Are There Still So Many Jobs? The History and Future of Workplace Automation,” *Journal of Economic Perspectives* 29(3), pp. 3–30.
- Autor, David H. (2019), “Work of the Past, Work of the Future,” *AEA Papers and Proceedings* 109, pp. 1-32.
- 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.
- 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.
- Brussevich, Mariya, Era Dabla-Norris, Christine Kamunge, Pooja Karnane, Salma Khalid, and Kalpana Kochhar (2018), “Gender, Technology, and the Future of Work,” IMF Staff Discussion Note, SDN/18/07, October 2018.
- 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.
- Cerra, Valerie, Barry Eichengreen, Asmaa El-Ganainy, and Martin Schindler (Eds.) (forthcoming), *How to Achieve Inclusive Growth*, Oxford University Press and IMF.
- Charlton, Andrew, and Joseph E. Stiglitz (2005), *Fair Trade for All*, Oxford University Press.
- Choi, Eleanor J., and James R. Spletzer (2012), “The declining average size of establishments: evidence and explanations,” BLS Monthly Labor Review, March 2012. <https://www.bls.gov/opub/mlr/2012/03/art4full.pdf>
- 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.

- Colmer, Jonathan, Ian Hardman, Jay Shimshack, and John Voorheis (2020), “Disparities in PM2.5 air pollution in the United States,” *Science*, Vol. 369, Issue 6503, pp. 575-578.
- Comin, Diego, and Bart Hobijn (2010), “An Exploration of Technology Diffusion,” *American Economic Review* 100(5), pp. 2031-59.
- 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.
- Ding, Jeffrey, Paul Triolo, and Samm Sacks (2018), “Chinese Interests Take a Big Seat at the AI Governance Table,” *New America* blog, June 20, 2018.
<https://www.newamerica.org/cybersecurity-initiative/digichina/blog/chinese-interests-take-big-seat-ai-governance-table/>
- El-Ganainy, Asmaa, Ekkehard Ernst, Rossana Merola, Richard Rogerson, and Martin Schindler (2021), “Inclusivity in the Labor Market,” IMF Working Paper No. 2021/141. Also forthcoming as Chapter 3 in Cerra et al. (forthcoming).
- 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.
- Grace, Katja, John Salvatier, Allan Dafoe, Baobao Zhang, and Owain Evans (2018), “When Will AI Exceed Human Performance? Evidence from AI Experts,” *Journal of Artificial Intelligence Research*, <https://arxiv.org/abs/1705.08807v3>.
- Greenwald, Bruce and Joseph Stiglitz (1986) “Externalities in Economies with Imperfect Information and Incomplete Markets,” *Quarterly Journal of Economics*, Vol. 101, No. 2, May, pp. 229-264.
- Greenwald, Bruce and Joseph Stiglitz (2006) “Helping Infant Economies Grow: Foundations of Trade Policies for Developing Countries,” *American Economic Review: AEA Papers and Proceedings* 96(2), pp. 141-146.
- 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.
- Harari, Yuval Noah (2017), *Homo Deus: A Brief History of Tomorrow*, Harper.
- Harris, Karen, Austin Kimson, and Andrew Schwedel (2018), “Labor 2030: The Collision of Demographics, Automation and Inequality,” Bain and Company Reports.
- Kennedy, Charles (1964), “Induced Bias in Innovation and the Theory of Distribution,” *Economic Journal* LXXIV, pp. 541– 47.
- Klinova, Katya and Anton Korinek (2021), “AI and Shared Prosperity,” *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES '21)*.
- Korinek, Anton (2019), “The Rise of Artificially Intelligent Agents,” working paper, University of Virginia.
- 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 (2021), “Preparing for the (Non-Existent?) Future of Work,” in preparation for Justin Bullock et al. (Eds.), *Oxford Handbook of AI Governance*, Oxford University Press.
- 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 (2020), “Steering Technological Progress,” working paper.
- Korinek, Anton and Joseph E. Stiglitz (2021a), “Artificial Intelligence, Globalization, and Strategies for Economic Development,” NBER Working Paper w28453.
- Korinek, Anton and Joseph E. Stiglitz (2021b), “Will COVID-19 Drive Advances in Automation and AI that Exacerbate Economic Inequality?” *BMJ* 372:n367.
- Kurzweil, Ray (2005) *The Singularity is Near*, Penguin Group, pp. 135–136.
- Louie, Edward P., and Joshua M. Pearce (2016), “Retraining investment for U.S. transition from coal to solar photovoltaic employment,” *Energy Economics* 57, pp. 295-302.
- 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.
- McKinsey Global Institute (2017), “Jobs lost, jobs gained: What the future of work will mean for jobs, skills, and wages,” McKinsey Reports.
- McMillan, Margaret S., and Dani Rodrik (2011), “Globalization, Structural Change and Productivity Growth,” NBER Working Paper 17143.
- 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.
- OECD (2019), *The Future of Work*, OECD Employment Outlook 2019.
- Ostry, Jonathan, Prakash Loungani, and Andrew Berg (2019), *Confronting Inequality: How Societies Can Choose Inclusive Growth*, Columbia University Press.
- Partnership on AI (2021), *Redesigning AI for Shared Prosperity: An Agenda*. Technical Report. Available at: partnershiponai.org/shared-prosperity
- Peralta-Alva, Adrian, and Agustin Roitman (2018), “Technology and the Future of Work,” IMF Working Paper WP/18/207. Also published as IMF Staff Note to Group of Twenty: April 2018.
- Peszko, Grzegorz; van der Mensbrugghe, Dominique; Golub, Alexander; Ward, John; Zenghelis, Dimitri; Marijs, Cor; Schopp, Anne; Rogers, John A.; Midgley, Amelia (2020). *Diversification and Cooperation in a Decarbonizing World: Climate Strategies for Fossil Fuel-Dependent Countries*. Washington, DC: World Bank.
- Peterson Institute for International Economics (PIIE) (2020), “How to Fix Economic Inequality?” <https://www.piie.com/sites/default/files/documents/how-to-fix-economic-inequality.pdf>.
- Pouokam, Nathalie (2021), “Sharing Resource Wealth and Addressing Fragility,” IMF Working Paper No. 2021/097. Also forthcoming as Chapter 6 in Cerra et al. (forthcoming).
- 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>
- Rodrik, Dani (2011), “Unconditional Convergence,” NBER Working Paper w17546.
- Romer, Paul M. (1986), “Increasing Returns and Long-Run Growth,” *Journal of Political Economy*, Vol. 94, No. 5 (October), pp. 1002-1037.
- Rubin, Richard (2020), “The ‘Robot Tax’ Debate Heats Up,” *Wall Street Journal*, January 20, 2020. <https://www.wsj.com/articles/the-robot-tax-debate-heats-up-11578495608>.
- Sacks, Samms (2018), “Beijing Wants to Rewrite the Rules of the Internet,” *The Atlantic*, June 18, 2018. <https://amp.theatlantic.com/amp/article/563033/>
- 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.
- Shapiro, Carl and J. E. Stiglitz (1984) “Equilibrium Unemployment as a Worker Discipline Device,” *American Economic Review*, 74(3), June 1984, pp. 433-444.
- Solow, Robert (1987), “We’d Better Watch Out.” *New York Times Book Review*, July 12, pp. 36.
- Solow, Robert (2009), “Does Growth Have a Future? Does Growth Theory Have a Future? Are These Questions Related?,” *History of Political Economy* 41, pp. 27-34

- Stiglitz, Joseph E.(1975a) “Information and Economic Analysis,” *Current Economic Problems*, J.M. Parkin and A.R. Nobay (eds.), Cambridge: Cambridge University Press, pp. 27-52. (Proceedings of the Association of University Teachers of Economics, Manchester, England, April 1974.)
- Stiglitz, Joseph (1975b) “The Theory of Screening, Education and the Distribution of Income,” *American Economic Review*, 65(3), June 1975, pp. 283-300
- Stiglitz, Joseph E.(1987a) “On the Microeconomics of Technical Progress,” *Technology Generation in Latin American Manufacturing Industries*, Jorge M. Katz (ed.), New York: St. Martin’s Press, pp. 56-77. (Presented to IDB-CEPAL Meetings, Buenos Aires, November 1978.)
- Stiglitz, Joseph E.(1987b) “Learning to Learn, Localized Learning and Technological Progress,” *Economic Policy and Technological Performance*, P. Dasgupta and Stoneman (eds.), Cambridge University Press, 1987, pp. 125-153.
- Stiglitz, Joseph E. (2006), “Samuelson and the Factor Bias of Technological Change,” *Samuelsonian Economics and the Twenty-First Century*, M. Szenberg et al, eds., Oxford University Press: New York, pp. 235-251
- Stiglitz, Joseph E. (2013) *The Price of Inequality: How Today’s Divided Society Endangers Our Future*, New York: W.W. Norton & Company.
- Stiglitz, Joseph E. (2014), “Unemployment and Innovation,” *NBER Working Paper* no. 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*, accessible at <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.
- Summers, Lawrence (2017), “Picking on robots won’t deal with job destruction,” *The Washington Post Opinion*, March 5, 2017.
- 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.

- 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. <https://edoras.sdsu.edu/~vinge/misc/singularity.html>
- Walsh, Toby (2016), "The Singularity May Never Be Near," <https://arxiv.org/abs/1602.06462v1>.
- World Bank (1998), *World Development Report 1998/1999: Knowledge for Development*, Oxford University Press. <https://openknowledge.worldbank.org/handle/10986/5981>
- 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>
- World Economic Outlook (WEO) (2018), "Manufacturing Jobs: Implications for Productivity," Chapter 3, April 2018.
- Wu, Tim (2018), *The Curse of Bigness: Antitrust in the New Gilded Age*, Columbia Global Reports.