Simulating Endogenous Global Automation

Seth G. Benzell[*] Laurence J. Kotlikoff[†] Guillermo Lagarda[‡] Victor Yifan Ye[§]

September 4, 2023

Abstract

We develop a dynamic, multi-region, multi-skill-group OLG model of global automation. Regions endogenously choose whether to produce with higher capital- and high-skill-labor intensive technologies. Automation and its induced capital flows significantly raise certain regions' GDPs while lowering those of others. Thus, Western Europe's 2050 output is 12.7 percent higher; Mexico's is 3.6 percent lower. Automation also greatly exacerbates intra-regional wage inequality. These results assume automation advances at recent historical rates. A faster pace, augured by recent AI advances, would deliver larger impacts. The rise in wage inequality, though, will, we show, be limited if automation engenders higher relative supply of high-skilled workers.

We thank the Alfred P. Sloan Foundation, the Gaidar Institute, Boston University, the Stanford HAI Digital Economy Lab, and the Goodman Institute for research support. We also thank Daniel Rock and Sarah Bana for very helpful comments and providing key data. Andrey Zubarev, Kristina Nesterova, and Rodrigo Razo provided analytical assistance. And we thank Ulrich Müller, James Stock, and Mark Watson for providing critical region-specific productivity-growth estimates.

[*]Chapman University, MIT Initiative on the Digital Economy, and Stanford Digital Economy Lab
[†]Boston University, Fiscal Analysis Center, and National Bureau of Economic Research
[‡]Boston University Global Development Policy Center and Inter-American Development Bank
[§]Opendoor Technologies, Boston University, and Stanford Digital Economy Lab

1 Introduction

Advancements in robotics, machine learning, and high-performance computing are revolutionizing production technologies in virtually all industries and altering the distribution of economic rewards, both within and across countries. Cross-country studies record a decline in labor's income share as well a rise in wage inequality associated with reductions in middle-class jobs (Autor, 2015). Karabarbounis and Neiman (2014) report a downward trend in labor's share in over 70 percent of countries with at least 15 years of data. As for the U.S., its labor share fell by 4.3 percentage points between 1980 and 2017. During the same interval, the share of wages of top-quintile U.S. workers rose from 48.5 to 59.1 percent (Census, 2019). Acemoglu and Restrepo (2020) and Acemoglu and Restrepo (2019, 2018b) connect these trends to automation and skill-biased technological change.

This paper develops a large-scale, multi-region, computable general equilibrium model of endogenous global automation to study its impact through time on the distributions of output across regions and of wages within regions. Our model can examine different scenarios concerning technology's spread, adoption, and induced factor-supply responses. It offers a framework for addressing a worst-case automation scenario that would greatly undermine the economic futures of low-income regions and low-skilled workers regardless of region.

Our model features 101 overlapping generations located in 17 regions. These regions are aggregations of over 150 countries, accounting for 99 percent of the world's population and 98 percent of its output. Agents comprise three skill groups – high, middle, and low – that are more or less abundant depending on the region. The model has 464 demographic, fiscal, preference, and technology parameters. We develop a novel technique to simultaneously calibrate these parameters using UN, IMF, and other macro data in conjunction with micro findings from Acemoglu et al. (2020), Altig et al. (2020), and Auerbach et al. (2016).

Technological change evolves, in part, exogenously and, in part, endogenously. Endogenous references the decision of whether or not to adopt the latest available, region-specific technology. Exogenous refers to the pace of U.S. factor-share innovation, which we extrapolated forward based on recent data. Foreign regions' production coefficients change, we assume, at the U.S. percentage pace. Yet regions are free to operate with older production functions – those with lower capital and high-skilled worker share coefficients – if doing so, as least short term, is more profitable.

The multi-region nature of our global model has critical implications for the consequences of technological change. We show that access to global capital significantly improves both economic and aggregate welfare gains from automating technologies. Absent global investments, most developed regions will never adopt share-shifting technologies. Similarly, our choice of utilizing an OLG model allows us to evaluate the intergenerational welfare impact of automating technologies. Indeed, we find that such technologies generally and disproportionately benefit a small subset of young, high-skilled workers in developed countries.

In our baseline simulation, the U.S. and other economically advanced regions choose to adopt frontier technologies as they emerge. Other regions adopt at a slower pace or not at all. Many factors, including the panoply of region-specific fiscal policies, demographics, and technological parameters, contribute to this decision by altering factor-price paths. In line with Zeira's (1998a) seminal contribution, regions with higher initial capital shares are more likely to automate, increasing cross-country differences in living standards. Since factor prices are determined in the global economy's dynamic general equilibrium, a decision by one region to automate will impact automation decisions in all others.

We calibrate initial factor-input and population skill-group labor-force shares based on IMF (2019a) data on the global distribution of capital as well as data from Alvaredo et al. (2020) on country-specific income inequality. Our calibration evidences important technological and factor-supply differences across regions, including initial levels of capital intensivity and the relative prevalence of high-skilled workers. We also calibrate region-specific total factor productivity (TFP) growth rates to long-term growth rates projected by Müller et al. (2019).

In our baseline, which projects innovations through 2050 consistent with post-1980 changes in U.S. input shares, the U.S. total labor share falls from 65.9 to 63.1 percent. The U.S. capital share increases accordingly. The Chinese total labor share falls from 61.7 to 59.1 percent. The U.S. high-skilled share rises from 22.0 to 26.4 percent. For China, the increase is from 20.6 to 24.8 percent. Middle-skilled labor shares remain roughly fixed in both regions. In contrast, the U.S. low-skilled labor share falls from 22.0 to 16.1 percent and the Chinese low-skilled labor share falls from 20.6

to 15.1 percent.

Automation has the largest impacts on Japanese, Western European, Canadian, UK, and U.S. experience output levels. In 2050, the regions' respective GDPs are 12.9, 12.7, 12.1, 10.0, and 5.1 percent higher than they would have been absent automation. China defers automating until 2027. Still, it's 2050 GDP is 3.6 percent larger than in the no-automation case. Automation makes slow- or non-automating regions, including India, Brazil, Mexico, and Sub-Saharan Africa, worse off. This reflects their loss of foreign investment. Take Mexico. It produces 3.6 percent less output in 2050 due to a 9.2 percent decline in the amount of capital that can profitably operate in that region. The impact on wage inequality is striking. Wages of high-skilled U.S. workers, absent automation, are projected to be 18.9 times those of the low-skilled by 2050. With automation, they are 31.5 times higher.

To better understand the interaction between capital mobility and automation, we simulate the model with the U.S. operating as a closed economy. The U.S. still automates under this scenario, but experiences a much smaller, indeed, near-zero change in output. Clearly, the benefit of a using a more capital-intensive technology increases when more capital is available. By forgoing trade, the U.S. forgoes substantial foreign investment, particularly from Asia and the Middle East. This counterfactual demonstrates the need to study automation in general equilibrium with capital mobility.

Suppose, as another counterfactual, that only the U.S. automates. As expected, the U.S. gains even more. Relative to no automation, U.S. GDP in 2050 increases by 5.2 percent in our baseline scenario. This figure is 8.3 percent if only the U.S. automates, but less than 0.1 percent if the U.S. can't import foreign capital. If automation were to proceed at twice the historical rate, U.S. GDP would be 11.9 percent higher in 2050. In this case, 2050 middle- and high-skilled wages are 1 percent and 31.1 percent higher, respectively, compared with a 1.7 percent decrease and 27.6 percent increase in our baseline. Low-skilled wages are 20.8 percent lower, compared with 22.9 percent in the baseline.

Finally, we study the transition assuming an increase in the relative supplies of middle- and high-skilled workers. We consider three cases. In these scenarios, the population share of high-skilled U.S. workers rises from our baseline calibration of 4 percent to between 4.6 and 5.5 percent by 2050. The mid-skilled population share increases from 20 to 23.5 to 26.5 percentage points. Under each scenario, automation raises U.S. output, increasing it by 15.9 to 29.0 percent by 2050. But the most interesting impact is on wage inequality, which is significantly reduced due to the supply-side response.

Across all scenarios, and as suggested by stylized models (e.g. Benzell et al. 2016; Sachs et al. 2015a; Acemoglu and Restrepo 2018a; Peretto and Seater 2013), automation increases the world interest rate. However, this effect is relatively modest. Absent automation, the global interest rate falls from 6.0 percent in 2017 to 2.8 in 2050. In our baseline scenario, the 2050 interest rate decreases to 3.2 percent. Two trends contribute to the global savings glut (Eichengreen, 2015). The first is population aging in both developed and emerging regions, which raises global wealth per worker. The second is the high saving rate of Chinese and Indian households coupled with their rising share of world GDP. Taken together, these factors keep capital abundant, raising the returns to automation. What if automation proceeds more rapidly? Could it materially alter macroeconomic outcomes? Yes, but. Automation would need to proceed at more than five times the historical rate to prevent interest rates from falling.

2 Background

Concern with technological change is longstanding. The historical record stretches from Prometheus, who was bound and tortured for giving ancient Greeks fire, to the followers of 'King Ludd,' who were hanged for smashing English textile machines, to Marx (1867), who ardently proclaimed technology's economic threat to the proletariat. To quote Marx,

> *"Within the capitalist system all methods for raising the social productivity of labour are put into effect at the cost of the individual worker; all means for the development of production undergo a dialectical inversion so that they become means of domination and exploitation of the producers; ...they alienate from him the intellectual potentialities of the labour process in the same proportion as science is incorporated in it as an independent power..."*

Marx's views resonated for a reason. As documented by (Katz and Margo, 2014), early industrialization redistributed to the elite at the expense of middle-skilled, craft-workers. Keynes (1930) also raised concerns with automation, although he believed its cost to workers was transitory.

> *"We are being afflicted with a new disease of which some readers may not yet have heard the name, but of which they will hear a great deal in the years to come – namely, technological unemployment. This means unemployment due to our discovery of means of economizing the use of labor outrunning the pace at which we can find new uses for labor."*

Today, the world faces new concerns about technological breakthroughs, like Chatgpt, displacing labor across a wide range of tasks. McAfee and Brynjolfsson (2017) note how advances in machine learning are overcoming Polanyi's Paradox, an alleged critical barrier to automation. Polanyi observed that most human behaviors, such as driving a car or recognizing a face, harness tacit or intuitive knowledge which cannot be expressed mathematically, making it impossible for an expert system to replicate. However, neural networks are overcoming this barrier by training machines how to act in ways that don't require explicit algorithms.

The advance of machine vision technologies is particularly impressive. In the well known 'ImageNet' image-classification contest, the 2012 top performer had a 15 percent error rate. In 2017, the top twenty contestants had less than a 5 percent error rate (Wikipedia contributors, 2021). Estimates of how widely new technologies may impact employment are daunting. Frey and Osborne (2013), in a controversial article, estimate that 47 percent of the U.S. work force is at risk of automation. Recent estimates find that almost all occupations include some tasks conducive to automation (Brynjolfsson et al., 2018). Large language models, such as the one powering chatGPT, have the potential to augment or automate an even larger share of tasks (Eloundou et al., 2023).

These prognostications must be taken seriously given the dramatic shifts in factor shares in recent decades documented by Karabarbounis and Neiman (2014). There is strong evidence that the decline in the labor share in developed countries is connected to technological change. Autor et al. (2020) claims that innovations, especially in information and communication technologies, are making the U.S. economy more 'winner take all' and thereby reducing labor's share. Acemoglu and

Restrepo (2020) and Acemoglu et al. (2020) show that robot adoptions by U.S. counties and French firms also led to decreases in labor's income share. Goos et al. (2011) offer supporting evidence for Europe.

2.1 Stylized Models of Automation

Our study builds on a number of stylized macro models of technical change. Our model exhibits many of the effects posited in this theoretical literature, but allows for their relative importance to be quantified, showing some to be of more practical importance than others.

One of the most important papers in this literature is Zeira (1998b). That paper considers the introduction of labor-substituting technology, showing that countries with high labor costs and low interest rates will adopt such technology, i.e., automate, more rapidly than others. Countries that adopt grow faster, making automation a potential source of international growth disparities. This force for international inequality plays an important role in our model's baseline scenario.

Zuleta (2008), Peretto and Seater (2013), and Acemoglu and Restrepo (2018c) build on Zeira (1998b). Zuleta (2008) evaluates an economy's decision to invest in automation technology rather than traditional capital. He shows that countries will use either the most or least capital-intensive technology available – never an interior option. Zuleta (2008) further demonstrates that, in the absence of productivity growth, automation is necessary to maintain long-term growth in per capita output. Economies in his model either converge to AK growth, in which output grows at a constant rate or stop growing on a per capita basis. The reason automation enables long-term per capita growth is that capital per-capita is an elastic factor that can be increased, but labor per-capita is limited in the absence of improvements in workers skills (which plausibly cannot continue indefinitely). The larger labor's input share in the production function, the more this asymptotically scarce factor of production lowers asymptotic output.

Peretto and Seater (2013) more fully develops the micro foundations of Zuleta (2008), showing conditions under which firms will invest in automation technologies under imperfect competition. Acemoglu and Restrepo (2018a) endogenize the choice of technology in which to invest. In their setup, scientists endogenously decide whether to automate existing labor tasks or invent new ones for humans to perform. This decision depends on prevailing factor prices. When wages are high and interest rates are low, it makes more sense to automate. When the opposite holds, scientists invent new jobs to take advantage of the relatively abundant labor. Thus, technological change has a self-correcting mechanism preventing complete automation. However, depending on parameters, the response to automation-inducing shocks can take a very long time to raise wages.

An important point of these papers is that, when profitable, firms will invest to reduce their dependence on labor. Automation is profitable when capital is abundant and cheap and when labor is scarce and expensive. The result is generally not less aggregate employment, as these models assume that workers are willing to work for lower pay. There is, though, a rise in capital's income share. This motivates our choice to model automation as the adoption of technologies that are more capital-intensive and, in our baseline scenario, more high-skilled labor intensive.

Another critical factor in all of these papers is that the interest rate determines whether and to what extent a region gains from automation. The interest rate is, in turn, determined by the rate of saving. In Zeira (1998b) the rate of saving is constant, and in Acemoglu and Restrepo (2018a) the rate of saving is pinned down by a representative agent's Euler condition. Both of these saving rules tend to guarantee automation has positive social effects, because they keep the rate of investment constant, or even increase it, after an increase in automation technology.

Our life cycle model's saving behavior is different. Workers save to finance their retirements whereas retirees dissave, i.e., spend down their accumulated assets. Hence, the path of national saving is a complex function of the path of factor incomes and demographics. Automation may, therefore, lower the total saving rate and have an offsetting negative effect. Benzell et al. (2015b) and Sachs et al. (2015b) illustrate the important interaction of saving and automation. In their models, the adoption of new automation technologies redistributes from young savers to old spenders. This can reduce the economy's long-run supply of capital and, potentially, lower long-run output and welfare. In other words, automation can be immiserating.

In our calibrated model, we find this theorized negative effect of automation to be relatively unimportant. While automation does lower wages – which Benzell and Brynjolfsson (2019) shows to be a general, short-run prediction of a broad class of automation models – this effect is not large enough to dramatically redistribute between young savers and old consumers. Indeed automation in our baseline scenario leads to higher levels of global capital accumulation over time.

Our research is also related to the theoretical analyses of Korinek and Stiglitz (2019) and Nordhaus (2015). The former paper analyzes when technological improvements will fail to Pareto-improve. If all skill-groups and generations could participate in an Arrow-Debreu market behind the veil of ignorance, growth-inducing technological change would pose no downside. However, in the real world, markets may not sufficiently insure against technological risk to routine workers and future generations. In our calibrated model of the world's demographic transition, we find large inter-generational redistributions caused by technological change. Nordhaus (2015) evaluates a number of tests for whether the 'technological singularity' is near, ultimately rejecting this hypothesis.

2.2 Previous Computable General Equilibrium Models

Our model is, to our knowledge, the first large-scale, multi-region, global OLG model of automation. Dawid (2006) surveys the use, through 2006, of CGE models for understanding technological change, emphasizing the deployment of such models to understand microeconomic and industry level dynamics in R&D, innovation, spillovers, and firm-size distribution, none of which are directly relate to our study. In the years since, there have been a handful of papers simulating technological change, but these have been narrow applications to specific industries such as agriculture (Dietrich et al., 2014) and residential heating (Knobloch et al., 2021).

Other general equilibrium models of skill-biased technological change or the geography of automation (e.g. Krusell et al. 2000, Jaimovich et al. 2021, Deen Islam 2021) are useful, but do not feature detailed demographic or fiscal calibrations. Our model descends from Auerbach and Kotlikoff (1981), Auerbach and Kotlikoff (1983), Judd (1985), Auerbach and Kotlikoff (1987), Fullerton and Rogers (1996), Altig et al. (2001), Auerbach et al. (1989), Fehr et al. (2003), Fehr et al. (2013a), Fehr et al. (2013b), and Benzell et al. (2015a). This is a long literature that expands the original Auerbach-Kotlikoff model to include demographic change, multiple skill groups, multiple goods, multiple regions, an energy sector, and other key features. Our model's closest antecedent is the

Global Gaidar Model (GGM), presented in Benzell et al. (2020) and Benzell and Lagarda (2017), which simulates a variety of business cash-flow tax reform policies. Similar versions of this model have also been used to investigate the impact of fossil fuel rents on fiscal transitions Benzell et al. (2015a), to simulate long-term U.S. fiscal sustainability (Nelson and Phillips, 2019), and project the future of global economic power Benzell et al. (2022).

Compared to Benzell et al. (2020), we make four advances. First, we develop a novel calibration algorithm that accommodates a major increase in the complexity of CGE models. Instead of manually adjusting the model's parameters to match target data, the cost of which grows exponentially with the number of parameters, our approach, described in section 3.3.2, permit rapidly calibrating very large numbers of parameters. Absent this feature, our model would be too complex to calibrate.

Second, we calibrate region-specific production technologies, modeled as Cobb-Douglas functions with region-specific coefficients. We utilize region-specific capital stock data from the IMF Investment and Capital Stock Database (ICSD; IMF 2019a) and adjust the capital-use intensivity of each region's representative firm to match the ratio of capital to GDP in each region. This captures the fact that countries have different mixes of industries and technologies and, therefore, should not be modeled with the same production function.

Third, we use recent U.S. microeconomic data in Auerbach et al. (2016) and Altig et al. (2020) to calibrate our model's initial U.S. age-skill asset distribution, the progressivity of the U.S. tax system, and the skill distribution of U.S. Social Security benefits. We assume the model's other regions have the same initial age/skill asset distribution as the U.S., the U.S. degree of tax progressivity, and the U.S. skill distribution of state pensions.[1]

Fourth, we use micro estimates in Acemoglu et al. (2020) and other assumptions to calibrate factor-specific productivities. These productivities, in turn, tell us how output changes in response to changes in factor shares given factor prices. Determining these productivities is critical to ensure that the analysis of technological change arising from changes in factor shares is invariant to the

[1] While we would ideally calibrate to specific regions' age-asset distribution, data on this are not generally available. The initial age-asset distribution is important for understanding short-term welfare implications of shocks, but relatively unimportant for projecting long-term dynamics.

scaling of inputs and outputs. Controlling for scaling is not important in standard analyses of technological change, which involve only changes in TFP.

3 Model

Our model has seventeen regions, each with three distinct sectors: a household sector whose agents work, save, and consume, a government, which taxes, transfers, and spends, and a business sector, which hires labor, rents capital and produces a single output that can be consumed or invested. Regions also receive fossil fuel rents that accrue to both private investors and governments. The seventeen regions are listed in table 1.[2]

Acronym	Region	Acronym	Region
USA	U.S.	MENA	Middle East and North Africa
WEU	Western Europe	MEX	Mexico
JKSH	Japan, South Korea, Singapore and Hong Kong	SAF	South Africa
CHI	China	SAP	South Asia Pacific excluding Australia
IND	India	SLA	Latin America excluding Mexico and Brazil
RUS	Russian Federation	SOV	Former Soviet Central Asia
BRA	Brazil	SSA	Sub-Saharan Africa excluding South Africa
GBR	The U.K.	EEU	Eastern Europe
CAN	Canada, Australia and New Zealand		

Table 1: Regions in the model and their acronyms.

3.1 The Demographic Transition

The model has 101 overlapping generations of agents of ages 0 to 100. Children are born with the same skill as their parents and retain that skill for life. There are three skill levels – high, middle, and low. Children receive consumption from their guardians (either parents or grandparents) before entering the workforce. At a region- and skill-group k specific age, a_{w_k} between 18 and 21, agents enter the workforce and receive labor income until a region-specific retirement age, after which

[2] A list of countries in each region is provided in table A1.

they receive pensions.[3] Agents give birth to fractional children between age 15 and 49. Children of agents who are too young to work (those from 15 to $a_{w_k} - 1$) are raised by their grandparents. Henceforth, we use the term "guardians" to refer to the parents and grandparents responsible for a given cohort of children and "dependents" to refer to their children and grandchildren. Agents can die at any age through age 100, after which they die with certainty. Figure 1 summarizes the agent's life-cycle.

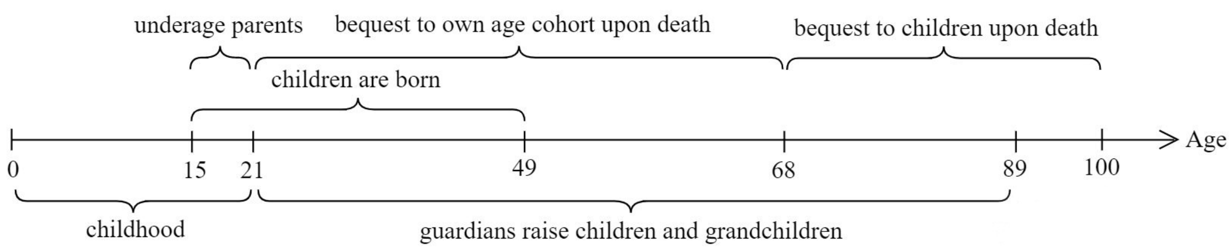

Figure 1: The individual life-cycle

The paths of demographics in each region, including year- and age-specific births and deaths are set to match United Nations (UN) projections through 2100. The UN projects births, deaths, and total population by country in every fifth year for five-year age groups. We aggregate these country-specific projections to generate region-specific projections. Next, we smooth the region-specific projections to generate births, deaths, and total population projections of agents in each region at each age in each future year. We calculate annual net immigration by age and year to reconcile the smoothed age- and year-specific UN births and deaths to the smoothed age- and year-specific UN population totals. These age- and year-specific migrant "residuals" can be positive or negative.

There is striking disparity across regions in population trends. As shown in figure 2, China's population is projected to fall from 1.4 billion now to approximately 1 billion by 2100 and, as birth rates stabilize, to 450 million by 2250. Much of the developed world – Europe, Japan and Korea – experience major population declines through century's end. Other developed regions age

[3] We set the age of entry into the labor force to roughly match the fraction of individuals in a region that receive some college education. Low- and mid-skilled agents of developing and least-developed regions enter the labor force at 18. For the U.S. and other developed regions, all three skill groups enter the labor force at 21. See section 3.2.3 for a detailed discussion of treatment of pension benefits.

Figure 2: UN and Model Population Projections

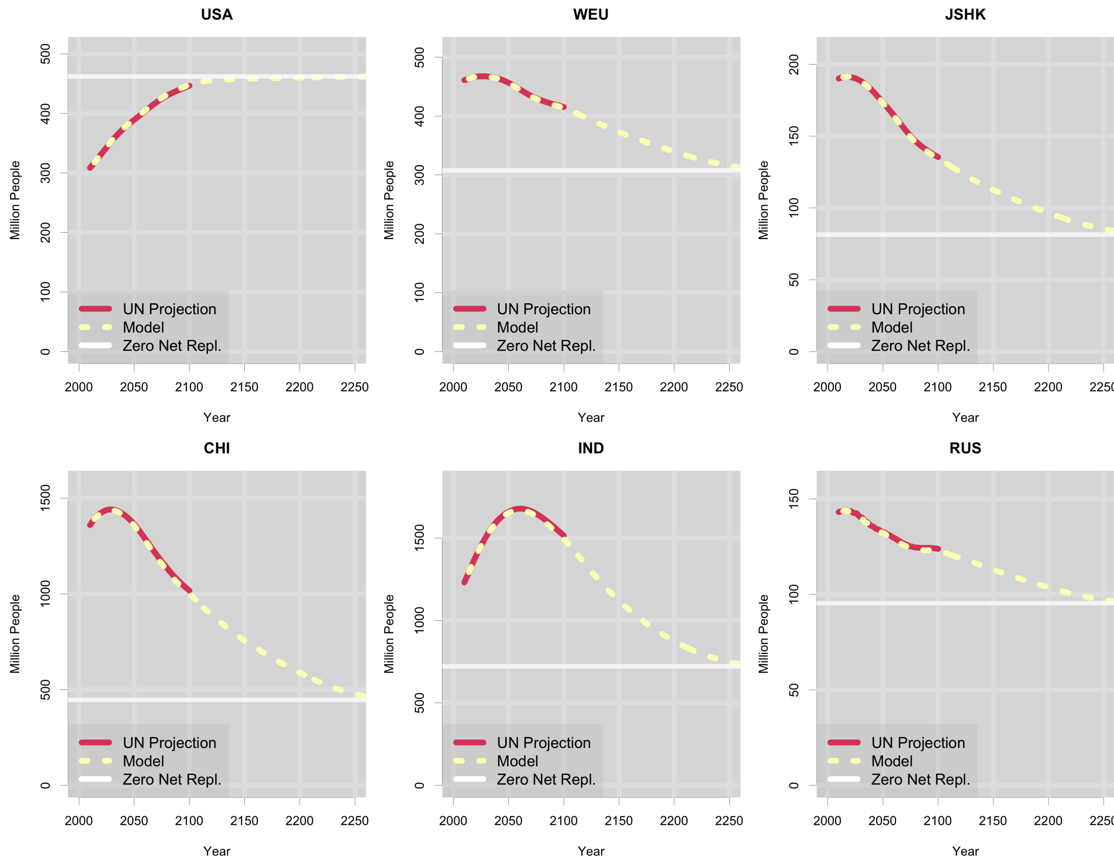

significantly. A total of 21.2 percent of Americans and 23.6 percent of Britons will be older than 70 by 2100, up from 9.0 and 11.6 percent in 2017, respectively. Sub-Saharan Africa and the Middle East, including North Africa, grow by nearly 3 billion people by 2100. This is due to persistently high fertility rates in the region – 3.1 births per women in 2050, for example – combined with rapidly declining mortality. Yet these regions also experience significant population again over the century.

These demographic dynamics have major implications for the adoption of technologies that substitute away from low- and mid-skilled labor. As the world ages, the pool of labor shrinks, tending to raise wages. Further, retirees have more assets relative to young workers. Hence, the global aging lowers interest rates. The changes in factor prices favor automation, particularly in developed regions with high wage rates and aging populations.

3.2 Sectors

3.2.1 Households

Households are composed of agents who decide how much to work, save, and consume each year. Agents receive a fixed time endowment, h, in each period and maximize lifetime utility given by a nested, time-separable CES utility function. For a working agent of age a in time t and skill-group k, dropping regional subscripts, we have

$$U_{a,t,k} = V_{a,t,k} + H_{a,t,k}, \tag{1}$$

where $U_{a,t,k}$ is the agent's utility, the sum of his or her utility $V_{a,t,k}$ derived from consumption and leisure, and $H_{a,t,k}$, is the utility derived from the consumption of their dependents, including both children and grandchildren of underage parents. The function $V_{a,t,k}$ is given by:

$$V_{a,t,k} = \frac{1}{1-\frac{1}{\gamma}} \sum_{i=a}^{100} \left(\frac{1}{1+\delta_t}\right)^{i-a} P_{a,i,t} \left[c(i,t+i-a,k)^{1-\frac{1}{\rho}} + \varepsilon_{t,k}\ell(i,t+i-a,k)^{1-\frac{1}{\rho}}\right]^{\frac{1-\frac{1}{\gamma}}{1-\frac{1}{\rho}}}. \tag{2}$$

This equation represents discounted CES utility over consumption and leisure over the agent's remaining lifetime. $P_{a,i,t}$ is the probability that an agent of age a at time t survives to age i, which is not specific to skill group. $c(i,t+i-a,k)$ is the consumption of an agent of age-a, skill-group k, at time $t+i-a$, the period when they reach age i, and $\ell(i,t+i-a,k)$ is leisure for this group in the same period. The function $H_{a,t,k}$ is given by:

$$H_{a,t,k} \quad = \quad \frac{1}{1-\frac{1}{\gamma}} \sum_{i=a}^{100} \left(\frac{1}{1+\delta_t}\right)^{i-a} P_{a,i,t} D_{a,i,t,k} c_{D_{a,i,t,k}}^{1-\frac{1}{\gamma}}, \tag{3}$$

which represents discounted utility over consumption $c_{D_{a,i,t,k}}$ of all of the agent's dependents $D_{a,i,t,k}$. δ_t denotes the time preference rate, which is assumed to be constant across age and skill but varies across regions and with time. ρ and γ are, respectively, the intratemporal elasticity of substitution

(ELS) between consumption and leisure and the intertemporal ELS of utility over time. Both are assumed to be fixed across age-skill groups, regions, and time. $\varepsilon_{t,k}$ is the time- and skill-group specific subjective leisure preference parameter.

Assets of surviving agents evolve according to

$$A_{a+1,t+1,k} = (A_{a,t,k} + I_{a,t,k})R_{t+1} + (h - \ell_{a,t,k})w_{a,t,k} - T_{a,t,k} - C_{a,t,k}. \tag{4}$$

Agents earn a pre-tax rate of return, R_{t+1}, on their time-t assets, $A_{a,t,k}$, and on total inheritances received by age-a individuals $I_{a,t,k}$. $h - \ell_{a,t,k}$ represents the fraction of the fixed time endowment, h, spent on labor, and $w_{a,t,k}$ is the pre-tax wage of the sub-scripted agent. $T_{a,t,k}$ references an agent's personal taxes net of pensions and other transfer payments received. And $C_{a,t,k} = c_{a,t,k} + D_{a,t,k} \cdot c_{D_{a,t,k}}$.

The model's bequests arise due to an assumed lack of annuity markets; i.e., they are not intentional. Agents give birth and leave bequests exclusively to agents in their own skill cohort. We assume that agents below age 68 bequeath to their spouses, represented as other agents in their own age cohort. Agents at or above 68 bequeath proportionally to all of their adult children.[4] Note also that when grandparents raise children of their "underage" children, bequests always go to the grandparents' own adult children.

3.2.2 Production

Regional gross domestic product is the sum of output from a production sector, Q_t, and a natural resource endowment flow, X_t

$$Y_t = Q_t + X_t. \tag{5}$$

Energy endowments are modeled as an annual flow of the consumption good net of extraction and transportation costs. They accrue, in part, to households as a private asset and, in part, to

[4]For example, an agent dying at 68 bequeaths to his or her age and skill cohort's distribution of working children, all of whom are between ages a_{w_k} and 53.

the government as revenue. We calibrate the size of the flow in each region, the fraction of which accrues to the public sector, and the date of exhaustion based on World Bank (2021) fossil-fuel and IMF (2017) fiscal data.

We define regional gross national income, N_t, as the sum of GDP and net foreign asset income $r_t F_t$,

$$N_t = Y_t + r_t F_t, \tag{6}$$

where r_t is the global rate of return and F_t is net foreign assets.

Production is governed by

$$Q_t = \phi_t (\varphi_K K_t)^{\alpha_t} L_{t,1}^{\beta_{t,1}} L_{t,2}^{\beta_{t,2}} L_{t,3}^{\beta_{t,3}}, \tag{7}$$

where $L_{t,k} = \varphi_{l_k} l_{t,k}$, $k \in (1, 2, 3)$ are productivity-adjusted labor inputs, $l_{t,k}$ is labor supply of group k, φ_{l_k} is a skill-group specific labor productivity term, and φ_K is capital productivity.[5] Both K_t and output are denominated in real 2017 dollars. The term ϕ_t captures time-varying total factor productivity, the calibration of which is detailed in Appendix section A.

In our no-automation calibration, each skill group earns an equal share of total labor income. I.e.,

$$\beta_{t,k} = \frac{1 - \alpha_t}{3}, \tag{8}$$

for $k \in (1, 2, 3)$. Total labor supply, measured in efficiency units, of skill group k at time t, $l_{t,k}$ satisfies

$$l_{t,k} = \sum_{a=a_{w_k}}^{100} (h - \ell_{a,t,k}) E_{a,t,k} \cdot N_{a,t,k}, \tag{9}$$

[5] we use $k = 1$ to represent low-skilled workers, $k = 2$ for the mid-skilled, and $k = 3$ for the high-skilled.

where $N_{a,t,k}$ is the population of skill-group k at age a and time t. $E_{a,t,k}$ is an age- and time-specific productivity profile. This profile is that in Auerbach and Kotlikoff (1987), but modified to allow for different ages of labor force entry. To be precise,

$$E_{a,t,k} = e^{4.47+0.033(a-a_{w_k}-1)-0.00067(a-a_{w_k}-1)^2}(1+\lambda)^{a-a_{w_k}}. \tag{10}$$

We assume that $\lambda = 0.01$, i.e. the time-endowment, rises at 1 percent per year. This assumption that cohorts' effective time endowments grow at a fixed rate through time ensures the global economy will reach a steady state given our CES preferences.

Our perfectly competitive firms hire factors to maximize profits. Thus,

$$w_{t,k} = \frac{\beta_{t,k}Q_t}{l_{t,k}}, \; k \in (1,2,3) \tag{11}$$

$$r_t = (1-\tau_t^K)\left(\frac{\alpha_t Q_t}{K_t} - \delta_K\right), \tag{12}$$

where $w_{t,k}$ are wages per efficient unit of labor for the respective skill groups, τ_t^K is the marginal effective tax rate (METR) on capital, and δ_K is a region-specific depreciation rate.

3.2.3 Government

Each region's government (which includes, in our calibration, super-national, national, and sub-national governments) collects taxes, borrows, and spends on both consumption and transfer payments. Spending includes direct government purchases of goods and services, C_t^g, a fraction, ϱ, of state-pension payments, e_t, which are not covered by payroll taxes, a variety of transfer payments to the private sector, f_t, and interest payments, $r_t B_t$, on debt, B_t.

Government receipts comprise the government's share of energy sector rents, X_t^g, net borrowing, ΔB_t, corporate income taxes, T_t^K, and personal taxes collected from each cohort within a given year, $T_{a,t,k}$. The government's flow budget is

$$\sum_{k=1}^{3}\sum_{a=a_{w_k}}^{100} T_{a,t,k}N_{a,t,k} + T_t^k + X_t^g + \Delta B_t = C_t^g + \varrho e_t + f_t + r_t B_t. \tag{13}$$

Households face taxes on capital, wage, and total income, as well as consumption. Firms face corporate income taxes. Capital income, corporate income and payroll tax rates are fixed, with the payroll tax financing a fixed, region-specific fraction of state pensions. In our calibration, governments endogenously set income and consumption tax rates to maintain a constant level of debt-to-GDP, with the level of income tax progressivity fixed through time.[6] Time-t income tax revenue, R_t, is given by

$$R_t = \sum_{k=1}^{3}\sum_{a=a_{w_k}}^{100}\left(\tau_t w_{a,t,k} + \frac{\xi_t w_{a,t,k}^2}{2}\right) l_{a,t,k}, \tag{14}$$

where τ_t endogenously adjusts to keep regional debt-to-GDP fixed. ξ_t is a progressivity term calibrated to the ratio between the U.S. METR of the top four percent of households by labor earnings and the U.S. average income tax rate of working households. $w_{a,t,k} = w_{t,k} \cdot E(a,t,k)$ is the age-skill group's wage rate. Corporate income taxes are assessed on all domestic output net of labor costs and depreciation.

$$T_t^K = \tau_t^K\left(Y_t - \sum_{k=1}^{3}\sum_{a=a_{w_k}}^{100} w_{a,t,k} l_{a,t,k} - \delta_K K_t\right), \tag{15}$$

where δ_K is the depreciation rate.

In each region, the payroll tax finances a fraction of pension payouts, and the remainder, ϱ, is financed through the general account.[7] Payroll taxes are assumed to be proportional up to a region-specific contribution ceiling, after which the marginal payroll tax rate is zero.[8]

Turning to government expenditures, education and healthcare are treated as a combination of direct government consumption (public schools, government purchases of medical infrastructure

[6]We set the rates endogenously such that the ratio of revenues raised by the two taxes is held constant at their initial regional values. This ratio is calibrated to fiscal data.

[7]In some regions $\varrho < 0$, in which case the payroll tax subsidizes other government spending.

[8]As discussed in section 3.3, our calibration implies that the highest skill group in all regions faces a marginal payroll tax rate of zero. This precludes kinks in budget frontiers.

and supplies, etc.) and transfer payments to households. Both categories are calibrated separately.

Program-p transfers are given by

$$Tr_{p,t} = \zeta_{p,t} Q_t v_p \sum_{k=1}^{3} \sum_{a=a_{w_k}}^{100} Z_{p,a} N_{a,t,k}, \tag{16}$$

where $Z_{p,a}$ is the region- and spending-program specific age-expenditure profile, and $\zeta_{p,t}$ is a region- and program-specific shift term. This term is calibrated to match program expenditure as a share of GDP and assumed to evolve through time to keep transfers constant as a ratio of per capita GNI. v_p is the share of government expenditure on p that is treated as a transfer. General welfare transfers include disability insurance and similar programs and are evenly distributed across adults. In contrast, education and healthcare transfers disproportionately benefit the young and old, respectively.

Pensions and general welfare transfers are fungible. Healthcare and education transfers are not. This prevents young agents from borrowing against healthcare benefits received later in life. In-kind consumption is inframarginal and, thus, doesn't impact utility maximization. Pension payouts are a time-varying and region- and skill-group specific fraction of an agent's average career earnings. An agent who retires at age a_r in year t receives pension benefits as a constant fraction of their average lifetime wage income.

$$PB_{a_r,t,k} = \nu_k \frac{1}{a_r - a_{w_k}} \sum_{a=a_{w_k}}^{a_r} w_{a,t-a_r+a,k}, \tag{17}$$

where ν_k is skill group k's replacement rate.

3.3 Solution Method and Calibration

This section describes our iterative solution method for precisely solving or model's over four million equations in an equal number of unknowns. It then turns to calibrating our 400 plus parameters using our automated calibration method.

3.3.1 Solution Method

We solve the model using a variant of Auerbach and Kotlikoff (1987)'s Gauss-Seidel method. Solving the model means finding region-specific paths of skill-specific wages and a global path of the return to capital such that region- and skill-specific supplies of labor equal their corresponding demands and the path of global capital supply equals the path of global capital demand.

The model's factor demands are governed by equations (11) and (12). The model's factor supplies are governed by equations (1) through (4). The model's solution begins with a guess of the path (i.e., time path) of the U.S. capital stock and guessed region-specific paths of skill-specific labor supplies. Given the guessed paths of U.S. skill-specific labor supplies, the guessed path of U.S. capital, and the U.S. METR on capital, we determine the path of the world interest rate, r_t. From the world interest-rate path and region- and skill-specific leisure guesses, we determine region-specific capital demand paths.

The implied paths of region-specific wages and the r_t path as well as guessed paths of region-specific average and marginal household tax rates lets us update region- and skill-specific paths of leisure and consumption demands and, thus, region- and skill-specific labor and region-specific asset supplies. Subtracting off the paths of the values of non-capital assets (debt and claims to energy endowments), yields a path of the world supply of capital. Subtracting the sum of paths of all non-U.S. capital demands from the path of the world supply of capital produces our new guess of the path of the U.S. capital stock.

We also update all endogenous region-specific taxes to satisfy region-specific annual government budget constraints and the paths of market values of non-capital assets.[9] A weighted average of the new and prior guesses of the U.S. capital stock path is then used in the model's next iteration.

The initial year of the simulation is set at 2017, the most recent year for which comprehensive

[9] To be precise, region-specific paths of payroll tax rates endogeneously adjust to pay for their predesignated region-specific shares of government pension benefits per equation (17). Region-specific paths of pension benefits are calculated in each iteration based on workers' past average wage earnings and an assumed progressive replacement rate function discussed in the Appendix. Regional-specific paths of non-pension transfer payments, both fungible and in-kind, are also calculated at this stage based on equation (16). As the equation indicates, non-pension transfer payments grow with gross national income (GNI) per capita and population. The same is true of region-specific government consumption. As for income tax functions, the region-specific linear components are adjusted to maintain region-specific debt to GDP ratios. This routine also updates region-specific paths of debt stocks.

macro and fiscal data are available. The initial conditions consist of asset holdings by age, skill group, and region. We allow a maximum of 500 years for the model to converge.[10] Updating stops when, for all years, the year-specific difference between global output demand and supply converges to within 0.01 percent of that year's global output.

3.3.2 Calibration

We calibrate the model's parameters to fiscal aggregates, macro aggregates, and microeconomic estimates. This subsection summarizes these targets and their sources. Appendix sections A and B provide details. Tables A12, A13, A14, and A15, which compare region-specific calibration targets with model output, show the extraordinary ability of our calibration method to simultaneously match, to a high degree of precision, hundreds of parameter values to their associated targets.

Region-specific shares of population in each skill-group are estimated using the World Inequality Database (Alvaredo et al., 2020). Capital shares in each region are calibrated to IMF capital stock data (IMF, 2019a). Production input shares and population shares by skill-group are reported in table A16. Region-specific macro and government fiscal targets are obtained from a number of data sources.[11] We use the microeconomic estimates of Acemoglu et al. (2020) and Altig et al. (2020) to calibrate, respectively, φ_K, and the 2019 progressivity of the U.S. federal income tax. We assume all regions maintain this degree of income-tax progressivity through time.[12]

Region-specific pension replacement rates in 2017 are calibrated to IMF (2017) data. Growth rates of pensions are calibrated to OECD projections of pension expenditures in 2050. Region-specific 2017 TFP values are based on each region's 2017 GDP. As mentioned, their growth rates are calibrated to long-term GDP estimates from Müller et al. (2019), who use a Bayesian model to predict country-specific output levels through 2100. The authors graciously provide us with two projections. The first projects GDP paths of all countries jointly using a multivariate (cross-country) estimation strategy. The second projects each country's GDP path using a univariate (country-specific) estimation strategy. We grow TFPs in each region at a constant, region-specific

[10]This is 250 years beyond the year each region's demographics stabilize.

[11]Most are sourced from the IMF, the World Bank, and the International Labour Organization (ILO). The remainder are sourced from country-specific treasury department (or equivalent) websites.

[12]Marginal tax rate data is generally unavailable outside of the U.S. While statutory rates are available, they do not reflect the actual tax burden on workers or fiscal progressivity (Auerbach et al., 2016).

rate between 2017 and 2100 such that, assuming no further automation occurs after 2017, each region's GDP in 2100 equals the average of these two 2100 projections.

We calibrate private consumption as a fraction of GDP by adjusting δ_t, the time preference rate, to match IMF fiscal data. Each region's δ_t is assumed to slowly converge to the GDP-weighted average time-preference rate of developed regions as of 2017, with full convergence occurring in 2100.[13] Conceptually, it is unlikely that China and India maintain their strong desire to save as their economies develop.[14] Hence, we assume time preferences of all regions gradually evolve toward our best guess of a single, global steady-state rate.

The shape of the initial age-skill asset distribution is calibrated to U.S. Survey of Consumer Finances (SCF; Bricker et al. 2017) data. Fossil fuel rent flows and the shares of rents that accrue to governments are calibrated to World Bank (2021) data. We interpret leisure as measuring labor force participation and calibrate leisure preferences to a combination of SCF and ILO data. Finally, we calibrate distributional aspects of each region's pension system (e.g., benefits by skill-group and contribution ceilings) to a combination of IMF and Social Security Administration (SSA) data.

A key advance over prior large-scale life-cycle CGE models is our automated calibration algorithm, which jointly solves for our model's 464 parameters to produce the closest match between simulated outputs and calibration targets. Calibrating each parameter by hand is infeasible given our model's scale. We sketch our method here and provide details in Appendix section B.

Some parameters are directly observable. But over 400 are not so identified. They are joint, non-linear, complex functions of the data. Their interdependence makes the standard parameter-by-parameter calibration method infeasible. Our solution partitions parameters into groups whose macroeconomic targets are most influenced by their modification. This is achieved in three steps. First, we simulate the model for a large grid of non-calibrated parameters. Second, using the results, we run LASSO regressions (see Tibshirani 1996) of each targeted output against the entire set of parameters to rank them by a unique measure of explanatory power. LASSO is advantages as it sets less important regression coefficients to zero. This produces equations relating subsets of targets to a same sized set of parameters, which, in turn, permit solving for those parameters.

[13] We define developed regions as the U.S, Western Europe, JSHK, the U.K, Canada, Australia, and New Zealand.
[14] Similarly, the U.S. strong preference to consume is, fiscally, difficult to sustain in the long run.

Third, starting from a "naive" approach – attempting to hit each calibration target with a single parameter – we run a number of candidate parameter partitionings, each of which attempts to calibrate the model by jointly updating a subset of parameters and individually updating the rest.[15] Each successive candidate parameter partitioning introduces one additional jointly updated set of parameters, with the order in which they are added determined by the ranking in step two. This process, which can be parallelized, continues until a candidate parameter partitioning is able to precisely calibrate the model.[16] In all cases, we are able to successfully calibrate our model to our desired precision over all targets in less than 200 CPU-hours.

3.4 Modeling Endogenous Automation

As indicated, we model automation and skill-biased technological change as permitting regions to produce using their initial Cobb-Douglas technology or a year-specific frontier technology, which is also Cobb-Douglas, but with larger capital and high-skilled labor coefficients.

Our approach strikes a balance between resolution and computational complexity. On the producer side, the model captures the distinction between increased dependence on an abundant factor (capital), increased dependence on a scarce factor (high-skilled labor), and TFP growth. More importantly, it is calibrated to real-world macro and micro evidence, with parameters extrapolated from historical U.S. data. On the household side, the model captures automation's impact on the supply of labor, capital, and factor prices. While the development of frontier technologies is no doubt endogenous to R&D investment, among other factors, these relationships are difficult to measure and model without greatly increased complexity. Similarly, while the effect of automation is heterogeneous by industry and occupation, such details are generally unavailable outside of the most developed countries.

Assessing the impact of changes in production-share coefficients is more complex than simply considering improvements to TFP. To understand the issue, take the production of output, Q_e, based on a Cobb-Douglas function of capital, K_e, and labor, L_e, with TFP of ϕ_e, share coefficient,

[15] The procedure is detailed in Appendix section B.

[16] We define this as a maximum of 2 percent deviation or, for targets that can be either negative or positive, 2 percentage points.

α_e, and productivity coefficients (conditional on input scaling), φ_{K_e}, and φ_{L_e}, where e denotes example.

$$Q_e = \phi_e(\varphi_{K_e} K)^{\alpha_e}(\varphi_{L_e} L)^{1-\alpha_e}. \tag{18}$$

For given measured output, Q_e, given measured inputs, K_e and L_e, and given share parameter, α_e, the product $\phi_e(\varphi_{K_e})^{\alpha_e}(\varphi_{L_e})^{1-\alpha_e}$ is defined, but not ϕ_e, $(\varphi_{K_e})^{\alpha_e}$, or $(\varphi_{L_e})^{1-\alpha_e}$ separately. Were α_e fixed, it would suffice to call the product $\hat{\phi}_e$ and consider how output changes when this product, K, L, or some combination changes. But we need separate values for ϕ_e, φ_{K_e}, and φ_{L_e} to assess the impact on output of a change in α_e. Two normalizations are available to help pin down these three values. The first is TFP scaling, namely $\phi_e = 1$. The second is the normalization $(\varphi_{K_e})^{\alpha_e}(\varphi_{L_e})^{1-\alpha_e} = 1$. This gives us one equation in φ_{K_e} and φ_{L_e}. An empirical estimate or an assumed relationship can provide the second equation.

3.4.1 Calibrating Labor and Capital Productivity Coefficients

In our model, we have five productivity parameters. We use the same first and second normalization, i.e. $\phi = 1$ and $\varphi_K^{\alpha}\varphi_{l_1}^{\beta_1}\varphi_{l_2}^{\beta_2}\varphi_{l_3}^{\beta_3} = 1$, as in the example. A third is provided by Acemoglu et al. (2020). Using French factory-level data, the authors estimates that, holding inputs fixed, a one percentage point increase in capital's share, α, is related to a 0.3 percent increase in productivity.

The remaining two equations needed to determine the four unknowns are provided by two assumptions biased against our finding major impacts of automation on output. Specifically, we assume, in our base case, that a rise in the high-skilled or mid-skill labor shares coupled with an equal-sized decline in the low-skilled share lead to no change in output. These assumptions provide a lower-bound estimate for automation's impact.[17]

To generate what may represent an upper bound estimate, we also simulate, as part of our sensitivity analysis, the model assuming a 0.3 percent output increase for a 1 percentage point

[17] Unlike capital-use intensivity, there are no empirical estimates of the output effects of changes in labor skill-group coefficients. However, given the well-documented, global increase in the share of income of high earners, it stands to reason that replacing low-skilled tasks with higher-skilled ones should not lower output.

increase in the high-skilled (mid-skilled) labor shares offset, again, by a 1 percentage point reduction in the low-skilled labor share. As detailed in Appendix section C, this assumption slightly raises regional growth rates and speeds up the adoption of the frontier technology in some regions, but otherwise makes little difference to our results.

The following four equations make our calibration of the four productivity coefficients precise. Our first normalization sets U.S. TFP to 1.0 in 2017. Equation (19) is the aforementioned second normalization,

$$\varphi_K^{\alpha} \cdot \varphi_{l_1}^{\beta_1} \varphi_{l_2}^{\beta_2} \varphi_{l_3}^{\beta_3} = 1. \tag{19}$$

Differentiating equation (7), dropping time subscripts, and holding inputs constant yields

$$\frac{1}{Q}\frac{dQ}{d\alpha} \quad = \quad \log(\varphi_K K) - \frac{1}{3}\sum_{k=1}^{3} \log(\varphi_{l_k} l_{t,k}) = 0.3. \tag{20}$$

The two remaining equations come from differentiating Q with regard to β_2 and β_3 assuming β_1 declines pari passus.

$$\frac{1}{Q}\frac{dQ}{d\beta_2} = \log(\varphi_{l_2} l_2) - \log(\varphi_{l_1} l_1) \quad = \quad \chi_2 \tag{21}$$

and

$$\frac{1}{Q}\frac{dQ}{d\beta_3} = \log(\varphi_{l_3} l_3) - \log(\varphi_{l_1} l_1) \quad = \quad \chi_3, \tag{22}$$

where χ_2 and χ_3 are both set to 0 in our baseline and 0.3 in our sensitivity analysis.[18]

[18] To be precise, we set $\chi_2 = \chi_3 = \epsilon = 0.01$ in our baseline to weakly prefer the frontier production function under scenarios when either the mid-skilled or high-skilled labor-share coefficients increase with a commensurate decline in the low-skilled labor-share coefficient. Results assuming $\chi_2 = \chi_3 = 0.3$ are discussed in Appendix section C.

This system of four equations in φ_K, φ_{l_1}, φ_{l_2}, and φ_{l_3} is solved for the U.S. for 2017. The derived coefficients are then applied to the U.S. and all other regions through time.

3.4.2 Endogenous Automation Choice

Producers in each region choose between legacy and frontier technology to maximize output. They decide whether to use the frontier technology, in each period, based on whether

$$\frac{Q^*}{Q} = \frac{\phi(\varphi_K \hat{K})^{\alpha^*} \hat{L}_1^{\beta_1^*} \hat{L}_2^{\beta_2^*} \hat{L}_3^{\beta_3^*}}{\phi(\varphi_K K)^{\alpha} L_1^{\beta_1} L_2^{\beta_2} L_3^{\beta_3}}, \tag{23}$$

the ratio of output Q^* using the frontier technology (α^* and β^* and resulting inputs $\hat{L}$ and $\hat{K}$) to output Q using the legacy technology and corresponding equilibrium inputs, exceeds 1.[19]

Note that the frontier technology changes neither the TFP term, nor how it evolves, nor factor productivities. However, a region's adoption of the frontier technology impacts its capital and labor demands. Such changes to regional factor demands impact the course of region-specific wage-rates as well as the evolution of the global interest rate. Moreover, changes in the world interest-rate path as well as region- and skill-specific wage paths will change each region's annual saving as well as skill-specific annual labor supply. These supply-side changes will further impact the evolution of interest and wage rates. Thus, one region's adoption of technology, in one or more years, influences all others regions' paths of automation by altering the global, dynamic equilibrium.

4 Results

4.1 No Automation

We first simulate our model assuming no automation occurs post 2017. As table A2 shows, this no-automation scenario entails a very major change in world economic power over this century

[19] It is easy to show, as detailed in Benzell et al. (2016), in the two-factor case that if a range of Cobb-Douglas production functions are available, firms will only choose between functions with the most extreme β values available. Therefore, our algorithm compares only the 2017 technology and the frontier technology. The adoption of frontier technology is *not* a permanent decision. A region that automated in the past may choose, in any year, to return to its 2017 production method. And, of course, it can opt to never adopt.

– one, which we shall show, is little impacted by automation. Between 2017 and the end of the century, China's share of global output rises from 16.6 percent to 25.1 percent. India's share rises from 7.0 percent to 11.4 percent. In contrast, the U.S. share of world GDP falls from 16.4 percent to 12.5 percent, and Western Europe and the UK's combined share falls from 19.7 percent to 11.7 percent.

As indicated in table A3, by century's end, China catches up to within 15 percent of the U.S. in terms of per capita output. India's per capita output, however, reaches only 26.5 percent of the U.S. level by century's end. Other regions make even less headway. In 2100, per capita output in Sub-Saharan Africa is 5.3 percent of the U.S. level – *down* from 6.0 percent in 2017. Sub-Saharan Africa's 2100 per capita output barely exceeds that of China in 2017.

Given their exceptionally high saving rates, China's rapid productivity catch-up growth and India's rapid demographic growth result in a massive buildup of assets in these regions. Table A4 documents this. In 2017, 2050, and 2100, China accounts for 17.4 percent, 35.7 percent, and 34.9 percent of global assets, respectively. For India, the corresponding three shares are 6.6 percent, 11.7 percent, and 13.1 percent. China and India's asset holdings are invested globally, including in the U.S. In 2050 and 2100, 32.9 percent and 19.5 percent of China's assets are invested abroad. The corresponding figures for India are 15.4 percent and 12.5 percent.

The combined advanced economies' – U.S., U.K., CAN, WEU, and JSHK– global asset share falls from 46.8 percent in 2017 to 20.2 percent in 2100. For the U.S., the 2017 to 2100 decline is from 16.8 percent to 5.4 percent. Yet its relatively high capital share and high worker productivity levels make it a good place to invest. Hence, its 2100 share of world capital usage – 13.6 percent – is still relatively high. The U.S.'s (negative) net foreign asset position grows from 34 percent of GDP in 2017 to over 200 percent by 2050.[20]

The huge glut of savings drives down the global rate of return to roughly 2.8 percent by 2050 and 2 percent by 2100, with three major implications. First, revenue raised by capital income and corporate income taxes will fall substantially across all regions, requiring compensating increases in income and consumption taxes. Second, wages increase globally. By 2050, U.S. workers' wages

[20] While this increase is striking, it is in line with extrapolations of the net International Investment Position (IIP) data from 2000 to 2017 as well as the actual net IIP for 2018-2020.

increase by 45.5, 47.1 and 52.9 percent relative to 2017 levels for low-, mid-, and high-skilled workers respectively. In China, they are projected to increase by 223.5, 217.9, and 222.9 percent. Third, although there is no automation in this scenario, global capital abundance would tend to stimulate demand for the development of more capital-intensive technologies.

In addition to losing economic hegemony, the U.S. faces, per our simulation, major fiscal challenges through the remainder of the century. As table A5 shows, tax rates increase significantly over time in the U.S. even under our optimistic assumption that debt remains fixed relative to GDP. As the U.S. ages, pension expenditures grow from approximately 5 percent of GDP in 2017 to 10 percent in 2100. At the same time, the shrinking corporate income tax base and the U.S.' reliance on income as opposed to consumption taxation compounds the fiscal burden on successive young generations. It also significantly undermines work incentives. In 2017, American workers face, on average, a marginal effective tax rate on labor earnings of 40.8 percent. This rate reaches 55.4 percent by 2100.

4.2 Baseline Findings

Our baseline simulation incorporates automation with the U.S. frontier technology factor-share coefficients changing at the same pace through 2050 as they did between 1980 and 2017. As for the other regions, their frontier capital, high-skill, and mid-skill technology share coefficients change at the same annual percentage rate as the U.S. share coefficients. Recall, each region, in each year, can use its 2017 technology or its current frontier technology. Hence, as in Brezis et al. (1993), the model permits less advanced regions to leapfrog more advanced regions in their use of advanced technology.

Tables A6 through A10 contrast the frontier technology's input shares in 2050 with shares of the 2017 production function. Between 2018 and 2050, frontier technology parameters are assumed to evolve at a constant rate over time. In 2050 and thereafter the U.S. has the option of producing with the 2017 technology: $\alpha_t = 0.341$ and $\beta_{t,k} = 0.22$, or with the 2050 frontier technology: $\alpha_{2050} = 0.369$, $\beta_{2050,1} = 0.161$, $\beta_{2050,2} = 0.205$, and $\beta_{2050,3} = 0.264$.

Some regions adopt the frontier technology as soon as it is introduced. Some, such as India and the Middle East, never do. Others adopt in the future. For example, China automates in 2027 and Southeast Asia automates in 2064. Regions' automation choices are displayed in figure 3. Overall, the choice and timing of adopting frontier technologies strongly correlates with the level of economic prosperity. The five richest regions in our model, measured by GDP per capita, all automate immediately. On the other hand, three of the five poorest regions never automate, and the other two automate only after 2050. While they have the option to do so, no region in the baseline chooses to return to the 2017 technology after adopting the frontier technology.

Figure 3: Choice of Production Function by Region

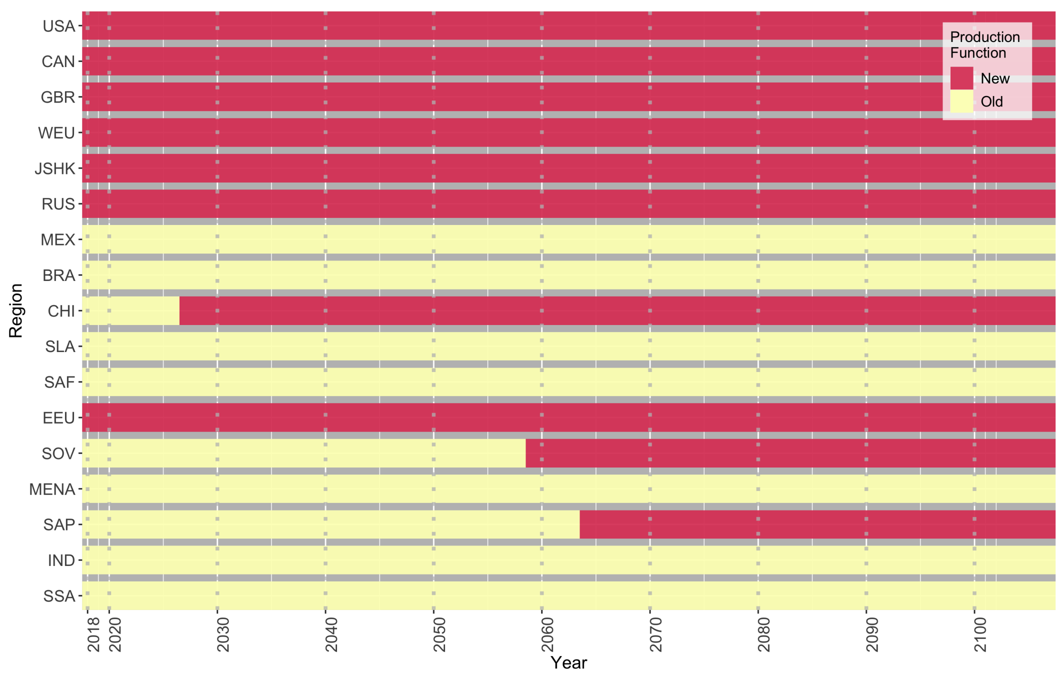

Why do some countries benefit more from automation than others, and why do some countries never choose to adopt? First note that, as derived in (20), the marginal product of automation (defined here as the partial derivative of output with respect to an increase in capital's share) is increasing in the ratio of effective labor to effective capital. Output is always positive, so a country will automate so long as the supply of effective capital is larger than the supply of effective labor. In other words, automation requires

$$\log(\varphi_K K) \geq \frac{1}{3} \sum_{k=1}^{3} \log(\varphi_{l_k} l_{t,k}), \tag{24}$$

i.e.

$$\varphi_K K \geq \prod_{k=1}^{3} \varphi_{l_k} l_k. \tag{25}$$

The question then becomes what determines the ratio of effective labor to effective capital. Three of the most important factors are regions' starting capital intensivity, share of high-skilled labor, and most importantly, TFP. To see why, we focus on a simplified case. Consider a small open economy with Cobb-Douglas production over capital and one type of labor l in fixed supply. Output is

$$Q = \phi(\varphi_K K)^{\alpha}(\varphi_l l)^{1-\alpha}, \tag{26}$$

where, setting the marginal product of capital equal to the interest rate r, rearranges to

$$\frac{\varphi_K K}{\varphi_l l} = \left(\frac{\alpha \varphi_K \phi}{r}\right)^{\frac{1}{1-\alpha}}. \tag{27}$$

A country will automate so long as the right hand side of (27) is greater than one. This ratio is increasing in $\varphi_K \phi$. φ_K is equal in each region as shown in (19). Hence, ϕ is the only productivity term that varies across countries. Countries with higher TFPs are better at making use of capital. Hence, they tend to have higher capital to labor ratios. With factor price equalization, they demand and, necessarily, import more capital per unit of labor and benefit more from automation.

The starting capital share of a region, α, plays a more complex role. For regions where $\frac{\alpha \varphi_K \phi}{r} > 1$, a high initial capital share unambiguously boosts the positive productivity effect of automation. On the other hand, when $\frac{\alpha \varphi_K \phi}{r} < 1$, a larger α worsens the negative effect on output. In other words, initial α intensifies the productivity effect of automation. A one percentage point increase in α from 50 to 51 percent increases the ratio of capital's share to labor's marginally. A one percentage point increase in α from 98 percent to 99 percent approximately doubles it.[21]

[21] To get an intuition for why larger initial α's intensify the output response of automation, consider the case of an small open economy moving from 99 percent capital intensity to 100 percent. Such a change would transforms it from a Solow into an AK economy with knife-edge growth characteristics. If the country is adept at using capital

Unlike capital, labor is not globally mobile. As φ_l is equal across countries and skill-groups receive equal shares of total labor income, equation (22) makes clear that, all else equal, countries with a larger share of high-skilled labor will benefit more from skill-biased automation, and those with a smaller relative quantity will benefit less.

Each of these analyses are complicated by multiple factors. Demographics, government policies and preferences also influence the decision to automate. Regions with aging populations such as Japan and Western Europe maintain large payroll taxes to sustain their pension systems, leading to higher post-tax costs of labor. Higher post-tax costs of labor incentivize firms to automate.[22] Regions with supply-elastic high-skilled workers who reduce their hours significantly in response to higher wages are less likely to automate. Similarly, regions with more inequality are disinclined to automate because they face larger wealth effects from skill-biased technologies.

Finally, the outcome of figure 3 is a consequence of the global distribution of capital. Corporate taxes introduce wedges between rates of return across countries. Countries with low corporate tax rates face lower post-tax costs of capital, and therefore firms in those countries are quicker to automate. When r_t is low, all countries use a higher ratio of capital to effective labor and benefit relatively more from automation. Over time, more countries automate due to decreasing interest rates.

If the increase to capital demand from automating exceeds the savings response, some regions will end up with less capital compared to a world where no region automates. The actual impact of automation on r_t is, however, relatively small. As figure 4 shows, automation drives up r_t, but by only 0.1-0.4 percentage points above the no-automation path. This weak response is a result of a limited number of regions automating and, more importantly, high savings rates of China, India, and, in the latter half of the century, the Middle East. It takes much more capital- and high-skilled labor-biased technological change to absorb the world's accelerating stock of assets.[23]

The weak response of r_t has two important consequences for welfare and policy. First, the

(i.e. $\varphi_K\phi$ is large), it could hypothetically absorb a huge share of the world's capital supply and increase output massively. The reverse is true if the country is very bad at using capital: automation will come at a huge cost to the country's output.

[22] Table A18 summarizes calibrated tax rates by region in 2017.

[23] Section 4.3.4 discusses these simulations in detail.

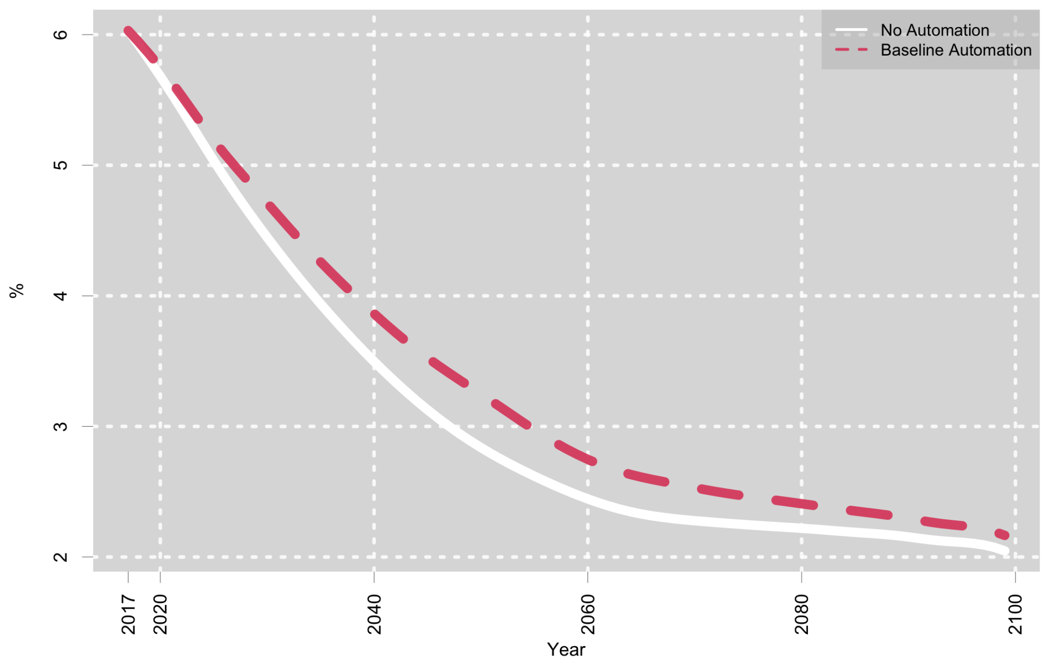

Figure 4: Global Post-Corporate Tax Rate of Return

primary channel through which automation could benefit households nearing or in retirement is by raising investment returns (Benzell et al., 2015b). But since automation doesn't matter much to the global interest-rate path, the current old gain little relative to high-skilled young workers. Second, automation does not substantially increase government borrowing costs. Hence, debt-financed redistributive policies are not made prohibitively expensive by global capital demand.

Automation produces winners and losers across regions in terms of GDP. Comparing regions in 2050, JSHK gains the most with a 12.9 percent GDP increase relative to no automation. The regions with the second, third, and fourth largest gains are Western Europe, CAN, and the U.K, which respectively experience gains of 12.7, 12.1, and 9.9 percent. The U.S. experiences the fifth largest GDP increase relative to no automation – 5.1 percent in 2050. Global output in the same period increases by 3.4 percent. Table A11 summarizes GDP gains and losses in the baseline relative to no automation. Figure 5 contrasts GDP percent changes, relative to no automation, of the U.S., Western Europe, China, Mexico, Sub-Saharan Africa, and the world average.

The economic windfall from automating technologies accrues disproportionately to a small fraction of workers. Sixty-three percent of all adults alive in 2017 experience an increase in remaining

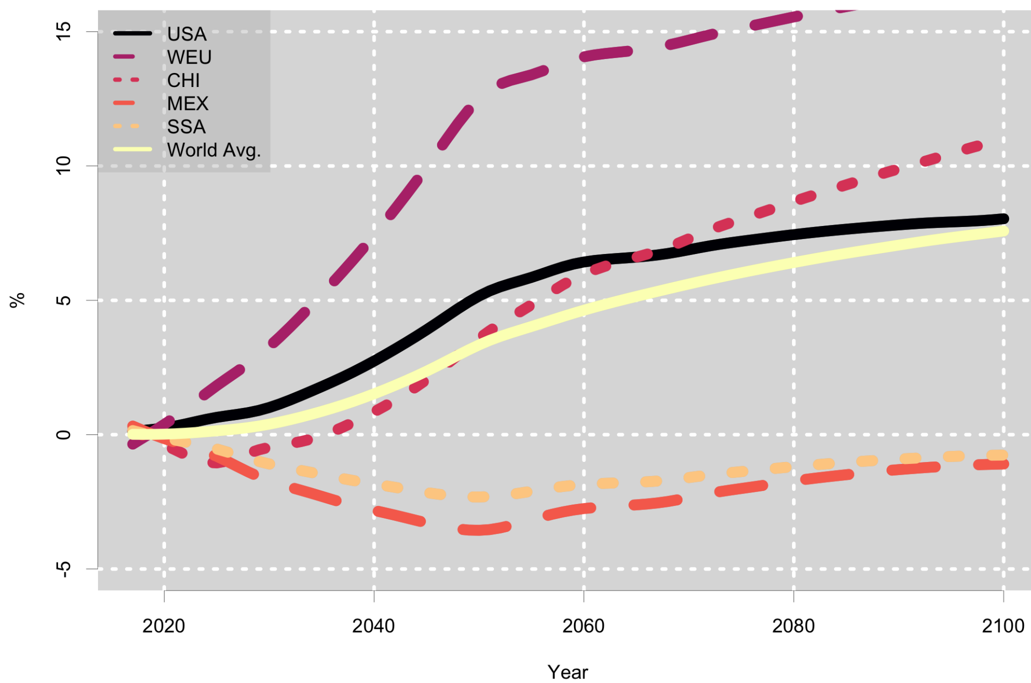

Figure 5: GDP Percent Change in Baseline Relative to 'No Automation' Scenario

lifetime welfare from baseline automation relative to no automation. However, this rate is deceptively high because of gains to retirees. Among all adults alive in 2050, only 28.5 percent experience an improvement, a fraction that decreases to 22.5 percent by 2100. On the other hand, automation makes nearly all high-skilled workers globally better off – 92.5 percent of those alive in 2050 and 78.2 percent in 2100 – but these workers only account for 5 percent of the world's adult population. [24] Wage changes are more striking: only 6.8 percent of global workers in 2050 – the high- and, in a few regions, mid-skilled workers – see wage growth under baseline automation relative to no automation.

Globally, the average worker is made slightly *worse* off by automation throughout most of the 21st century, by 0.9 percent in 2050, and 1.1 percent in 2100 relative to no automation. The global average wage increases by 65.3 percent between 2017 and 2050, only slightly higher than the 64.3 percent of the no-automation scenario. These trends are summarized in figure 6. By 2050 and relative to no automation, the global average low-, mid- and high-skilled workers experience, respectively, a lifetime welfare decrease of 2.3 percent, an increase of 2 percent, and an increase of

[24]The mid-skilled and low-skilled groups represent 19.7 percent and 75.3 percent, respectively.

10.3 percent.

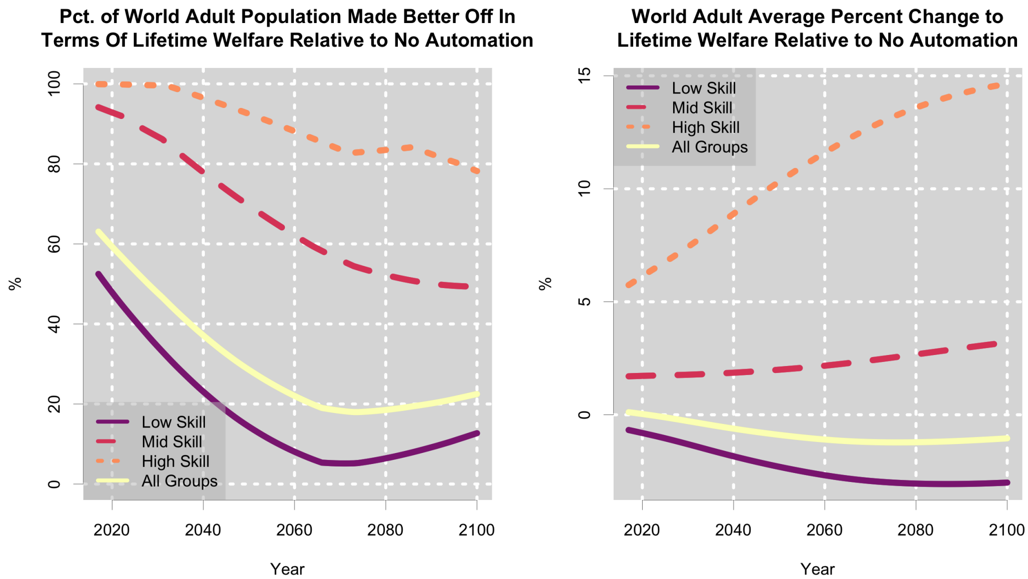

Figure 6: Global Lifetime Welfare Consequences of Baseline Automation

Not all skilled workers fare equally well from automation. Figure A3 presents the percent of adult population made better off in immediately-automating and never-automating regions.[25] As shown, virtually no mid- or high-skilled workers in immediately-automating regions are ever harmed by new technologies. In never-automating regions, some high-skilled workers are harmed, but only 20.2 percent in 2050 and 37.8 percent in 2100. This is because high-skilled workers, regardless of region, are disproportionate asset holders and, per our calibration, supply labor robustly in response to a reduction to their marginal product. Most future low-skilled workers are made worse off in both types of regions. However, in the very long run, increased growth from capital accumulation partially offsets their immiseration.

Figure A4 summarizes welfare changes by skill group and year of birth. Globally, no high-skilled workers and virtually no mid-skilled workers born before 2050 are made more than 1 percent worse off in terms of lifetime welfare. However, among workers born in 2020, only 47.7 percent of high-skilled and 27.3 percent of mid-skilled workers are made at least 1 percent better off. Workers in regions that delay or never adopt the frontier technology suffer regardless of skill level, but their welfare loss is tempered by consumption-leisure substitution and, for those with wealth, slightly

[25] Regions that delay automation (e.g. China, Southeast Asia) are not included in either sub-figure.

better returns on assets. Regions that do automate significantly lower the pre-tax wages of low-skilled workers. But the progressivity of tax and pension systems prevent such workers from bearing the full impact of rising wage inequality. For example, low-skilled workers in Japan born after 2027 are made better off assuming baseline automation, and no low-skilled cohorts in Western Europe are made more than 2.1 percent worse off.

We now focus on the U.S. Despite always choosing the frontier technology, the U.S's relative economic power is only marginally bolstered. Figure A5 presents U.S. GDP and capital stock as a share of the world's values with and without automation. Relative to no automation, the U.S's share of world GDP increases by only 0.24 percentage points or 1.7 percent in 2050 with baseline automation. Its share of world capital, thanks to capital inflows, increases by 0.7 percentage points to 15.8 percent. In the same year and relative to no automation, U.S. gross investment increases from 33.5 percent to 35.7 percent of GDP.

Automation primarily benefits high-skilled U.S. workers. Their wages increase by 28.8 percent in 2050 and by 37.5 percent in 2100 relative to no automation. Low-skilled U.S. workers see their wages fall by 22.9 percent in 2050 and 21.7 percent in 2100. Older workers also benefit from automation but by a much smaller amount. A high-skilled worker born in 1951 experiences, relative to no automation, an increase in remaining lifetime welfare of only 1.5 percent. Mid- and low-skilled workers born in the same year experience welfare increases of 0.8 percent and 0.3 percent.

As figure A6 shows, the average U.S. worker born in 1950, 1980, and 2020 are made, respectively, 0.4 percent better, 4 percent worse, and 8 percent worse off in terms of lifetime welfare from automation. However, the young and high-skilled enjoy extraordinary gains, with the average high-skilled worker born in 2020 and alive in 2020 being made, respectively, 24.2 and 6.2 percent better off.

The welfare gains to the high-skilled and the losses to the low-skilled are smaller than their corresponding wage changes. This is due to our calibration of leisure and consumption as gross complements, generating an income effect that dominates the substitution effect. When the high-skilled's wage rate increases, they demand more leisure, which is not made more productive by the new technology. Correspondingly, when the low skilled see their wages decrease, they work longer hours to partially offset decreased consumption.

4.3 Alternate Trade, Technology, and Labor Supply Scenarios

We simulate alternate scenarios to provide insights into possible futures and clarify the importance of our modeling assumptions to the baseline results. We consider five scenarios. Two are meant to illuminate the role of international capital markets in determining the economic consequences of automation technologies. Two consider alternate technology scenarios. The final scenario considers augmenting shares of high-skilled workers.

4.3.1 Automation in the US Without Trade

A natural question about our findings is the extent to which simulating entire the globe matters for the model's conclusions. To investigate this issue, we calibrated a version of the model with only the U.S. All demographic and productivity parameters are calibrated as in the baseline model, but the U.S's starting level of assets is modified so as to produce a domestic rate of return on capital in the initial period identical to the baseline simulation. Consistent with observed private sector consumption behavior, U.S. households are calibrated as being relatively impatient. Hence, as shown in Figure A11, the U.S's capital accumulation is much slower when it is isolated from the rest of the world.

Figure 7 reports U.S. GDP growth rates in four scenarios: with and without automation, and with and without global trade. In the single-country calibration, U.S. growth rates are slower with or without automation through 2050. Most striking, however, is the fact that automation does not significantly increase the rate of growth in the autarky scenario. U.S. GDP increases by less than 0.1 percent relative to no automation in 2050, and by approximately 1.1 percent by 2100.[26] Consequently, while the inequality impact of automation is broadly similar to that of the baseline, the welfare loss of future low-skilled U.S. workers from automation deepens by 2.9 to 4.5 percent.

[26] In the very short run, GDP is actually very slightly lower due to automation in the no-trade scenario. This is because we force technology adoption in this scenario. Output decreases in part due to a negative high-skill labor supply response. High-skilled wages and labor supply in 2050 are, respectively, 20.8 percent higher and 17.3 percent lower in the no-trade baseline scenario versus the no-trade no-automation scenario.

The welfare gain of high-skilled U.S. workers decreases by a similar amount.

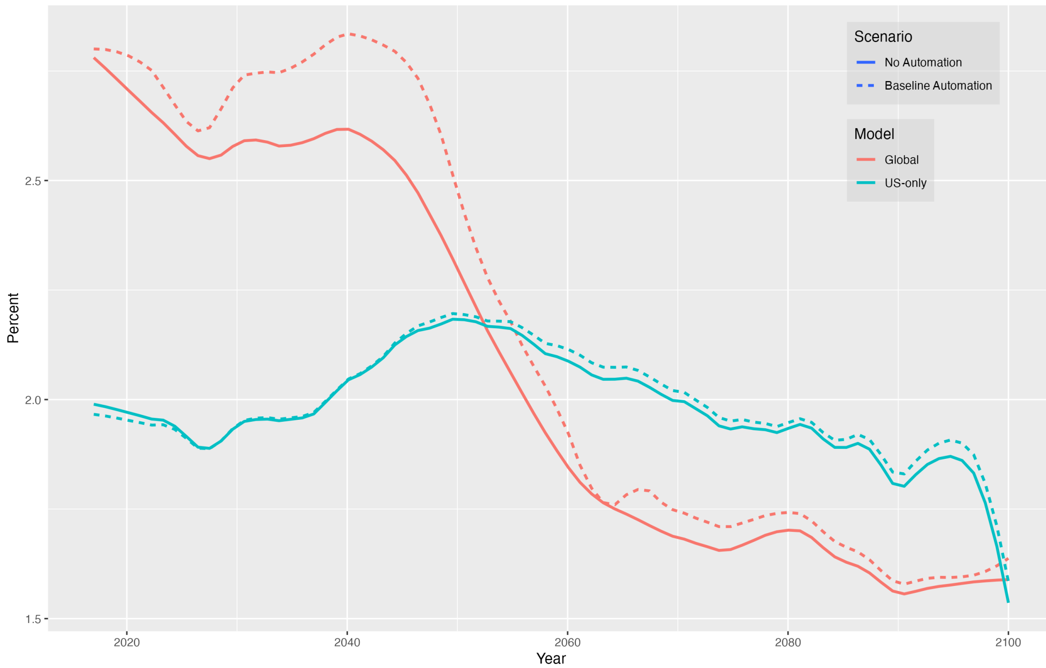

Figure 7: Annual GDP Growth Under Global and U.S-only Models

When capital is relatively scarce, and labor relatively abundant, the ratio of wages to the interest rate increases. To firms, this incentivizes automation, with labor-substituting technologies becoming relatively undesirable when rates are high. As Figure 8 shows, the no-automation interest rate under autarky is significantly higher than under trade even with automation, and the increase in interest rates due to automation is also higher under autarky. In 2050, the single-country model projects interest rates increasing from 6.6 to 7.6 percent due to automation. In comparison, rates are only 2.9 and 3.3 percent without and with automation, respectively, in the global model.

4.3.2 Simulating U.S-Only Automation

In this scenario, we utilize the global model but assume that automating technologies are made exclusive to the U.S. As in the baseline, the U.S. automates immediately. The global rate of return is now virtually unchanged from the no-automation scenario, with increased U.S. demand for capital driving up r_t by only 0.06 percentage points in 2050 and 0.02 percentage points in 2100.

Since other regions are not automating and, therefore, are not demanding more capital, the

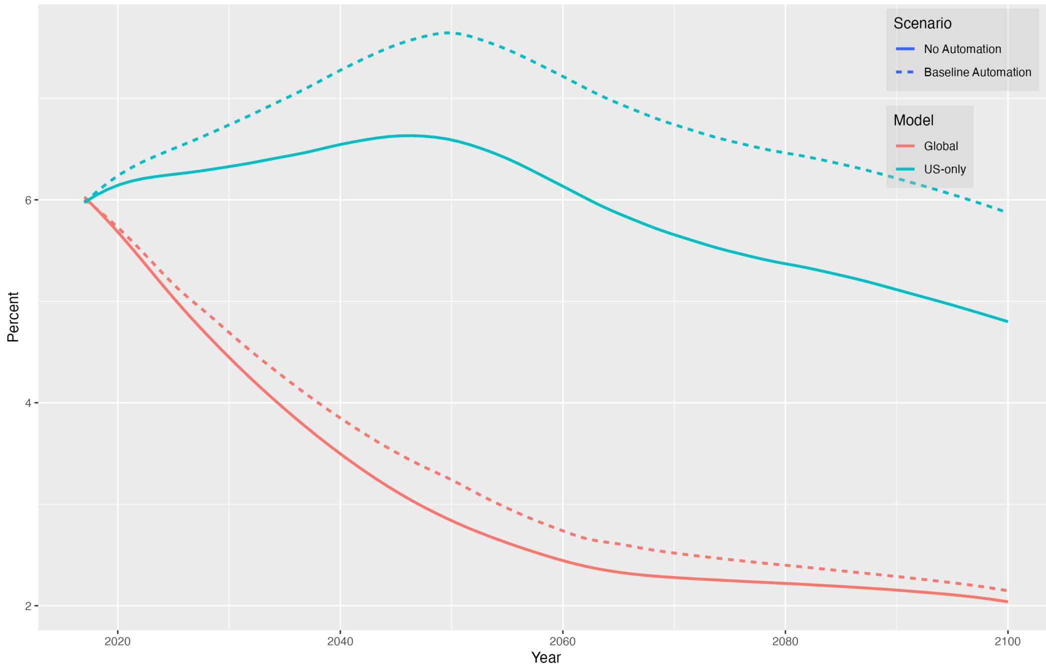

Figure 8: Post-Corporate Tax Rate of Return Under Global and U.S-only Models

U.S. benefits more. Relative to no automation, 2050 U.S. GDP is 8.3 percent higher and the 2050 U.S. capital stock is 15.5 percent higher. This increase to U.S. output has a tiny net positive effect on the rest of the global economy; relative to no automation, aggregate GDP of non-U.S. regions in 2050 increases by 0.1 percent. The U.S's gains are, however, insufficient to prevent it from declining as a share of the world economy. In 2050, the U.S's share of world GDP with U.S-only automation is 14.8 percent, higher than under global automation, but still lower than the 2017 share of 16.4 percent. Hence, even were the U.S. able to prevent other countries from automating, its share of global GDP would nonetheless fall.

Results for wages and welfare are largely consistent with section 4.2 and a reversal of findings in section 4.3.1. Gains to young, high-skilled workers are moderately larger – a lifetime welfare improvement of 17.9 percent for high-skilled U.S. workers born in 2000, in contrast to 16.7 percent under baseline automation – and losses to young, low-skilled workers are ameliorated through faster economic growth. 2017 retirees experience almost no change in lifetime welfare relative to no automation, as their pension payouts are not impacted by future economic growth.

4.3.3 Automation Without Increasing Wage Inequality

Our third alternate scenario considers the emergence of a new technology which only increases capital's share, while being skill neutral. For example, a general purpose AI algorithm might reduce demand for a wide range of occupational categories from customer service to management and executive positions. Indeed, research into large language models has suggested that these might disproportionately displace high-skilled labor, offsetting the skill-bias of other types of automation (Eloundou et al., 2023). We assume an identical share shift toward capital as section 4.2, amounting to a 4.25 percent decline in the share of labor by 2050. The new shares are specified in tables A6 through A10.

In this scenario, fewer regions with a relative abundance of high-skilled workers immediately adopt the new technology.[27] With the only productivity gain from automation coming from capital, a number of regions that immediately adopt in the baseline, such as Japan and Western Europe, delay adoption. Regions with more income inequality, on the other hand, prefer technologies which do not further enrich scarce high-skilled workers. Brazil and Mexico, for example, never adopt baseline automation but eventually adopt non-skill-biased automation. Globally, the capital accumulation rate falls between that of no automation and baseline automation. The same is true for global GDP, which increases by 0.5 percent relative to no automation in 2050, and by 3.2 percent in 2100.

Relative to baseline automation, the U.S. is made slightly better off with non labor-share dispersing technology, mostly through attracting a larger share of global capital. In 2050, this scenario increases U.S. GDP by 6.1 percent relative to no automation and by 0.9 percent relative to baseline automation. Welfare gains are comparable across skill groups. low-skilled workers born between 1990 and 2050 are 0.8-4.3 percent better off relative to no automation, and mid- and high-skilled workers born in the same periods are 2-7 percent better off.[28] Relative to baseline automation, U.S. wages in 2050 are 2.6 percent higher for low-skilled workers, and 2.7 percent lower for the high-skilled.

[27] Figure A7 summarizes the choice of production function under this scenario.

[28] Mid-skilled workers gain nearly as much as high-skilled workers. The high-skilled own more assets but also face higher marginal taxes.

These results suggest that the U.S. would be better off with automating technologies that do not further enrich high earners. This may contrast with the intuition that the most productive members of society should be given more tasks to perform. However, a non-skill biased technology economizes on the expensive labor of the high skilled and is therefore better both for economic growth and wage equality. An important additional reason is that workers in the model, as calibrated, have an income effect that dominates the substitution effect. Therefore, lower wages for the high-skilled increases their supply of labor.

4.3.4 Faster Rates of Automation

Our fourth alternate scenario keeps relative labor-skill shares of the frontier technology fixed relative to the baseline, but considers larger increases in capital's share. Specifically, we simulate a linear reduction in labor's total share culminating, in 2050, in 8.5, 21.3, and 42.5 percent declines. These declines reduce labor's overall share at rates twice, five, and ten times the recent U.S. experience. Recall that our baseline scenario has capital's share rising from 34.1 percent in 2017 to 36.9 in 2050. The three scenarios consider capital shares rising to 39.7, 48.1, and 62.1 percent of output by mid-century respectively.

These simulations produce three findings. First, while regional choices of whether to ultimately adopt the frontier technology are not greatly influenced, the timing of these choices does change. As shown in figures A8 and A9, all regions that adopt the frontier technology in the baseline scenario eventually adopt it under each of three more rapid technological change scenarios. But the U.K, for example, chooses to immediately automate in the baseline but only starting in 2028 when capital's share grows at the five times higher-than-baseline pace.

The reason for this is our second finding. The world interest rate falls by less, and can even grow above 2017's level, under these scenarios. Figure A10 displays the world interest rate under no automation, baseline, and faster capital-share growth scenarios. Strikingly, even a technology that, if adopted, has a capital input share of over 50 percent will not raise the interest rate in 2050 above the 2017 level. By mid-century, automation which advances at twice the baseline rate leaves r_t at 3.7 percent, rather than 3.3 percent with baseline automation and 2.9 percent absent all automation. At five times acceleration, r_t is 5.7 percent in 2050, still below the 2017 level of 6 percent. It takes an extremely large increase in capital's share to increase interest rates above 2017 levels in 2050. At ten times the rate, r_t is 8.7 percent in 2050. Despite the strong headwind of a global savings glut, a faster rate of increase in capital's maximum share is enough to raise the interest rate and discourage adoption of the frontier technology in some regions, at least for a time.

Finally, as a result of immediately automating, the U.S benefits increasingly from a faster rate of automation. With a relatively capital-intensive economy, high TFP, and low effective corporate taxes, the U.S. is among the regions best poised to adjust to more capital-intensive technologies. This is shown in figure 9. By 2050, U.S. output increases, relative to our no-automation baseline, by 11.9, 37.3, and 115.3 percent, respectively, for technologies that substitute away from labor at twice, five times, and ten times the historical rate. Lifetime welfare gains are similarly substantial, with the average U.S. worker born in 2020 experiencing 109.6 percent higher lifetime welfare relative to no automation, at the highest assumed rate of reduction in labor's composite share.

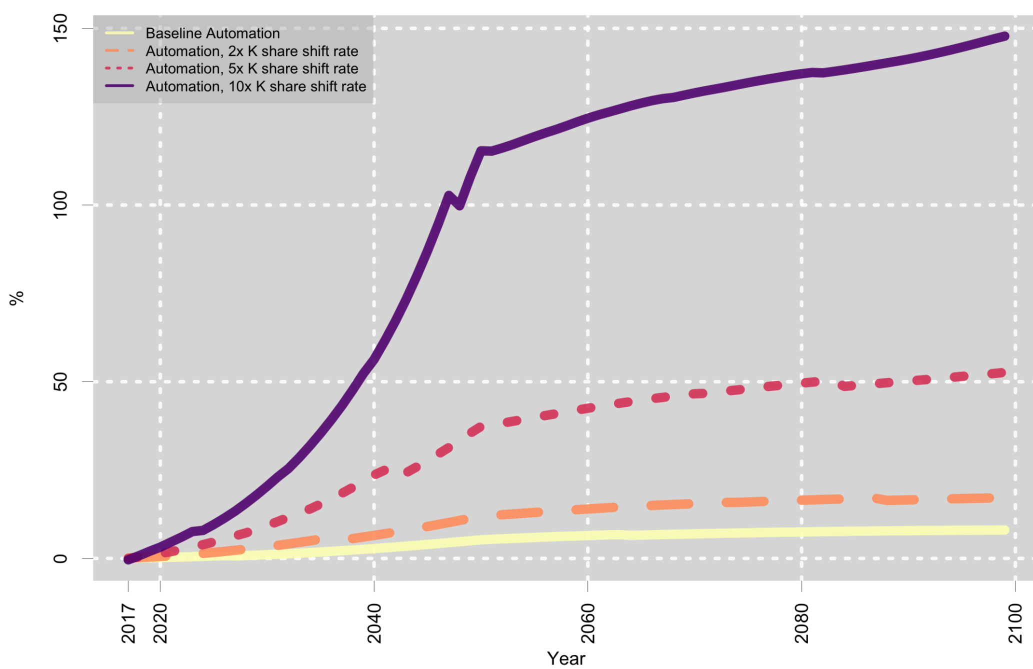

Figure 9: U.S. GDP Percent Change Relative To No Automation Under Alternate Scenarios

4.3.5 Changes in the the Relative Supplies of Skill Groups

One response to automation's induced increase in skill premiums is for more low-skilled young people to join the ranks of the mid- and high-skilled. This response is suggested by data on changes in the number of employees in skilled occupations (Deming, 2017; Benzell et al., 2019) and on workers holding advanced degrees. In the U.S., for example, the number of Masters' and Ph.D degrees conferred has increased from 0.4 million in 1991 to over 1 million in 2018 (NCES, 2019), while total nonfarm employment only increased from 109.1 million to 147.8 million (Bureau of Labor Statistics, 2022).

We consider skill-supply response simply by positing that an increasing share of new entrants to the work force are mid- and high-skilled. In these scenarios, in order to fairly contrast projections with our baseline, we adjust the rate at which frontier technologies become more skill biased. Had we assumed in the baseline that the share of workers providing high- and mid-skill labor was increasing, we would have estimated a different rate of change in their Cobb-Douglas input-share parameters. We therefore increase the rate at which the frontier technology changes to match the high- and mid-skill groups' observed relative wage growth. See Appendix section D for a detailed discussion of our approach.

We assume that population shares of skill-groups evolve gradually with the skill shares of newborns stabilizing in 2050. The rate of change is calibrated to past U.S. experience, utilizing three estimates based on alternate definitions of what constitutes a high- or mid-skilled worker. The three estimates are based on, respectively, a distributional assumption on wages, National Center for Education Statistics (NCES) data on degrees conferred, and BLS occupation-level data. All three skill-growth estimates suggest a significant rise in the population shares of both mid- and high-skilled U.S. workers by 2050. As Appendix table A22 reports, there will be a 4.1, 3.5, and 6.5 percentage point increase in the U.S.' share of workers who are mid-skilled in 2050, according to the wage-based, degree-based, and occupation-based approaches respectively. The high-skilled population share increases by, respectively, 1.1, 0.6, and 1.5 percentage points.

We first simulate increases in the relative supply of skilled labor for only the U.S.[29] Figure 10

[29]Non-U.S. regions are assumed to have fixed population shares and has access to the same frontier technology as

displays the change in U.S. GDP over time relative to no automation, under our baseline technology with no human capital accumulation and assuming that only the U.S. experiences an increase in the relative supply of skilled workers. Under all scenarios, the U.S.' gain from automation is considerably larger than under the baseline. The three scenarios entail a 15.9 to 29 percent higher U.S. level of mid-century GDP than the no-automation scenario. By comparison, our baseline automation scenario features only a 5.2 percent increase. By 2100, the three scenarios entail GDP increases of 33.8 to 41.3 percent relative to no automation – more than four times higher than the baseline's 8 percent increase.

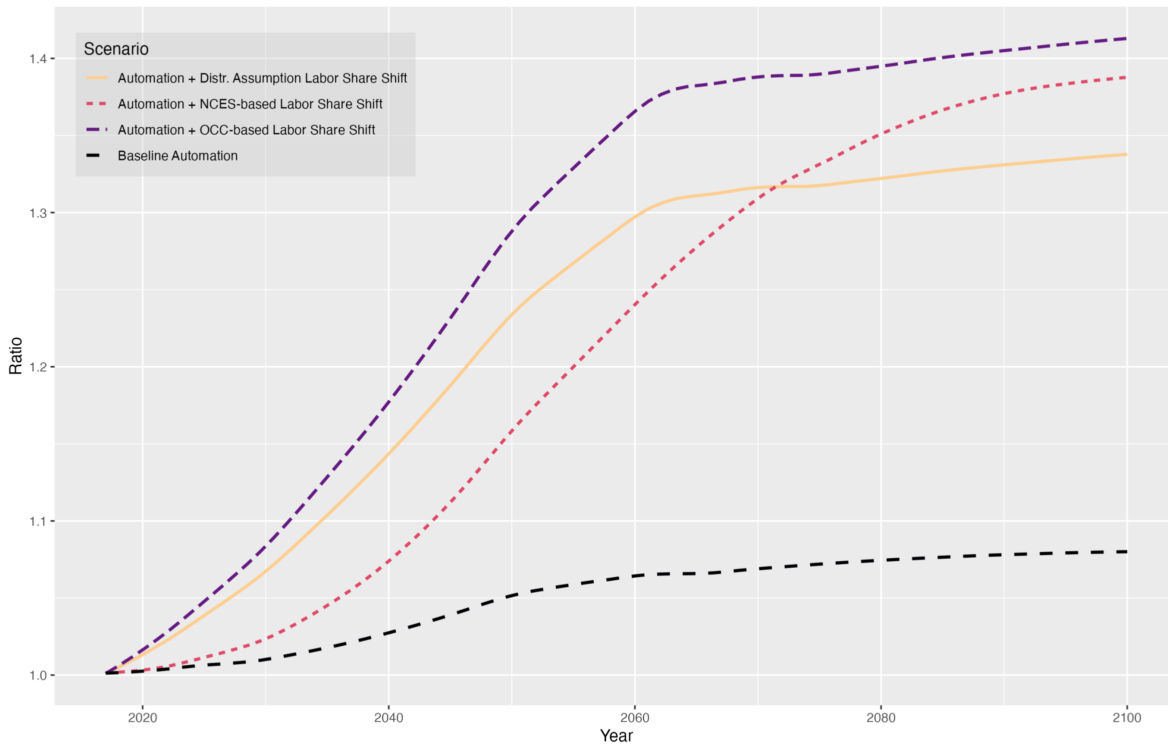

Figure 10: U.S. GDP Relative to No Automation, U.S-only Human Capital Accumulation Scenarios

Appendix figure A16 reports the welfare impact of automation with human capital accumulation. No matter the skill growth assumption, all age-skill cohorts born after 2025 are made better off, versus both baseline automation without skill share changes and no automation. For the high-skilled, the additional competition lowers wages slightly versus baseline automation for cohorts born between 1950 and 2000. However, in the long run, the strong additional growth created by the additional high-skilled workers is enough to offset the increased competition, especially for future

the baseline.

birth cohorts.

For those who remain low-skilled, the skill-growth changes are unambiguously good. Depending on the specific skill growth scenario, human capital accumulation is sufficient on its own to make almost all generations better off from automation. For the mid-skilled, skill growth has the most ambiguous effects. All mid-skilled generations after 2025 are made better off. Yet in the scenario with the fastest mid-skill employment growth, a number of generations born between 1950 and 2000 are made worse off.

Appendix figures A15 and A17 repeat these analyses, but assume that all regions experience the same labor share shifts.[30] The difference for US welfare or GDP growth is not substantial between these scenarios and scenarios where only the U.S. experiences a skill-share shift. Whether the growth in high-skilled workers is U.S-only or global, it has strongly positive effects on both reducing inequality and accelerating economic growth in the presence of automation.

5 Conclusion

This paper develops a large-scale, multi-region, computable general equilibrium model to evaluate the global consequences of automation. Automation is modeled as capital and high-skill biased technological change – an implication of micro-based structural analyses. Our general equilibrium model is highly detailed, incorporating region-specific demographics, production technologies, fiscal systems, and TFP growth rates. The model also features a highly detailed micro-foundation, including endogenous technology choice and labor supply, age-specific fertility and net immigration, idiosyncratic mortality risk, and progressive taxes and pensions.

Our baseline automation scenario extends the 1980-2017 observed trend in U.S. input shares. Specifically, the U.S. labor share falls from 65.9 to 63.1 percent, the high-skilled labor share rises from 22.0 to 26.4 percent, and capital's share rises from 34.1 percent to 36.9 percent. These changes may seem modest, but they materially impact inter- and intra-regional inequality. Through the mechanism posited by Zeira (1998b), automation benefits developed regions such as the U.S. and

[30]In the scenarios where all countries experience a rising share of mid- and high-skill workers, this growth is proportional to U.S. changes, starting from their respective initial levels.

EU, while leaving the least developed regions worse off.

In the U.S., high-skilled Americans born in 2020 are 24.2 percent better off due to automation; low-skilled Americans born in the same year are 12.8 percent worse off. Globally, 71.5 percent of people alive in 2050 are, on average, worse off thanks to automation, but those with high-skills are, on average, 10.3 percent better off. Relative to no automation, low- and high-skilled U.S. workers' wages decrease by 22.9 and increase by 28.8 percent, respectively, in 2050. Chinese low-skilled wages decrease by 22.7 percent, and high-skilled wages increase by 26.6 percent.

Regions with high TFPs, low corporate tax rates, and other structural features that lead to high wages, such as the U.S., Japan and Western Europe, benefit the most from automation. Consequently, they tend to immediately adopt frontier technologies. Less developed regions delay automation or never automate. Nations that fail to adopt receive less investment, as capital flows to automating regions, and produce less output. The region which is most negatively impacted is Mexico. Should global automation proceed at the projected rate, their GDP will be 3.6 percent lower in 2050 than it would have been. This is due, in large part, to a 9.2 percent relative decline in their capital stock. China, while not adopting the frontier technology immediately, benefits approximately as much as the U.S. from global automation by 2060. Therefore, automation has a relatively small impact on the future balance of economic power.

The U.S. would benefit more – by an additional 1 to 3 percent of GDP – were it the only region to automate or have access to the new technology. On the other hand, if the U.S. were to cutoff foreign investment, the gains from automation would be minimal, as the U.S. would not have enough capital to fully benefit from the new technologies. Automating technologies that do not disproportionately benefit high-skilled workers induce slightly more economic growth in the U.S. while reducing wage inequality.

A more rapid increase in capital's input-share, as augured by recent advances in AI, would increase benefits for automation adopting countries. By 2050, U.S. output increases, relative to our no-automation baseline, by 11.9, 37.3, and 115.3 percent, respectively, for technologies that substitute away from labor at twice, five times, and ten times the historical rate. Larger increases in automation also lead to larger increases in capital demand and interest rates, delaying automation in some regions. However, due to the headwind of a global saving glut, automation would need to proceed at over five times its historical rate in order to keep the interest rate in 2050 above its 2017 level. Finally, a modest increase in the relative supply of high-skilled workers would help to prevent the predicted major rise in wage inequality due to automation, while also increasing economic growth.

Appendix For Online Publication

Figure A1: UN and Model Population Projections (Continued)

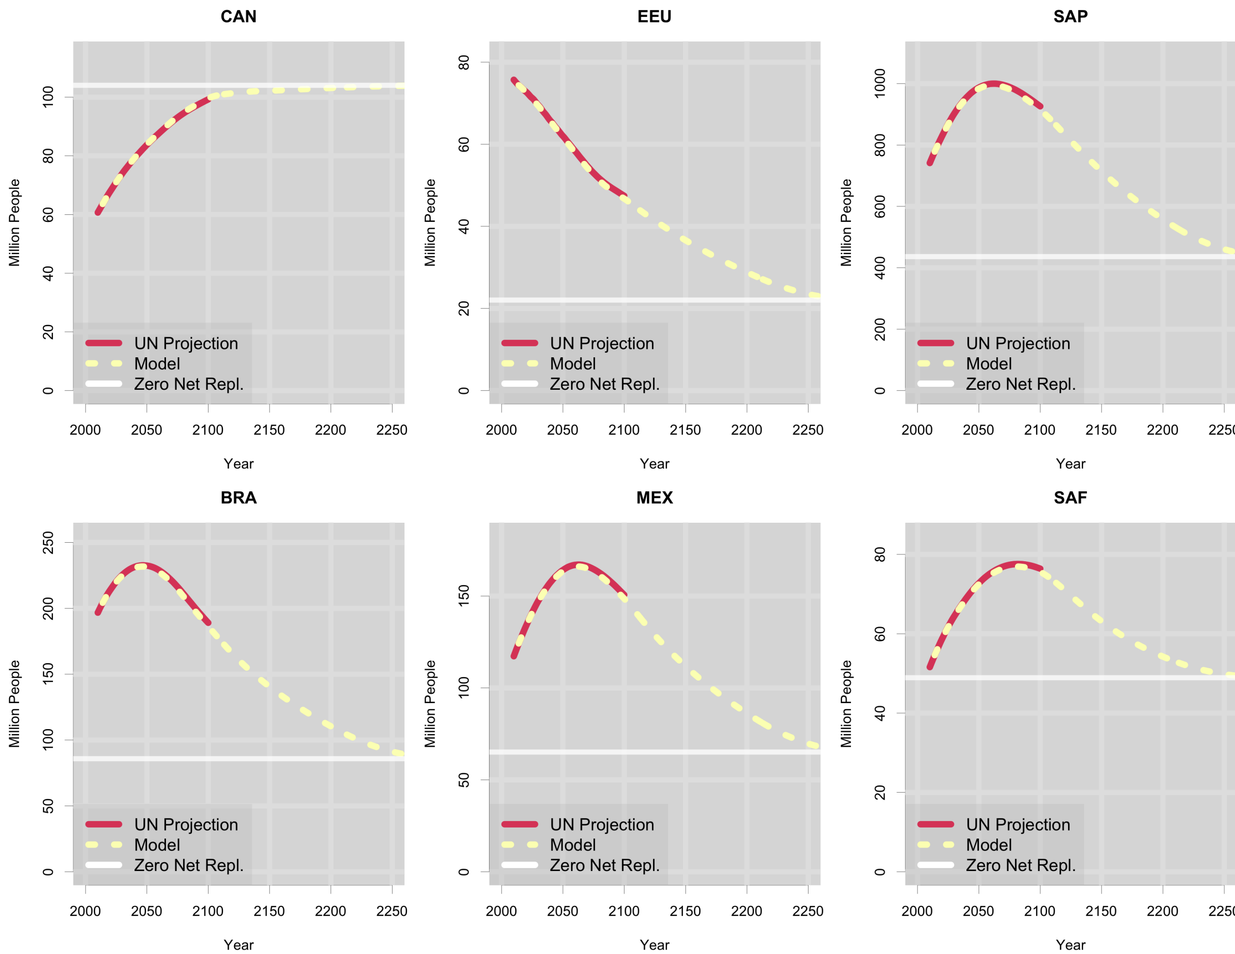

Figure A2: UN and Model Population Projections (Continued)

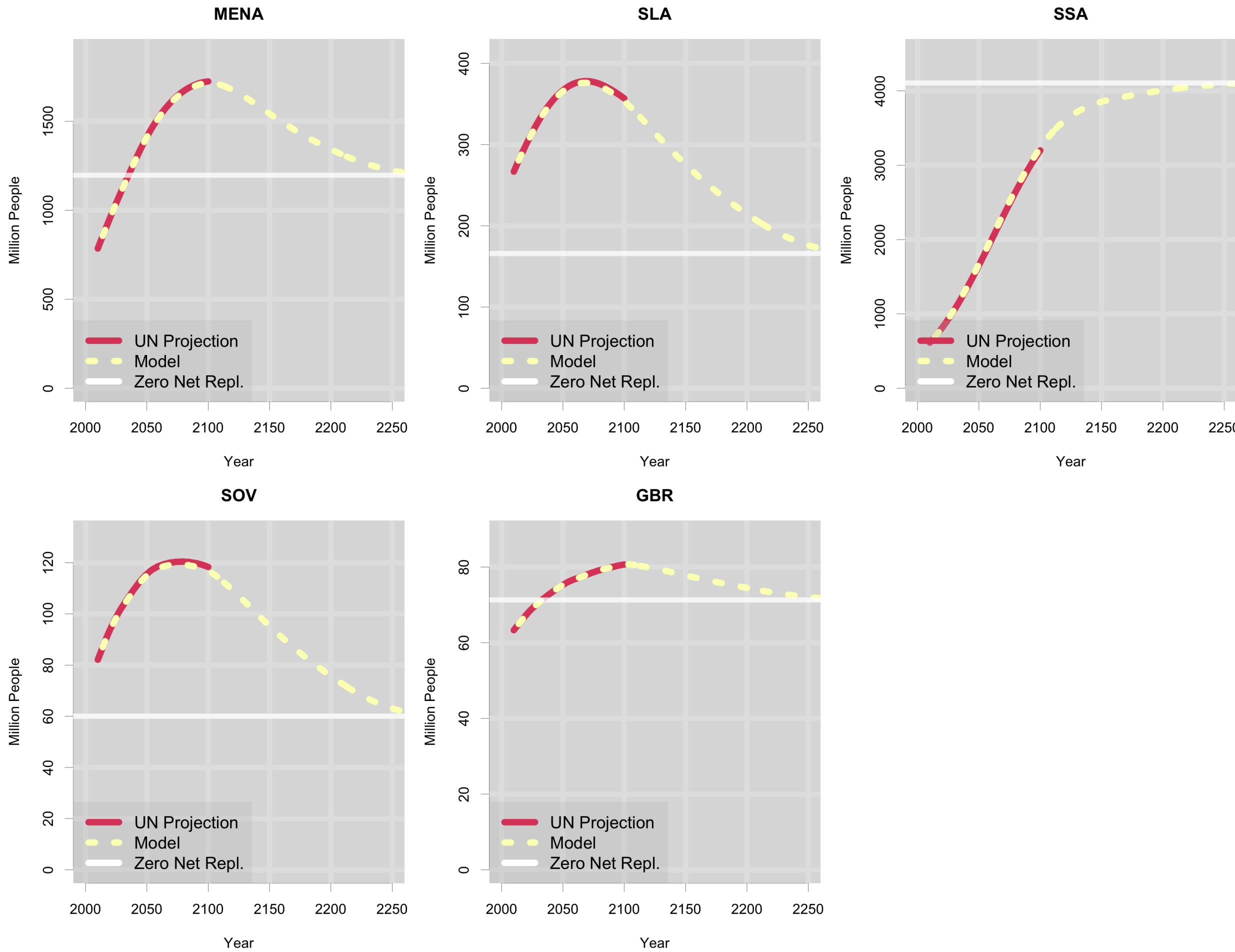

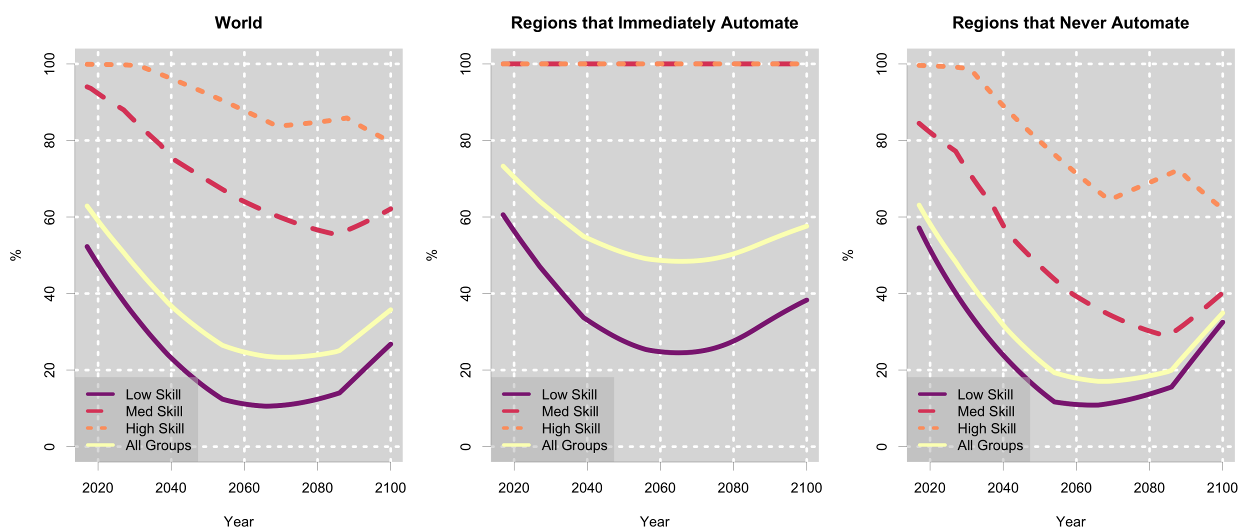

Figure A3: Percent of Adult Population Made Better Off In Terms of Lifetime Welfare, Relative to No Automation

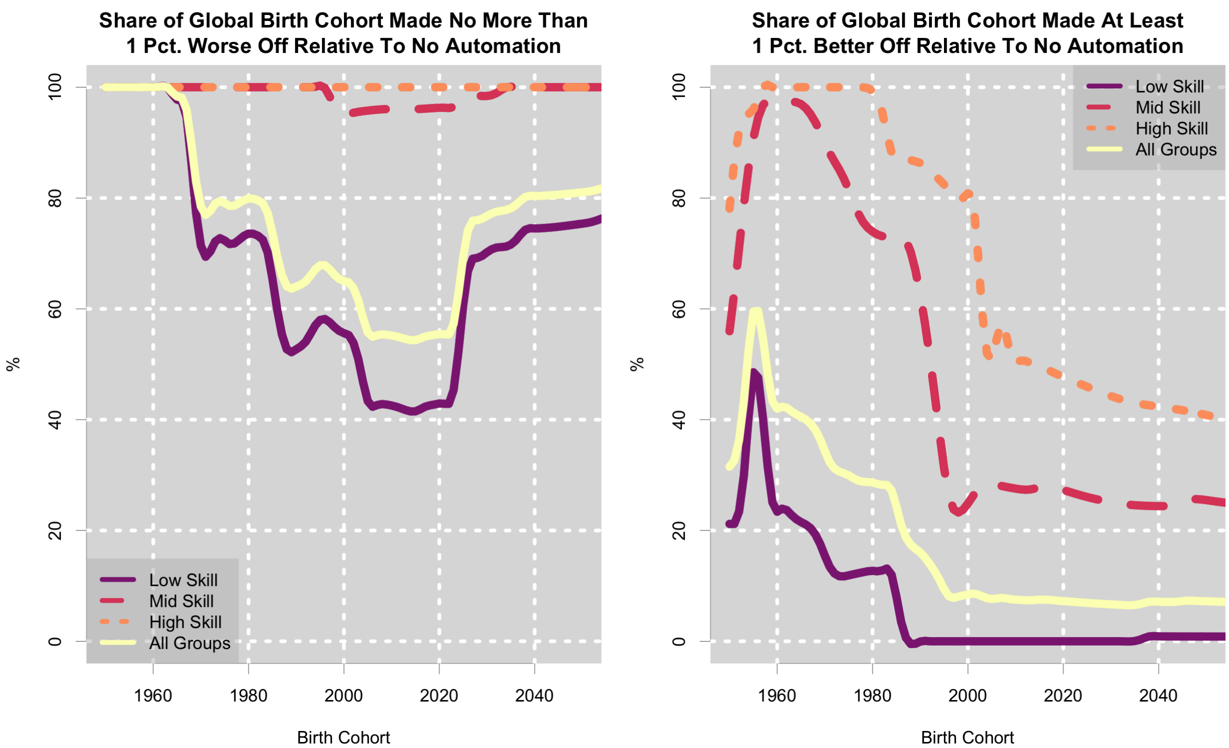

Figure A4: Global Lifetime Welfare Consequence of Baseline Automation By Birth Cohort

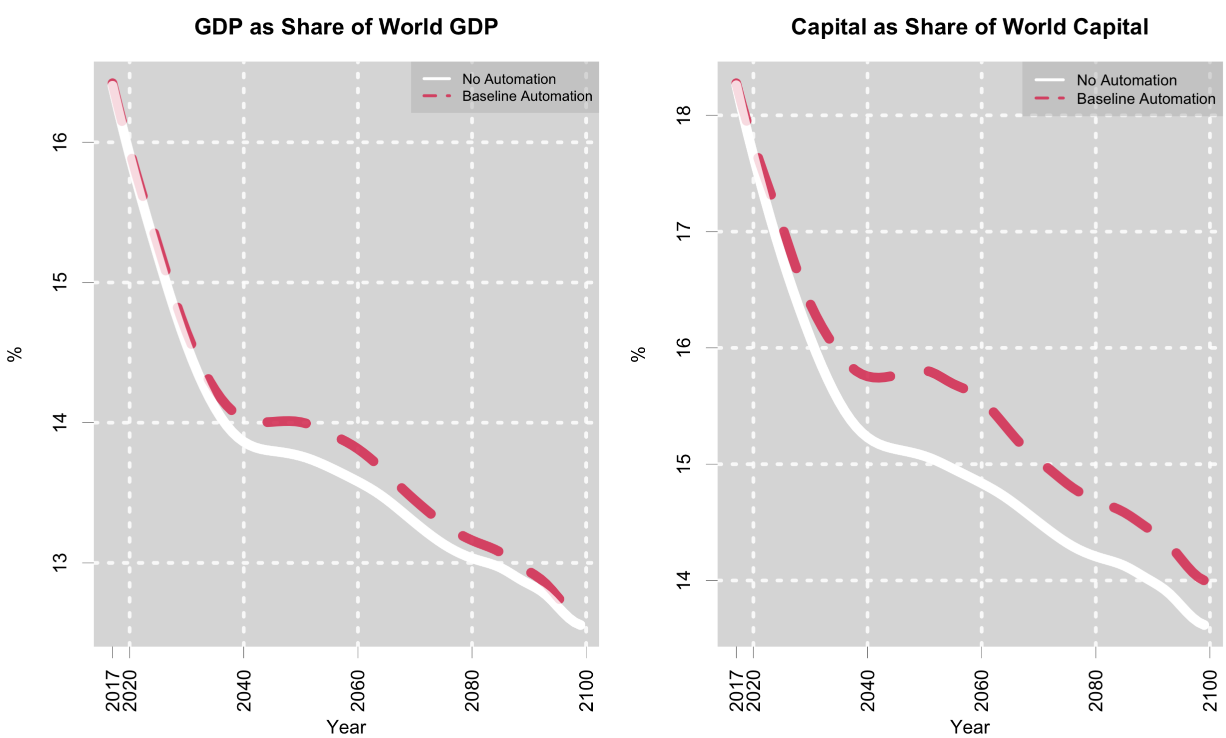

Figure A5: U.S. GDP and Capital as Share of World With and Without Baseline Automation

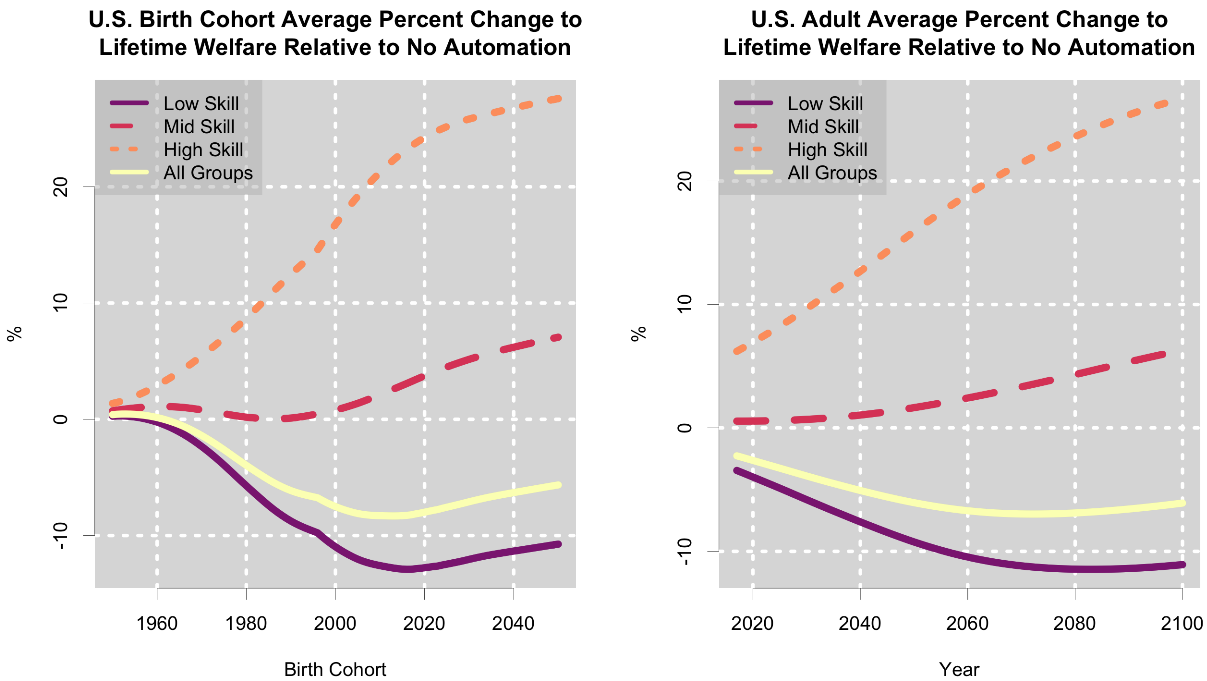

Figure A6: U.S. Lifetime Welfare Consequence of Baseline Automation By Year and Birth Cohort

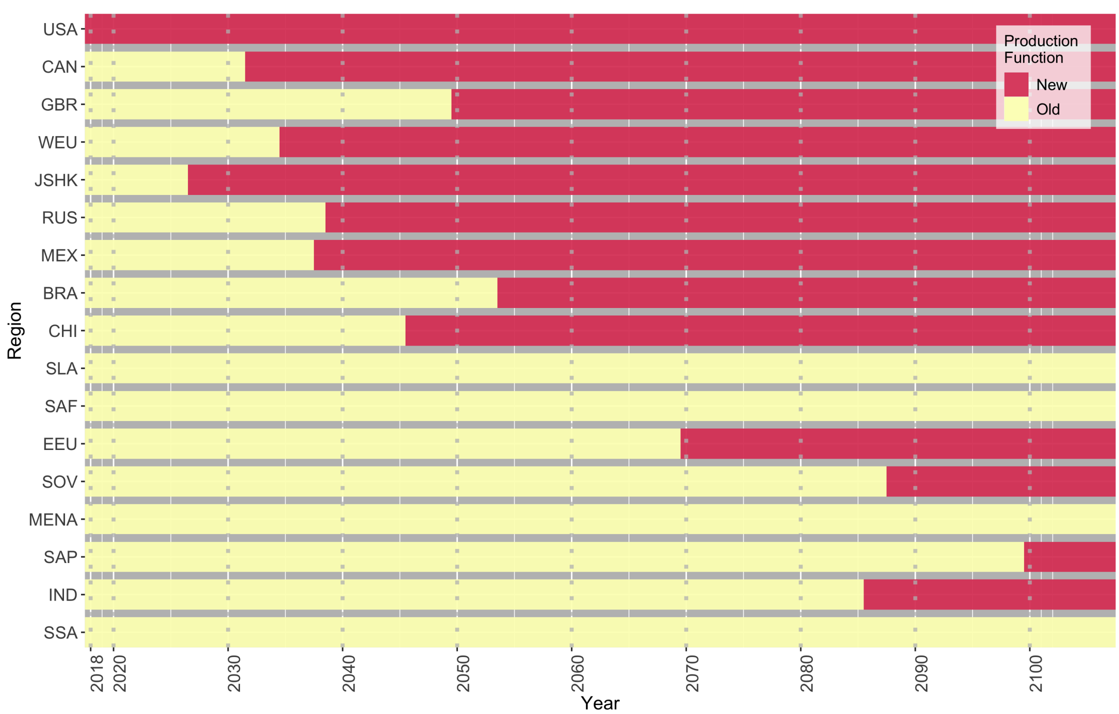

Figure A7: Choice Of Production Function By Region, Automation Without Rising Inequality

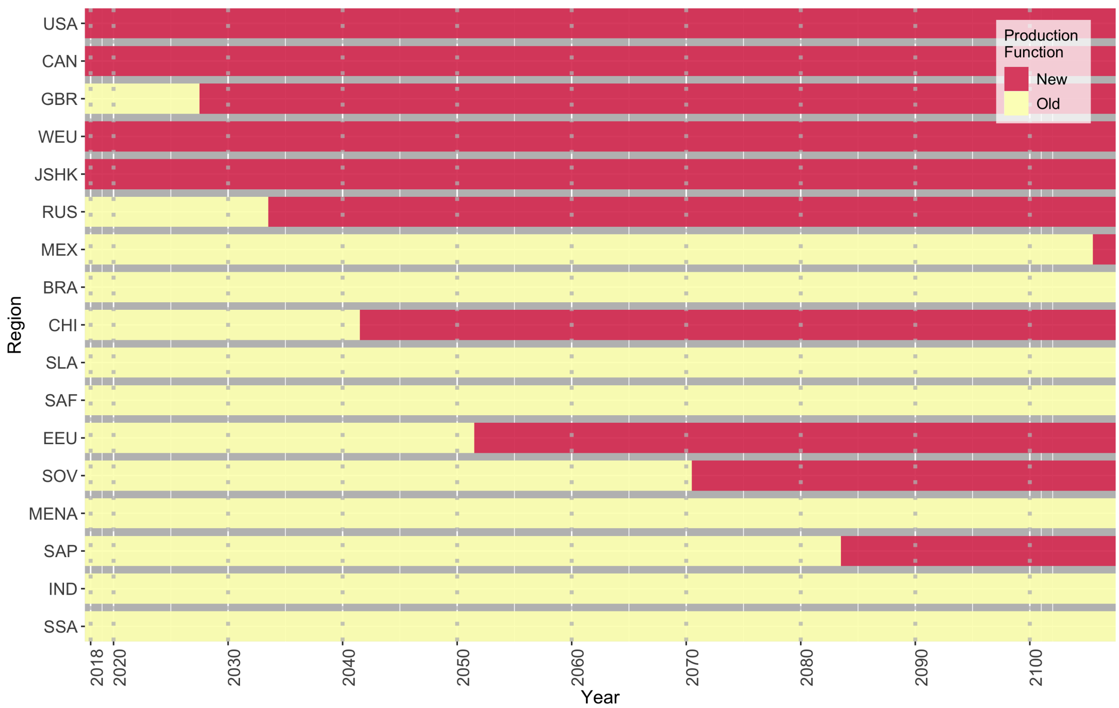

Figure A8: Choice Of Production Function, Automation With Five Times the Baseline Rate of Capital Share Increase

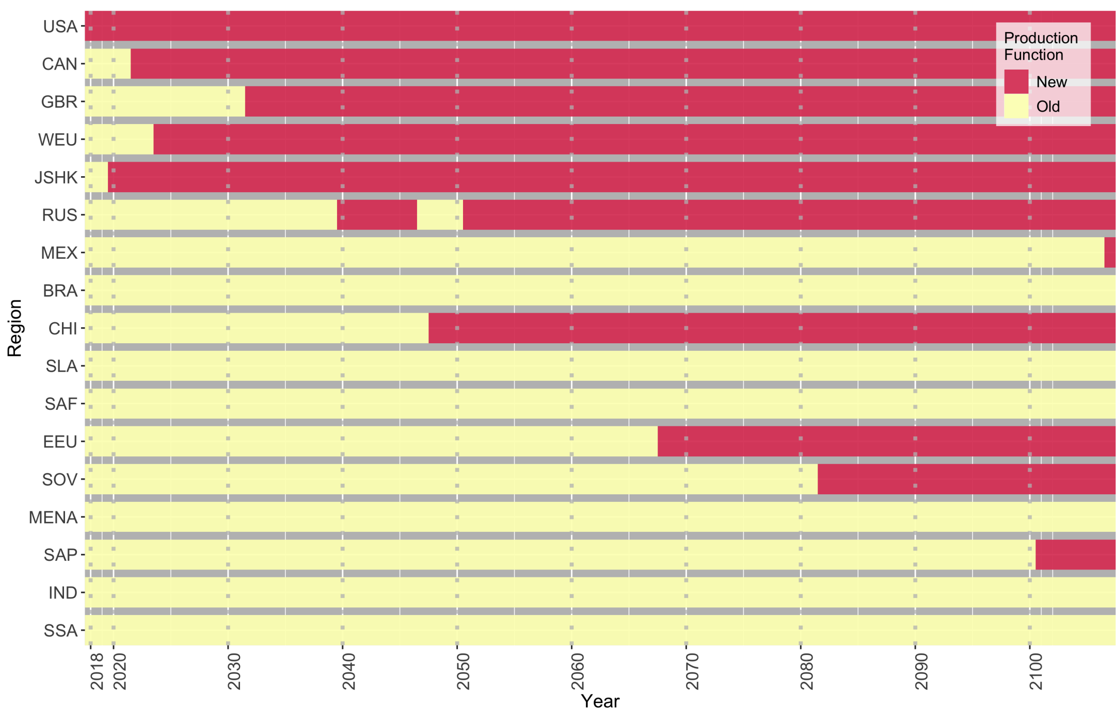

Figure A9: Choice Of Production Function, Automation With Ten Times the Baseline Rate of Capital Share Increase

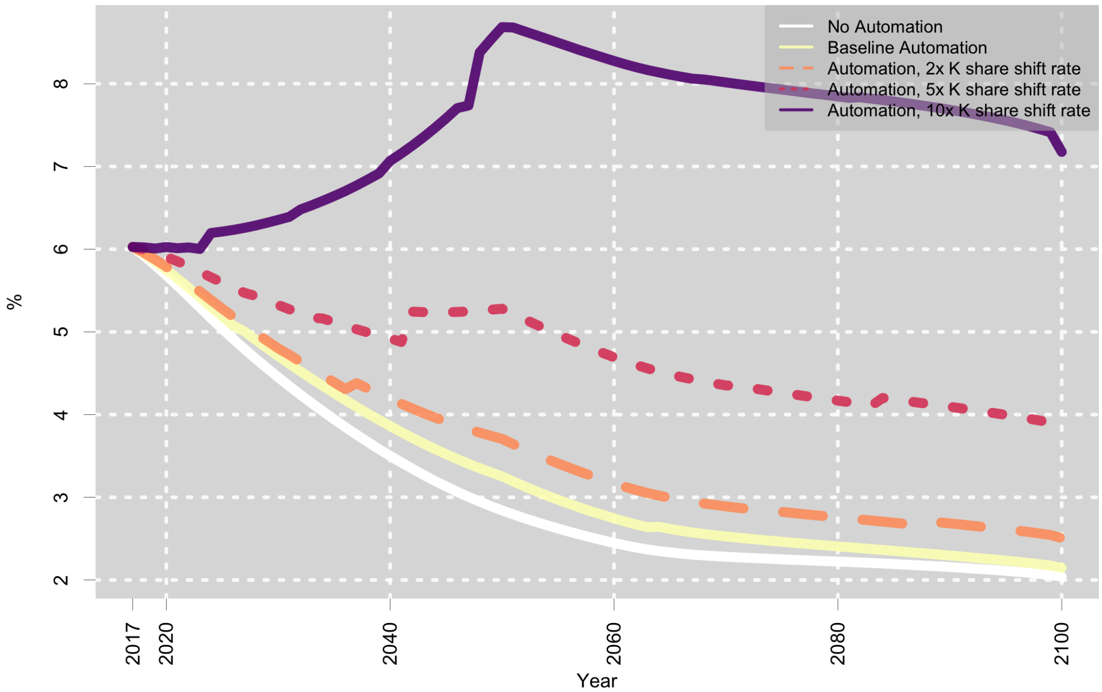

Figure A10: World Interest Rate By Automation Rate Scenario

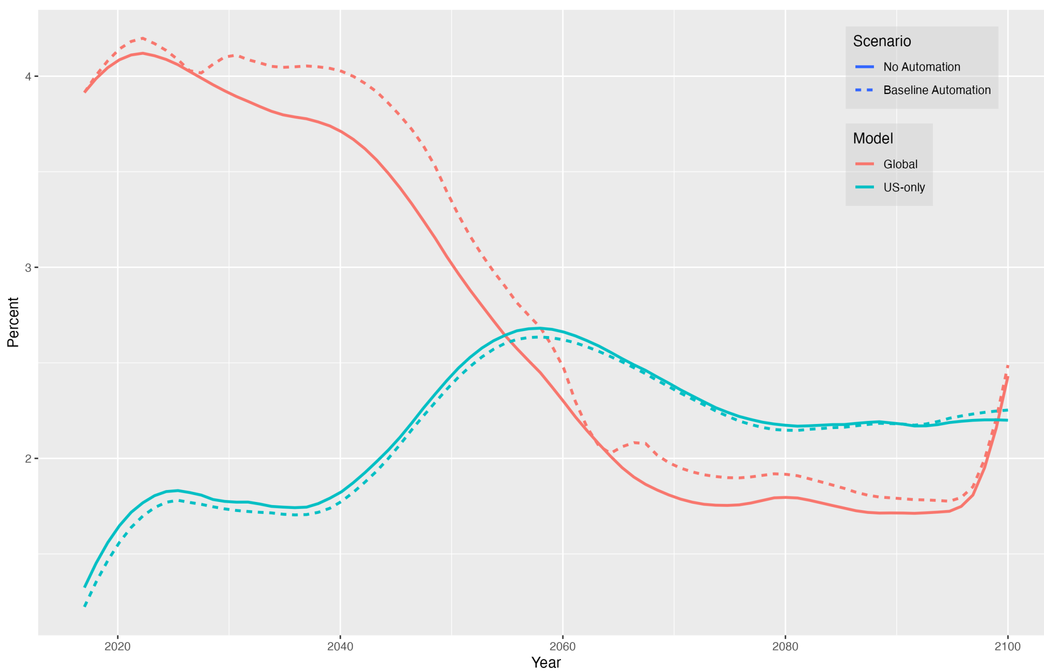

Figure A11: Annual Capital Stock Growth Under Global and U.S-only Models

Region	Countries
BRA	Brazil
CAN	Canada, Australia, New Zealand
CHI	China
EEU	Belarus, Republic of Moldova, Ukraine, Albania, Bosnia and Herzegovina, Montenegro, Serbia
GBR	United Kingdom
IND	India
JSHK	China, Hong Kong SAR, Japan, Republic of Korea, Singapore
MENA	Ethiopia, Algeria, Egypt, Libya, Morocco, Sudan, Tunisia, Western Sahara, Mali, Afghanistan, Iran (Islamic Republic of), Pakistan, Bahrain, Iraq, Jordan, Kuwait, Lebanon, Oman, Qatar, Saudi Arabia, Syrian Arab Republic, Turkey, United Arab Emirates, Yemen
MEX	Mexico
RUS	Russian Federation
SAF	South Africa
SAP	Bangladesh, Nepal, Sri Lanka, Brunei Darussalam, Cambodia, Indonesia, Lao People's Democratic Republic, Malaysia, Philippines, Thailand, Timor-Leste, Viet Nam
SLA	Sao Tome and Principe, Antigua and Barbuda, Bahamas, Barbados, Dominican Republic, Grenada, Haiti, Jamaica, Saint Lucia, Saint Vincent and the Grenadines, Trinidad and Tobago, Belize, Costa Rica, El Salvador, Guatemala, Honduras, Nicaragua, Panama, Argentina, Bolivia (Plurinational State of), Chile, Colombia, Ecuador, Guyana, Paraguay, Peru, Suriname, Uruguay, Venezuela (Bolivarian Republic of)
SOV	Mongolia, Kazakhstan, Kyrgyzstan, Tajikistan, Turkmenistan, Uzbekistan, Armenia, Azerbaijan, Georgia
SSA	Eritrea, Kenya, Madagascar, Mozambique, Rwanda, South Sudan, Uganda, United Republic of Tanzania, Zambia, Zimbabwe, Angola, Cameroon, Central African Republic, Congo, Democratic Republic of the Congo, Equatorial Guinea, Gabon, Botswana, Lesotho, Namibia, Swaziland, Burkina Faso, Côte d'Ivoire, Gambia, Ghana, Liberia, Niger, Nigeria, Senegal, Sierra Leone, Togo, Tonga
USA	United States of America
WEU	Cyprus, Israel, Bulgaria, Czechia, Hungary, Poland, Romania, Slovakia, Channel Islands, Denmark, Estonia, Finland, Iceland, Ireland, Latvia, Lithuania, Norway, Sweden, Croatia, Greece, Italy, Malta, Portugal, Slovenia, Spain, Austria, Belgium, France, Germany, Luxembourg, Netherlands, Switzerland

Table A1: Full List of Countries in Each Region

	2017	2050	2100
USA	16.4	13.8	12.5
WEU	17.1	11.9	10.1
JSHK	6.9	6.5	6.2
CHI	16.6	22.1	25.1
IND	7.0	9.6	11.4
RUS	3.2	1.8	1.4
CAN	2.7	2.4	2.2
EEU	0.8	0.7	0.7
SAP	6.6	7.5	6.9
BRA	2.5	2.2	1.5
MEX	2.1	2.0	1.3
SAF	0.6	0.5	0.4
MENA	8.7	9.5	10.1
SLA	3.2	3.1	2.4
SSA	2.2	3.3	4.7
SOV	0.9	1.2	1.4
GBR	2.6	1.9	1.6

Table A2: Share of GDP As Percent of Global Output, No Automation

	2017	2050	2100
USA	1.00	1.94	3.96
WEU	0.93	1.79	4.16
JSHK	0.72	2.05	6.43
CHI	0.23	0.89	3.44
IND	0.10	0.32	1.05
RUS	0.44	0.76	1.58
CAN	0.81	1.55	3.15
EEU	0.15	0.46	1.54
SAP	0.16	0.42	1.05
BRA	0.24	0.53	1.13
MEX	0.32	0.67	1.22
SAF	0.21	0.41	0.76
MENA	0.19	0.37	0.83
SLA	0.22	0.46	0.95
SSA	0.06	0.11	0.21
SOV	0.20	0.56	1.69
GBR	0.76	1.36	2.78

Table A3: GDP Per Capita Relative to 2017 U.S. Level, No Automation

	Asset Share (%)			Capital Share (%)		
	2017	2050	2100	2017	2050	2100
USA	16.8	6.1	5.4	18.2	15.1	13.6
WEU	16.1	10.4	7.9	16.7	10.9	9.0
JSHK	9.3	5.9	4.9	8.4	7.4	6.8
CHI	17.4	35.7	34.9	18.8	25.7	29.3
IND	6.6	11.7	13.1	6.2	9.0	10.9
RUS	4.1	2.4	1.5	3.8	2.4	2.0
CAN	2.7	1.4	1.2	2.8	2.5	2.3
EEU	0.7	0.1	0.3	0.8	0.7	0.7
SAP	5.5	7.9	8.8	6.1	7.1	6.5
BRA	3.0	2.3	1.9	2.5	2.4	1.7
MEX	1.9	1.5	1.3	2.2	2.2	1.4
SAF	0.6	0.6	0.5	0.5	0.5	0.4
MENA	9.0	9.0	10.2	6.4	7.1	7.6
SLA	2.0	1.8	2.2	2.1	1.9	1.5
SSA	1.9	1.7	3.8	1.8	2.8	4.1
SOV	0.6	0.8	1.4	0.7	0.9	1.1
GBR	1.9	0.5	0.8	2.0	1.3	1.1

Table A4: Shares of Assets and Capital as Percent of Global Assets and Capital, No Automation

	2017	2050	2100
Avg. Pension Tax Rate (%)	13.5	16.3	26.8
Pension Benefits as Share of GDP (%)	5.0	6.0	10.0
Consumption Tax Rate (%)	9.0	9.9	10.8
Avg. Income Tax Rate (%)	24.4	27.4	28.6

Table A5: U.S. Endogenous Tax Rates and Pension Benefits, No Automation

	No Automation	Baseline Automation	Automation W/O Rising Inequality
K	0.341	0.369	0.369
L High	0.220	0.264	0.210
L Med	0.220	0.205	0.210
L Low	0.220	0.161	0.210

USA

	No Automation	Baseline Automation	Automation W/O Rising Inequality
K	0.368	0.395	0.395
L High	0.211	0.254	0.202
L Med	0.211	0.197	0.202
L Low	0.211	0.154	0.202

WEU

	No Automation	Baseline Automation	Automation W/O Rising Inequality
K	0.487	0.509	0.509
L High	0.171	0.206	0.164
L Med	0.171	0.160	0.164
L Low	0.171	0.125	0.164

JSHK

	No Automation	Baseline Automation	Automation W/O Rising Inequality
K	0.383	0.409	0.409
L High	0.206	0.248	0.197
L Med	0.206	0.192	0.197
L Low	0.206	0.151	0.197

CHI

Table A6: 2050 Output Shares of Capital and Labor By Skill Group

	No Automation	Baseline Automation	Automation W/O Rising Inequality
K	0.267	0.298	0.298
L High	0.244	0.294	0.234
L Med	0.244	0.228	0.234
L Low	0.244	0.179	0.234

IND

	No Automation	Baseline Automation	Automation W/O Rising Inequality
K	0.405	0.431	0.431
L High	0.198	0.239	0.190
L Med	0.198	0.185	0.190
L Low	0.198	0.145	0.190

RUS

	No Automation	Baseline Automation	Automation W/O Rising Inequality
K	0.388	0.414	0.414
L High	0.204	0.246	0.195
L Med	0.204	0.191	0.195
L Low	0.204	0.150	0.195

CAN

	No Automation	Baseline Automation	Automation W/O Rising Inequality
K	0.416	0.440	0.440
L High	0.195	0.235	0.187
L Med	0.195	0.182	0.187
L Low	0.195	0.143	0.187

EEU

Table A7: 2050 Output Shares of Capital and Labor By Skill Group, Continued

	No Automation	Baseline Automation	Automation W/O Rising Inequality
K	0.326	0.355	0.355
L High	0.225	0.270	0.215
L Med	0.225	0.210	0.215
L Low	0.225	0.165	0.215

SAP

	No Automation	Baseline Automation	Automation W/O Rising Inequality
K	0.288	0.318	0.318
L High	0.237	0.286	0.227
L Med	0.237	0.222	0.227
L Low	0.237	0.174	0.227

BRA

	No Automation	Baseline Automation	Automation W/O Rising Inequality
K	0.365	0.392	0.392
L High	0.212	0.255	0.203
L Med	0.212	0.198	0.203
L Low	0.212	0.155	0.203

MEX

	No Automation	Baseline Automation	Automation W/O Rising Inequality
K	0.292	0.322	0.322
L High	0.236	0.284	0.226
L Med	0.236	0.221	0.226
L Low	0.236	0.173	0.226

SAF

Table A8: 2050 Output Shares of Capital and Labor By Skill Group, Continued

	No Automation	Baseline Automation	Automation W/O Rising Inequality
K	0.335	0.363	0.363
L High	0.222	0.267	0.212
L Med	0.222	0.207	0.212
L Low	0.222	0.163	0.212

MENA

	No Automation	Baseline Automation	Automation W/O Rising Inequality
K	0.265	0.296	0.296
L High	0.245	0.295	0.235
L Med	0.245	0.229	0.235
L Low	0.245	0.180	0.235

SLA

	No Automation	Baseline Automation	Automation W/O Rising Inequality
K	0.281	0.311	0.311
L High	0.240	0.289	0.230
L Med	0.240	0.224	0.230
L Low	0.240	0.176	0.230

SSA

	No Automation	Baseline Automation	Automation W/O Rising Inequality
K	0.356	0.383	0.383
L High	0.215	0.259	0.206
L Med	0.215	0.201	0.206
L Low	0.215	0.158	0.206

SOV

Table A9: 2050 Output Shares of Capital and Labor By Skill Group, Continued

	No Automation	Baseline Automation	Automation W/O Rising Inequality
K	0.335	0.363	0.363
L High	0.222	0.267	0.212
L Med	0.222	0.207	0.212
L Low	0.222	0.163	0.212

GBR

Table A10: 2050 Output Shares of Capital and Labor By Skill Group, Continued

	2050			2100		
	No Automation	Baseline Automation	% Diff	No Automation	Baseline Automation	% Diff
USA	45.3	47.7	5.1	107.7	116.3	8.0
WEU	39.3	44.3	12.7	86.6	101.1	16.8
JSHK	21.4	24.2	12.9	53.4	63.6	19.1
CHI	72.7	75.3	3.6	216.5	240.4	11.0
IND	31.7	30.8	-2.9	98.2	97.4	-0.8
RUS	6.0	6.2	2.4	12.0	12.9	8.1
CAN	7.8	8.8	12.1	19.1	22.0	15.5
EEU	2.4	2.5	6.6	6.2	7.0	12.3
SAP	24.7	24.0	-3.0	59.6	62.6	5.0
BRA	7.4	7.1	-3.3	13.1	12.9	-1.0
MEX	6.6	6.4	-3.6	11.2	11.1	-1.1
SAF	1.8	1.7	-3.6	3.5	3.5	-1.1
MENA	31.3	30.6	-2.3	87.4	86.8	-0.7
SLA	10.2	10.0	-2.0	20.7	20.5	-0.6
SSA	10.8	10.5	-2.4	40.4	40.1	-0.8
SOV	3.8	3.7	-2.7	12.1	12.7	5.4
GBR	6.2	6.8	10.0	13.7	15.5	13.2
World	329.5	340.6	3.4	861.3	926.6	7.6

Table A11: GDP in Trillions of 2017 Dollars

	GDP in 2017 (Trillion Dollars)		Capital Stock in 2017 (Trillion Dollars)		Net Intl. Investment Position in 2017 (Trillion Dollars)		GDP in 2100 (Trillion 2017 Dollars)	
	Target	Model	Target	Model	Target	Model	Target	Model
USA	19.49	19.49	59.44	59.44	-6.71	-6.70	112.37	108.03
WEU	20.32	20.35	54.27	54.12	1.87	2.18	88.95	86.67
JSHK	8.24	8.25	27.27	27.22	7.02	7.16	54.97	53.44
CHI	19.76	19.75	62.05	61.82	3.85	3.89	226.04	215.64
IND	8.26	8.27	20.40	20.33	-1.22	-1.25	102.12	97.59
RUS	3.81	3.81	13.90	13.88	0.76	0.83	12.59	12.02
CAN	3.20	3.19	9.18	9.10	-0.25	-0.22	19.90	19.12
EEU	0.92	0.92	2.74	2.73	-0.43	-0.41	6.45	6.18
SAP	7.82	7.83	20.34	20.27	-1.36	-1.32	61.89	59.58
BRA	3.01	3.01	8.19	8.17	-0.83	-0.79	13.77	13.08
MEX	2.47	2.46	7.47	7.44	-1.04	-1.02	11.66	11.26
SAF	0.72	0.72	1.76	1.75	0.08	0.08	3.70	3.51
MENA	10.28	10.29	24.37	24.29	-0.37	-0.38	89.80	87.35
SLA	3.82	3.82	6.88	6.84	-0.74	-0.73	21.56	20.62
SSA	2.59	2.59	6.16	6.17	-0.80	-0.78	41.11	40.62
SOV	1.06	1.06	2.50	2.49	-0.34	-0.34	12.72	12.05
GBR	3.03	3.03	6.42	6.40	-0.39	-0.37	14.07	13.71

Table A12: Calibration Targets: Macro Aggregates

	Private Consumption (% of GDP)		Government Consumption (% of GDP)		Health Expenditure (% of GDP)		Education Expenditure (% of GDP)	
	Target	Model	Target	Model	Target	Model	Target	Model
USA	68.20	68.08	14.13	14.08	9.28	9.28	6.05	5.99
WEU	53.36	53.43	20.32	20.21	6.75	6.72	4.73	4.77
JSHK	53.25	53.40	17.75	17.96	6.38	6.45	3.46	3.49
CHI	39.00	39.06	14.62	14.53	2.98	2.96	3.73	3.71
IND	58.90	58.92	10.76	10.70	1.00	1.00	3.80	3.79
RUS	52.87	52.93	18.22	18.17	3.02	3.02	3.48	3.48
CAN	57.28	57.28	19.65	19.59	7.36	7.36	5.32	5.32
EEU	67.28	67.14	18.07	18.06	3.87	3.87	5.57	5.56
SAP	59.06	59.03	10.76	10.72	1.33	1.33	2.88	2.88
BRA	64.48	64.04	20.17	20.24	3.80	3.82	6.00	6.02
MEX	65.27	65.41	11.63	11.64	2.85	2.84	4.91	4.90
SAF	59.43	59.48	20.93	20.85	4.32	4.31	5.59	5.58
MENA	55.46	55.35	15.88	15.99	3.94	3.97	3.16	3.19
SLA	66.76	66.77	14.91	14.91	4.75	4.74	4.94	4.93
SSA	70.16	70.17	10.84	10.82	1.36	1.36	4.09	4.08
SOV	57.62	57.37	11.85	11.80	2.27	2.26	3.57	3.56
GBR	64.50	64.42	18.72	18.87	7.48	7.52	4.95	4.99

Table A13: Calibration Targets: Consumption and Expenditure

	Pension Benefits in 2017 (% of GDP)		Pension Benefits in 2050 (% of 2050 GDP)		Gov. G&S Purchases Excl. Health & Educ. (% of GDP)		Welfare Expenditure Excl. Pensions (% of GDP)		Net Payment on Debt/Assets (% of GDP)	
	Target	Model	Target	Model	Target	Model	Target	Model	Target	Model
USA	5.00	5.00	6.02	6.01	11.05	11.03	2.66	2.66	3.94	4.03
WEU	10.05	10.05	10.14	10.04	13.64	13.55	8.86	8.82	1.84	1.89
JSHK	8.04	8.03	7.11	7.06	10.01	10.11	5.47	5.52	1.50	1.53
CHI	4.90	4.90	11.35	11.37	14.62	14.53	3.59	3.57	0.86	0.88
IND	4.30	4.30	4.30	4.31	7.34	7.30	1.56	1.56	5.00	5.11
RUS	8.75	8.73	11.89	11.72	16.05	16.03	4.41	4.41	0.92	0.94
CAN	4.02	4.01	4.64	4.64	13.52	13.51	5.80	5.80	2.14	2.20
EEU	6.91	6.91	9.43	9.43	13.10	13.07	7.59	7.59	2.98	3.04
SAP	0.43	0.43	0.58	0.57	10.71	10.68	1.27	1.27	1.39	1.43
BRA	12.00	12.01	17.81	18.00	10.37	10.40	7.12	7.15	8.95	9.16
MEX	2.10	2.10	3.50	3.52	7.31	7.30	5.40	5.40	2.93	3.00
SAF	1.51	1.51	2.27	2.28	20.93	20.85	5.73	5.72	3.94	4.02
MENA	4.63	4.62	5.87	5.82	15.88	15.99	3.83	3.86	1.35	1.39
SLA	2.12	2.12	2.74	2.75	14.08	14.04	3.45	3.44	0.96	0.98
SSA	2.10	2.10	3.76	3.75	10.84	10.82	7.41	7.42	2.54	2.60
SOV	4.02	4.02	5.73	5.75	11.85	11.80	2.77	2.76	0.82	0.84
GBR	8.33	8.34	8.84	8.83	10.94	10.99	6.78	6.82	2.69	2.76

Table A14: Calibration Targets: Consumption and Expenditure (Continued)

	Corporate Income Tax Revenue (% of GDP)		Natural Resource Rents (% of GDP)		Contributions to Pension System (% of GDP)		Income Tax Revenue (% of GDP)		Consump. Tax Revenue (% of GDP)	
	Target	Model	Target	Model	Target	Model	Target	Model	Target	Model
USA	1.98	1.98	0.20	0.20	6.68	6.67	17.95	21.15	5.21	6.12
WEU	3.65	3.66	0.10	0.10	13.61	13.61	14.06	14.37	12.99	13.26
JSHK	6.34	6.35	0.00	0.00	10.37	10.35	8.89	9.25	7.90	8.21
CHI	5.51	5.52	1.20	1.20	5.20	5.18	3.41	3.82	12.49	13.97
IND	3.30	3.31	1.60	1.60	1.90	1.90	2.50	3.63	5.60	8.01
RUS	4.11	4.13	7.30	7.29	7.04	7.05	5.62	5.83	11.74	12.18
CAN	6.12	6.13	0.50	0.50	2.37	2.37	20.23	20.04	8.09	8.01
EEU	3.40	3.42	0.10	0.10	8.16	8.14	7.23	6.87	20.69	19.60
SAP	4.59	4.60	1.20	1.19	0.31	0.31	1.81	1.93	8.36	8.91
BRA	3.98	3.99	1.40	1.39	10.85	10.86	7.19	8.00	16.55	18.06
MEX	4.29	4.30	2.80	2.74	2.22	2.22	5.32	6.10	7.12	8.13
SAF	6.92	6.94	1.60	1.60	0.56	0.56	14.27	15.60	13.80	15.05
MENA	1.76	1.76	9.60	9.73	4.36	4.35	2.78	3.35	10.71	12.90
SLA	3.90	3.91	0.80	0.80	3.11	3.12	5.52	7.32	10.93	14.50
SSA	3.09	3.11	2.80	2.80	2.20	2.20	4.45	6.02	9.16	12.38
SOV	3.96	3.97	6.80	6.81	1.85	1.84	3.20	3.34	8.23	8.59
GBR	3.90	3.91	1.00	1.06	7.79	7.79	13.67	14.63	12.07	12.89

Table A15: Calibration Targets: Revenue[31]

[31]Note that income and consumption tax revenue are set to balance the budget and therefore may deviate from each region's reported revenue level.

A Calibration Data and Assumptions

We first consider the Cobb-Douglas coefficients of the production function. In each region, we define skill groups such that $\beta_{t,k} = (1 - \alpha_t)/3$, $k \in (1, 2, 3)$, i.e., that each skill group, in 2017, collects a third of of total labor income. In the absence of automation, αs and βs are fixed through time. The population share of each group is estimated with the World Inequality Database (Alvaredo et al., 2020) by interpolating country-specific income shares by population quintile (or decile, if data is available). Shares are summarized in table A16.

	USA	WEU	JSHK	CHI	IND	RUS	CAN	EEU	SAP	BRA	MEX	SAF	MENA	SLA	SSA	SOV	GBR
Capital Income Share																	
	0.341	0.368	0.487	0.383	0.267	0.405	0.388	0.416	0.326	0.288	0.365	0.292	0.335	0.265	0.281	0.356	0.335
No Automation Labor Income Share Per Skill Group																	
	0.220	0.211	0.171	0.206	0.244	0.198	0.204	0.195	0.225	0.237	0.212	0.236	0.222	0.245	0.240	0.215	0.222
Share of Population by Skill Group																	
High	0.040	0.109	0.109	0.060	0.030	0.050	0.109	0.093	0.052	0.020	0.020	0.040	0.040	0.020	0.048	0.050	0.090
Med	0.200	0.255	0.255	0.210	0.170	0.230	0.255	0.226	0.193	0.180	0.180	0.166	0.166	0.180	0.170	0.230	0.250
Low	0.760	0.636	0.636	0.730	0.800	0.720	0.636	0.681	0.755	0.800	0.800	0.795	0.795	0.800	0.782	0.720	0.660

Table A16: Input Shares by Region and Population Shares by Skill Group and Region

Capital income shares are calibrated such that, in 2017, each region's capital demand matches its capital stock as calculated from the IMF Investment and Capital Stock Database (ICSD; IMF 2019a). To achieve this, we scale the overall model so that one unit of model output is equal to 100 billion real U.S. dollars in 2017 value. The total global stock of capital is set such that the amount of global capital in a given period is equal to the period's global supply of net private assets (i.e. assets net of private debt), less total global non-capital assets. These are the sum of world-wide government debt and the present value of non-government owned fossil fuel rent flows, each of which have the same rate of return as other assets. The depreciation rate in each region is calibrated to the average difference between its net capital accumulation rate, from 2012 to 2017, and gross domestic investment rate for the same period. Both estimates are obtained from ICSD data; we assume that depreciation rates are fixed through time.

The total factor productivity term, ϕ_t, is used to calibrate output in each region. We calibrate ϕ_t in 2017 to each region's GDP in the same year, adjusting the overall scaling factor to express the model's output in real dollar units. The path of ϕ_t is calibrated to long-term economic growth projections from Müller et al. (2019). The authors use a Bayesian model to provide two sets of estimations for GDPs through 2100, one estimating growth paths of countries jointly in a multi-variate approach, and the other estimating each country individually. In the no automation case, we adopt a simple average of the two output paths as the calibration target, and adjust ϕ_t in 2100 to match output in each region.[32] ϕ_t is assumed to growth linearly between 2017 and 2100. After 2100, TFP in all regions grow at a fixed rate of 1 percent per year. TFP values by region in 2017, 2050, and 2100 are summarized in table A17.

	2017	2050 W/O 1 Pct Tech. Growth	2100 W/O 1 Pct Tech. Growth	2050 With 1 Pct Tech. Growth	2100 With 1 Pct Tech. Growth
USA	1.00	1.15	1.37	1.60	3.13
WEU	0.69	0.83	1.05	1.15	2.40
JSHK	0.62	0.84	1.16	1.17	2.65
CHI	0.34	0.60	1.00	0.83	2.28
IND	0.29	0.43	0.65	0.60	1.48
RUS	0.47	0.53	0.61	0.74	1.39
CAN	0.68	0.78	0.92	1.08	2.10
EEU	0.34	0.49	0.72	0.68	1.64
SAP	0.31	0.42	0.59	0.58	1.35
BRA	0.46	0.56	0.72	0.78	1.64
MEX	0.61	0.66	0.73	0.92	1.67
SAF	0.44	0.45	0.46	0.62	1.05
MENA	0.43	0.48	0.54	0.67	1.23
SLA	0.49	0.58	0.72	0.81	1.64
SSA	0.18	0.18	0.18	0.25	0.41
SOV	0.37	0.53	0.77	0.74	1.76
GBR	0.76	0.87	1.05	1.21	2.40

Table A17: Calibrated TFP By Region and Year

[32] While we assume 1 percent long-term technological growth in the model, individual regions are allowed to grow at a rate of less than 1 percent in the short run if required to match projections.

The initial age- and skill-distribution of assets is calibrated to the distribution of total household assets in the 2016 Survey of Consumer Finances (SCF; Bricker et al. 2017). For the age profile, the age-asset distribution is spline smoothed and weighted so that, depending on the region-specific age of labor force entry, agents in their first working year own zero assets. For the skill-group profile, we sort SCF households by annual wage income and calculate the fraction of assets owned by each bin corresponding to the population shares of the U.S's skill groups. Because of the lack of large-sample household-level data for non-U.S. regions, we apply the U.S's age-skill-group asset profile to all regions.[33]

The total amount of initial assets in each region is calibrated such that the model's net foreign asset stock, as a fraction of GDP, matches data from the International Investment Position (IIP) Database (IMF, 2019b). We measure a region's net foreign asset position as the difference between its net national wealth and its stock of capital. The total stock of global assets in 2017 is set to generate an initial world interest rate of approximately 6 percent.

Government consumption, taxes, and expenditures are calibrated to IMF macro-fiscal data. We set the rate of the (proportional) corporate income tax to raise, as a percentage of GDP, the amount of corporate income tax revenue in each region. We set the rate of the payroll tax to raise the share of pension benefits financed by payroll taxes. Calibrated tax rates are presented in Table A18. Income and consumption taxes adjust endogenously to balance the budget with the ratio between the two revenue sources calibrated to IMF data. We do not directly incorporate any other tax types in our model. Hence, we assume that non-tax government revenue (e.g. user fees and permit sales) is raised from our three main types of taxes – income, consumption, and corporate income – proportionate to their revenue and increase the revenue-to-GDP ratio targets accordingly.

[33] Non-U.S. regions are assumed to take the U.S profile, but scaled to their specific age of labor force entry, i.e. workers in regions where labor force entry occurs at age 18 have zero assets at age 18. We assign the U.S. profile globally for two reasons. First, while OECD collects data on wealth inequality (Balestra and Tonkin, 2018), this data is limited with regard to our model's scope and does not include countries such as China and India. Second, the initial asset distribution has very little influence on long-term household dynamics. Beyond the immediate future, asset stocks are determined by agents' consumption and leisure preferences, which are calibrated to region-specific macro data.

	Consumption Tax	Income Tax	Payroll Tax	Corporate Income Tax
USA	9.0	24.4	13.5	9.8
WEU	24.8	17.5	23.9	18.6
JSHK	15.4	11.8	22.4	24.2
CHI	35.8	4.6	10.7	22.9
IND	13.6	3.9	3.6	18.5
RUS	22.9	7.4	17.2	17.3
CAN	14.0	24.8	4.3	26.5
EEU	29.2	8.9	16.1	16.0
SAP	15.1	2.3	0.6	23.2
BRA	28.2	8.4	21.9	19.8
MEX	12.4	7.5	5.1	19.5
SAF	25.3	17.4	1.1	32.6
MENA	23.3	4.3	10.4	12.7
SLA	21.7	8.6	6.1	27.0
SSA	17.6	7.0	4.3	18.7
SOV	15.0	4.6	4.1	23.5
GBR	20.0	18.1	13.8	23.7

Table A18: Average 2017 Tax Rates By Region

On the expenditure side, we calibrate spending levels on five different categories: healthcare, education, general welfare expenditures, pensions, and other general public purchases of goods and services ("miscellaneous" expenditures, e.g. expenditure on military and federal administration). We treat miscellaneous expenditures as government consumption. The share of healthcare and education expenditure that is not government consumption is rebated to households according to equation (16). We set each region's initial debt-to-GDP ratio to match the model's net interest payments on debt with 2017 data. With the exception of the UBI scenarios, which involve some debt financing, we assume that debt is fixed as a share of GDP.

We calibrate fossil fuel rent flows to World Bank (2021) data and match rents as a share of GDP in the initial year. The share of this flow owned by the government is calibrated to the share of natural resource revenue that accrues to the government. We estimate the region-specific exhaustion date using data on the average rate of extraction and proven reserves. If a region's fossil fuel stock is projected to last beyond 2150, we exhaust the flow at the end of 2150.

We interpret leisure in our model as a measure of labor force participation (LFP). The intratemporal ELS, ρ, is calibrated to roughly match the ratio of LFP between 21-year-old and 60-year-old Americans in 2017, which is estimated to be approximately 1.17.[34] We adjust the leisure preference parameter, $\varepsilon_{t,k}$, to match the LFP of each skill group in each region. These targets are derived by combining a non-region-specific distribution target of relative LFP by skill-group and a region-specific target of the average LFP.[35] The ratios of relative leisure taken by skill groups in a given region are estimated by calculating the fraction of non-participants of the labor force among low-, middle- and high-skilled American households – defined (as above) as those in the bottom 76, middle 20, and top 4 percent by wage income – in SCF 2016. In each region, the target for the weighted average level of leisure taken across all skill groups is estimated with ILO data (ILOStat, 2020).[36] Table A19 summarize leisure targets by skill-group and region.

We calibrate the income tax system by using ξ_t, the progressivity term, to adjust the ratio between the METR of the high-skilled group and the average tax rate across all workers. The target for this ratio is the ratio between the METR of the top 4 percent of households by labor earning (i.e. high-skilled households) and the average effective tax rate of U.S. households as estimated by Altig et al. (2020) and Auerbach et al. (2016). This ratio, approximately 1.72, is set as the income tax progressivity target for all regions. Calibrated average and marginal income tax rates by skill-group are summarized in table A20. As shown, low-skilled and high-skilled U.S. workers face, respectively, an average income tax rate of 20.6 and 31.1 percent in 2017. These

[34] Age 21 is chosen as the year that Americans enter the labor force. Age 60 is chosen to represent an age when LFP declines, but is still sufficiently far away from age 66 when the model assumes that all Americans retire. Since the actual LFP-age curve does not exhibit a kink at 66, it is difficult to match leisure taken by individuals in the years immediately prior to turning 66. LFP by age is estimated from 2016 Survey of Consumer Finance data.

[35] Specifically, all regions share the same ratio of LFPs between skill-group as estimated by SCF data, but each region's average LFP across skill-groups, which we draw from ILO data, is region-specific. In other words, SCF data determines the (non-region-specific) shape of the LFP profile by income, and ILO data determines the overall, region-specific level of participation. Leisure preferences coefficients are larger for agents with higher income. This is because LFP does not vary substantially across skill groups. Hence, high-skilled individuals must have a stronger preference for leisure. Otherwise they would spend a much larger than observed amount of their time endowment on leisure.

[36] We aggregate ILO data using the same working age interval as the model's assumption. For example, in the U.S. we calibrate leisure to the fraction of non-participants among 21 to 65-year-old Americans.

	High-skilled	Mid-skilled	Low-skilled
USA	11.7	8.5	17.1
WEU	14.9	10.9	21.8
JSHK	15.0	10.9	21.9
CHI	17.1	12.5	25.0
IND	28.2	20.6	41.2
RUS	8.2	6.0	12.0
CAN	7.3	5.3	10.7
EEU	24.0	17.5	35.1
SAP	11.9	8.7	17.4
BRA	17.5	12.8	25.6
MEX	20.8	15.1	30.3
SAF	23.0	16.8	33.6
MENA	26.3	19.2	38.5
SLA	14.9	10.9	21.8
SSA	9.1	6.6	13.3
SOV	20.3	14.8	29.7
GBR	8.3	6.1	12.1

Table A19: Leisure Targets As Percent of Time Endowment (%)

rates are comparable, albeit slightly above actual rates. The IRS (2019) estimates an average 2017 federal rate of 14.4 percent for all U.S. taxpayers, and a rate of 24.3 percent for the top 4 percent. Based on authors' calculations using data from Altig et al. (2020), average 2016 state rates are, respectively, 2.5 percent for all households and 4.1 percent for the top 4 percent of households by earning.[37]

The region and skill-group specific pension replacement rate ν_k is calibrated to the progressivity of the Social Security Administration's payout schedule assuming maximum-taxable earnings since age 22 and withdrawal at age 66 in 2017.[38] The pension system is highly redistributive, with high-skilled and middle-skilled U.S. workers only receiving 2.3 and 2.0 times the pension benefits of low-skilled workers, respectively.[39] We set these ratios as the pension progressivity targets in all

[37] This discrepancy is a result of budget balancing utilizing a combination of income and consumption taxation. Specifically, in the model U.S. income taxes raise 21.2 percent of GDP. Yet the actual ratio is only 18 percent, which includes the portion of non-tax government revenue accounted for through income taxation. We choose to explicitly target the METR for high-skilled workers because, among skill types, their labor-leisure decisions are most important to the consequences of skill-biased automation. Guner et al. (2014) and Holter et al. (2019) propose alternate, parametric tax progressivity calibrations. However, Guner et al. (2014) focus exclusively on federal income taxation. Holter et al. (2019) estimates tax progressivity parameters for the U.S. and a number of OECD countries, but only based on data from 2000 to 2007.

[38] See `https://www.ssa.gov/oact/cola/examplemax.html` for benefit examples of workers with maximum-taxable earnings since 22. We translate the model's average wage rates for each skill group into Averaged Indexed Monthly Earnings (AIME).

[39] The average middle- and high-skilled worker's wage rates are 3.5 and 18 times that of the average low-skilled worker, respectively.

	High-skilled		Mid-skilled		Low-skilled	
	Avg.	Marg.	Avg.	Marg.	Avg.	Marg.
USA	31.1	42.1	22.2	24.5	20.6	21.2
WEU	21.6	30.2	16.7	20.4	14.5	16.0
JSHK	14.5	20.1	11.3	13.7	9.8	10.8
CHI	5.9	8.0	4.3	4.9	3.9	4.0
IND	5.0	6.8	3.5	3.8	3.3	3.4
RUS	9.5	12.9	6.8	7.5	6.3	6.5
CAN	30.7	42.8	23.8	29.0	20.7	22.8
EEU	11.0	15.3	8.5	10.1	7.4	7.9
SAP	2.9	4.0	2.1	2.4	1.9	2.0
BRA	10.8	14.5	7.5	7.8	7.2	7.2
MEX	9.6	12.9	6.7	7.0	6.4	6.5
SAF	22.2	30.1	16.1	17.8	14.7	15.0
MENA	5.4	7.3	4.0	4.4	3.6	3.7
SLA	11.1	14.9	7.6	8.0	7.3	7.4
SSA	8.9	12.0	6.5	7.3	5.9	6.0
SOV	5.8	7.9	4.2	4.6	3.9	4.0
GBR	22.5	31.0	17.1	20.2	15.2	16.4

Table A20: 2017 Income Tax Rates By Region and Skill-group

regions.

The contribution ceiling on payroll taxes is calibrated to the ratio between the 2017 Old Age, Survivor and Disability Tax (OASDI) tax contribution limit and the average 2017 U.S. wage rate of the model.[40] This ratio, approximately 2.72, is set as the target for all regions. Consequently, all high-skilled workers and some middle-skilled workers in the U.S. experience a marginal payroll tax of zero percent, but some high-skilled workers globally experience a non-zero marginal payroll tax depending on their age and the level of local wage inequality.

We additionally calibrate non-region-specific factor productivities for each type of labor and capital utilizing estimates of cost-saving productivity gains from automation Acemoglu et al. (2020). Details of this procedure are discussed in Section 3.4.

B Calibration Algorithm

Our model has a large number of interrelated parameters to calibrate. To do so, we developed a calibration algorithm. To understand the algorithm, consider the parameterization of a model with l parameters, $\mathrm{P} = \{\mathrm{p}_1, \ldots, \mathrm{p}_l\}$ that govern model outputs $\mathrm{T} = \{\mathrm{t}_1, \ldots, \mathrm{t}_l\}$. We seek a value for P

[40]The limit is set at $127,200 following https://www.ssa.gov/oact/cola/cbb.html

such that the model outputs correspond to their target values, e.g. $\mathrm{T} \approx \mathrm{T}^*$ where T^* is a vector of targets drawn from macro-fiscal data (e.g. targets discussed in Appendix section A). Assume that the relationship between T and P can be linearly approximated in a small neighborhood around T^*. Then calibration entails finding M such that

$$\mathrm{T} = \mathrm{M} \cdot \mathrm{P} = \begin{pmatrix} \mathrm{m}_{11} & \dots & \mathrm{m}_{1l} \\ \vdots & \ddots & \\ \mathrm{m}_{l1} & & \mathrm{m}_{ll} \end{pmatrix} \cdot \mathrm{P}, \tag{28}$$

where m_{ij} is the local gradient of t_i with respect to p_j. Given M, the model can be parameterized simply by solving (28) at T^*.

This is difficult for two reasons. First, M depends on our choice of targets and must, therefore, be estimated on a relatively small interval – less than 50 percent of each output-t value of absolute deviation – around T^*. Hence, we need an iterative procedure to calculate M. Second, generating data for the estimation of parameters by running the model is, at 1 to 2 CPU-hours per run, costly. Were l small, it one could estimate M by exhaustively searching over a grid of parameter values. But our model has 416 parameters that ex-ante may need to be jointly calibrated. Using Newton-Raphson is computationally infeasible since deriving requisite cross-partials would require running the model hundreds of thousands of times.

An alternate approach is to set $\mathrm{m}_{ij} = 0$ for $i \neq j$, i.e. ignore off-diagonal elements in M. However, assuming that we update parameters simultaneously, this is feasible only if each target can be matched by adjusting exactly one parameter. If one or more of the off-diagonal, m_{ij}, terms is large, guesses of P will persistently overshoot. This is the case for our model. Clusters of t_i terms – pension expenditures in 2017 and 2050, for example, – are highly dependent on the same set of parameters. A third alternative is to ignore off-diagonal elements, but update only one parameter for each convergence of the model. The sequential updating of over 400 parameters is, however, no more feasible than comprehensively estimating M.[41]

Fortunately, most off-diagonal elements in M are essentially zero. For example, a given tax

[41] Even if we update every region simultaneously, adjusting 17 parameters at a time, fully updating the 24 sets of region-specific parameters requires 24 runs of the model.

multiplier should have little effect on revenues from other taxes, let alone tax revenues in different regions. In other words, M can be approximated by a (relatively) sparse matrix. On-diagonal m_{ii} terms can't be zero – otherwise the system is under-identified. Therefore, before estimating off-diagonal elements, we establish a bijective correspondence between T and P by choosing on-diagonal elements that are most influential in determining t_i.

Our procedure involves four steps prior to applying our automated calibration algorithm. First, we generate a large number of permutations of parameters – approximately 10,000 – and run the model for each permutation. Permutations are generated by random draws over uniform distributions symmetric over an initial guess of parameters. As draws are made, the distributions' supports are gradually increased until the model fails to converge. This process guarantees sufficient distributional coverage of each iteratively calibrated parameter. It also provides a rough estimate of the range of each parameter within which the model will successfully converge.

Second, using these results, we regress each output t_i against the entire set of parameters P,

$$\mathbf{t}_i = \mathrm{B}_i \cdot \mathbf{P}, \tag{29}$$

where $\mathbf{t}_i = \{\mathrm{t}_i^1, \dots\}$ is a vector of model outputs of i generated by permuting P, and $\mathbf{P} = \{\mathrm{P}^1, \dots\}$ is the corresponding matrix of permutations of parameters. Specifically, we non-parametrically bootstrap over draws of $\mathbf{t}_i$ using a LASSO regression (Tibshirani, 1996). L1 regularization is particularly useful for our purposes as it penalizes most coefficients in B_i to exactly zero. Because of computational constraints, we set a highly aggressive, t_i-specific shrinkage parameter so that nearly all coefficients are penalized to zero. For most targets, the shrinkage parameter is increased until less than 5 coefficients are consistently nonzero.

Third, for each t_i we identify the parameter that corresponds to the coefficient in B_i with the largest number of non-zero estimates (our definition of most influential), and set it as p_i. We address ties by gradually increasing the number of bootstrap runs until we observe a clear ranking. If two outputs share a top-ranking parameter, we prioritize production-side targets over consumption-side (e.g. saving and labor-leisure decisions) targets and consumption-side targets over fiscal targets.[42]

[42]We tiebreak within each of these groups by prioritizing the target that is, conceptually, more clearly linked to

Fourth, after sorting P, the off-diagonal b_{ij} coefficients, $l^2 - l$ in total, are sorted by the fraction of bootstrap coefficient values that are non-zero.[43] Next, we identify the most important off-diagonal element b_{ij} as determined in this step, and identify its corresponding target, t_i, and parameter, p_j. As described momentarily, we explicitly update b_{ij} by establishing two equations and two unknowns, with the second target having parameter j as its diagonal element. Although we're considering the example of solving for two parameters with two equations, our algorithm works for solving n parameters with n equations. Having established the first two-by-two partition by considering the most important off-diagonal coefficient, we follow the same procedure to establish the next two-by-two (or n-by-n) partition) by determining the second most important off-diagonal coefficient. The number of jointly updated parameters and the size of n depend on the desired calibration precision.[44]

Having determined our partitioning equations and unknowns, we apply our calibration algorithm, which has three steps. We start with a guess of P and a given set of m_{ij} (and therefore m_{ji}) elements that are explicitly estimated. First, a number of initializing runs of the model are made, incrementing P linearly proportional to each b_{ii}, the diagonal coefficient, to match targets to a rough level of precision.[45] Specifically, we use 50 runs for versions of the algorithm with clusters of up to two jointly updated parameters.[46] In each new run, each parameter is incremented by 2 percent of $\hat{p}_i - p_i$, where p_i is the starting guess, and $\hat{p}_i = t_i^*/b_{ii}$ is a rough guess of the parameter that calibrates target i based on b_{ii}. All parameters are updated simultaneously.

Second, using the 50 sets of parameters and model outputs, we estimate each non-zero element

the parameter. For example, GDP and private consumption are both strongly influenced by TFP, but TFP is clearly the primary determining factor of economic output. It is possible that a parameter's coefficient is penalized to zero for all bootstrap runs of all targets. In this case, the model is under-identified, and one target must be replaced with an alternative that is directly influenced by this parameter.

[43] We don't rank parameters by the size of their corresponding b_i coefficients. This is because some parameters can take on a much larger range of values relative to others. Hence, even if the absolute value of b_{ij} is extremely small, m_{ij} can still be important for calibrating the model, so long as we can reject the null hypothesis, e.g. $m_{ij} = 0$, with a high level of confidence. The non-parametric bootstrap is equivalent to a null hypothesis test against m_{ij} being precisely zero, where the fraction of non-zero values of each coefficient is an estimate of the confidence that the null can be rejected.

[44] This is achieved by running a number of candidates of the algorithm, as described below, with progressively larger numbers of estimated off-diagonal m_{ij} elements, until the model is calibrated to the desired precision.

[45] Approximately 25-50 percent of the target's absolute value depending on the target.

[46] We use 200 runs for versions with three jointly updated variables.

in M. For parameter i, if $m_{ij} = 0\ \forall j \neq i$, then no other parameter needs to be jointly updated with i. Therefore, new guesses of p_i can be proposed with a simple, univariate Newton-Raphson approach where m_{ii} is updated as t_i converges to t_i^*. If $m_{ij} \neq 0$ for some $j \neq i$, then p_i and p_j must be updated jointly. The sequences of t_i's and t_j's from step one are individually regressed on both parameters, with parameters being weighted inversely to the order in which the run is performed.[47] This allows us to grid search over combinations of p_i and p_j to propose a joint update that minimizes the sum of predicted squared deviations of t_i and t_j with respect to t_i^* and t_j^*.[48] In the case of three jointly updated parameters, each target is regressed on all three parameters.

We re-run the model with the proposed parameters and check for precision. If the model is not yet sufficiently calibrated, a third step is required. We add the latest, 51$^{\text{st}}$ run's output and parameters to the data, discard the observation of the first run – keeping the number of observations at 50 – and re-estimate each of the regressions in step two. Because of the inversely proportional weighting of the observations, new parameter guesses proposed with joint updating receive the largest weights. Early runs, where $|T - T^*|$ values are relatively large, are gradually eliminated from the estimation process. We continue to perform additional runs, adding their output to the moving window of observations, until all targets are matched to within 2 percent (or 2 percentage points for targets that can take on negative or positive values).[49]

By repeating steps one through three, the model can be recalibrated at the same level of precision to any alternate T^*. Higher levels of precision – below 1 percent for all targets – can be reached by allowing for more complex interactions involving sets of three output-parameter pairs. However, this comes at substantial computation cost. With respect to the 2 percent target, we find that, at a minimum, three pairs of parameters must be jointly updated for each region. First, TFP and the saving preference rate are jointly updated to match GDP and private consumption as a share of GDP. Second, the wage multiplier of current retirees and the pension replacement rate are

[47] A number of candidate regressions are utilized. We select for best fit among linear and log transformations of the independent and response variables, with and without an intercept, and with and without quadratic terms. Data points are weighted by the inverse of their location among the 50 runs, i.e. output set to 1 is weighted at $1/50 = 0.02$ and output set 50, from the most recent converged run of the model, is weighted at 1.

[48] For two jointly updated variables, we use a grid size of 1,000. The grid is centered on the current guesses of the parameters. The size of the grid is set to 25 percent of the range of each variable, as estimated in step one of the setup process. A multi-stage grid search is used for updating groups of three variables.

[49] This usually requires 5 to 20 convergences of the model in addition to the 50 initializing runs.

jointly updated to match pension expenditures in 2017 and 2050. Third, the corporate tax rate and the capital share are jointly updated to match corporate tax revenue as a share of GDP and the ratio of capital stock to GDP. Additionally, the global asset multiplier is jointly updated with each region's asset multiplier to match our 6 percent global interest target and region-specific net foreign asset positions.[50]

The algorithm requires no manual adjustment for re-parameterizations of the model – Appendix section C provides one such example – and can run in parallel to precisely calibrate the model to many sets of targets simultaneously. Major revisions – alternate production functions or additional sectors, for example – can be calibrate with a relatively small amount of human intervention.[51] Reducing the human factor in parameterizing large-scale OLG models introduces numerous research possibilities. Our ongoing work aims to comprehensively quantify the sensitivity of such models to the uncertainty of macro-fiscal data and projections of global economic growth.

C Sensitivity Analysis I: Alternate Labor Productivity Calibration

As a sensitivity analysis, we simulate a version of the model identical to that of section 4.1 with one major exception. Specifically, we calibrate the model such that $\chi_2 = \chi_3 = 0.3$ in equations (21) and (22), i.e. a unit share shift toward either middle- or high-skilled labor, at the cost of low-skilled labor, induces the same GDP change as a unit share shift toward capital. Hence, all else equal, a one percentage point increase to the share of high- or middle-skilled labor, accompanied by an equal-sized decline to low-skilled's share, causes a 0.3 percent increase to GDP. The automation scenario is otherwise identical to the baseline, entailing a new, region-specific frontier production function with the same shares as that of section 4.2. Henceforth, the "original calibration" refers to the version of the model with the original labor factor productivities as calibrated in section 4.1; "alternate calibration" refers to the version of the model as calibrated in this subsection.

[50] In this case, we average over the proposed global asset multiplier from each region-specific estimating equation.

[51] This is mostly related to tiebreaking when one parameter appears to be equally important in adjusting multiple outputs. Additionally, some "hyperparameters" of the algorithm – the choice of utilizing 50 initializing runs to estimate M, for example – are specific to our model, and may need to be adjusted for a different setup.

Conceptually, this alternate calibration assumes that the marginal output gain from substituting labor with capital is identical to that of substituting low-skilled labor with high-skilled (or middle-skilled) labor. This implies, naturally, that high- and middle-skilled workers are more productive than in our original calibration. Factor productivities are summarized in table A21. Relative to the calibration of 4.1, productivities of low-, middle-, and high-skilled workers are 16.9 percent lower, 11.1 percent higher, and 10.6 percent higher.

	Original Calibration	Alternate Calibration	% Diff
Capital	0.074	0.075	0.9
Low-Skilled Labor	0.964	0.801	-16.9
middle-skilled Labor	3.374	3.750	11.1
High-Skilled Labor	17.501	19.351	10.6

Table A21: Calibrated Factor Productivity Coefficients

In this scenario, skill-biased technological change is more attractive relative to the original calibration. Regions' new automation choices are displayed in figure A12. Regions that delay automation automate sooner in the alternate scenario. China, which automates in 2027 under the original calibration, now automates immediately. The two other regions that delay automation, former Soviet states (SOV) and Southeast Asia, both adopt the new technology roughly 20 years before doing so in the original calibration. However, with one exception – South Africa – regions that never automate with the original factor productivities also never automate with the alternates. This is in large part because less developed regions that use less capital also tend to have fewer high-skilled (and more low-skilled) workers, lowering the marginal benefit to output of an increase to high-skilled productivity.[52] Therefore, the frontier technology remains inefficient in these regions.

South Africa, which doesn't automate in the original model, automates in 2059, reverts to the legacy technology in 2074, and automates again, permanently, in 2102. However, their overall gain is small – at 0.1 to 1.5 percent of GDP between 2059 and 2073 relative to no automation – in the new calibration. Relative to other regions that never automate in the original calibration (e.g. Brazil, Mexico), South Africa has a larger share of workers that are high-skilled. This difference, alongside South Africa's relatively high TFP, means that the automating technology is marginally

[52] This can be seen from equation (22).

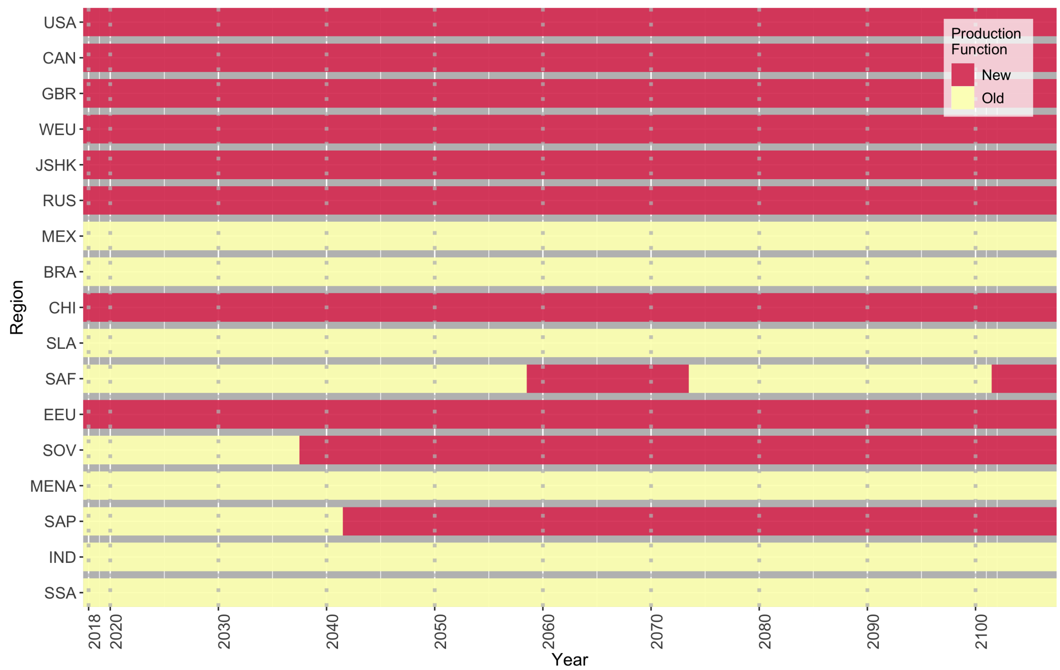

Figure A12: Choice of Production Function by Region, Alternate Labor Productivity Calibration

beneficial with the new calibration.

Interest rates increase with the alternate calibration, driven by higher capital demand and countries' decision to automate sooner, but only very slightly. Paths of the world interest rate are contrasted in figure A13. In 2050, the world interest rate is 3.3 percent with baseline automation under the original calibration, and 3.4 percent with baseline automation under the alternate. In 2100, these rates are, respective, 2.1 and 2.2 percent.

Regions that automate benefit slightly more, in terms of GDP, with the alternate labor productivity calibration. Figure A14 contrasts GDP percent changes, relative to no automation, of the U.S. and the world average.[53] Relative to no automation and by 2050, the U.S's GDP gain increases from the original model's 5.2 percent to 6.5 percent. By 2100, U.S. GDP is 10.1 percent higher than no automation, 2 percentage higher than in the original calibration. Regions that automate sooner with the new calibration experience larger gains to output. In 2050, China's GDP is 3.6 percent higher with the original calibration, and 5.3 percent higher with the alternate.

[53]Note that the no automation scenario is slightly different across calibrations. However, this difference is very small because the no-automation scenario's GDP in both 2017 and 2100 are calibrated to a target. In 2050, U.S. GDP is 0.2 percent higher, assuming no automation, in the alternative calibration relative to the original.

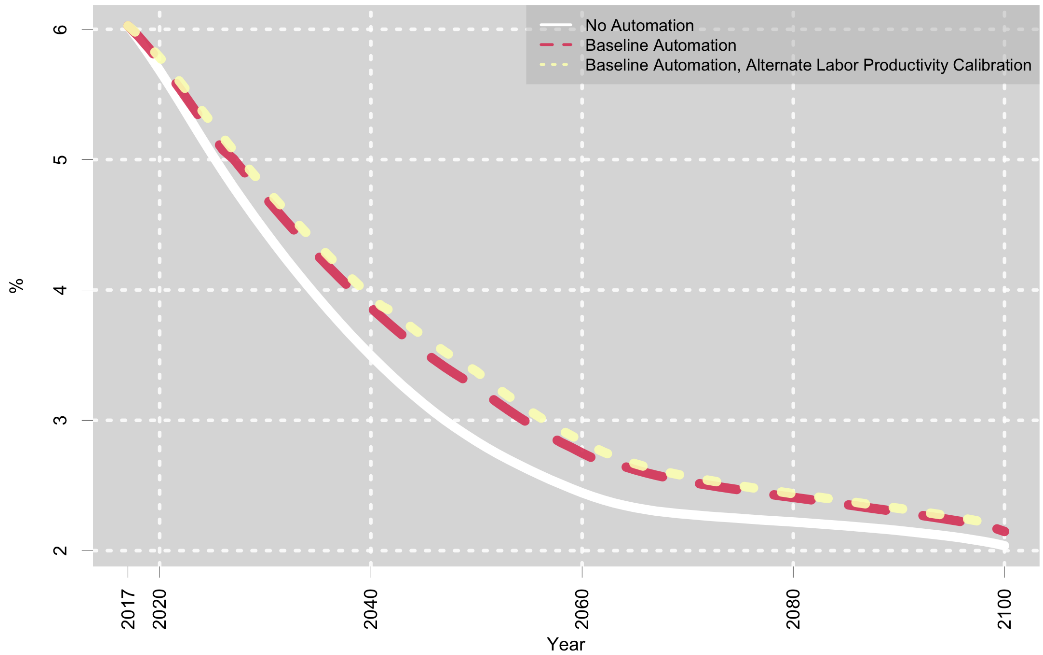

Figure A13: Comparison of Global Post-Corporate Tax Rate of Return

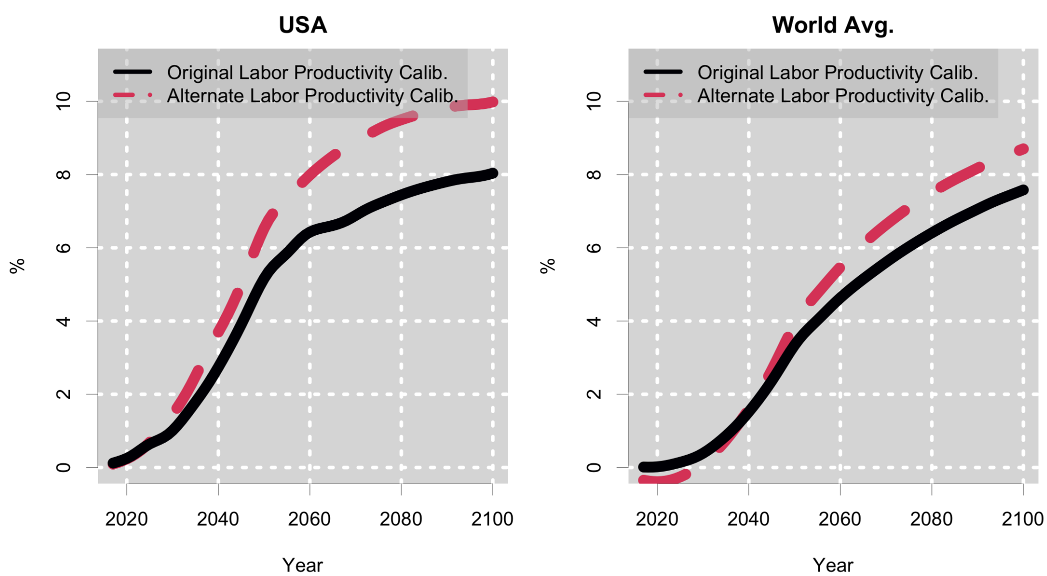

Figure A14: GDP Percent Change in Baseline Relative to 'No Automation' Scenario

It is not possible to precisely compare compensating-differential welfare changes across different calibrations of the model. This is because in order to match each region's consumption and labor force participation at different labor factor productivities, the parameters of the agents' utility functions must also be different. Hence, even relative welfare changes – those presented in figure 6, for example – are denominated in slightly different units. With this caveat in mind, we contrast relative welfare changes from automation with the two calibrations. In regions that choose to automate, the alternate calibration results in a slight increase to workers' welfare relative to the original calibration. In the U.S. and relative to no automation, low-, middle- and high-skilled workers born in 2000 are made 9.8 percent worse, 2.5 percent better, and 18.7 percent better off with baseline automation with the alternative calibration. In contrast, the same cohorts are made, respective, 11 percent worse, 0.8 percent better, and 16.8 percent better off with the original calibration. Other results regarding wage and welfare inequality are also broadly comparable to those of section 4.2. The average global worker alive in 2050 is 0.8 percent worse off, in terms of lifetime utility, and the average worker alive in 2100, is 0.5 percent worse off.[54]

D Sensitivity Analysis II: Increases in the Relative Supply of Skilled Workers

In our main simulation, population shares of each skill group are fixed over time. This is intentional. There is insufficient data to separately calibrate skill group factor shares and the share of the population within each group across all our regions. While the literature on skill-biased technological change does feature related calibrations to explain, for example, the college wage premium, a key issue is that we wish to model top-percentile inequality (which has seen the largest change in factor share globally).[55] To analyze the effect of human capital accumulation on top earners, one would need to measure, objectively, the number of workers across time and countries providing top-percentile labor as defined in our model. No such objective measure exists. Talent, rare skills, high effort, or particularly effective supply of a common skill could all lead to an individual receiving a high wage. Thus, it is difficult to define top-percentile skill by earnings, occupation or education alone. The different factors driving labor income all may change independently over

[54] This is to be contrasted with 0.9 percent and 1.1 percent, respectively, in the original calibration.

[55] Krusell et al. (2000), for example, calibrate a neoclassical aggregate production function where capital is more complementary to high-skilled labor than it is to low-skill labor. The framework of Krusell et al. (2000) is successful in explaining growth in the college/high school wage premium from 1962-1992. However, it is not useful for understanding more extreme inequality – wage changes since the 1980s have not primarily favored college educated workers, but rather top percentile workers.

time, further complicating the issue. Still, the effect of changes to the relative provision of high, medium and low-skill on the impact of automation are important to understand. Therefore, as a first step of a sensitivity analysis, we produce a range of estimates for the rate at which the skills provided to the labor market may change over time. We generated three different estimates, each based solely on U.S. data.

1. **Distributional Assumption:** The principle behind this approach is to assume that those providing high-skill labor are those making a fixed multiple of the median wage (or more). Those making a lower fixed multiple of the median wage, but not as much as the high-skill, are middle skill. These multiples are fixed based on 2017 data so that the high and medium skilled are the correct share of the workforce in that year – 4 percent and the next 20 percent, respectively.

   Our estimates of the share of the workforce with different multiples of the median wage are based on BLS OES data (U.S. Bureau of Labor Statistics). For each occupation in every year (where at least the 25th percentile income level is reported), we use stochastic gradient descent to estimate the parameters of a Burr distribution fitting the observed percentiles under the assumption that occupational incomes are bounded by a minimum of $1 and a maximum of five times the median income. While this functional form assumption may artificially compress the fat tails of highly skewed occupations, such as top executives, the precise nature of inequalities within the top 3 or 4 percent do not concern us. For occupations where the 25th percentile income data is available but higher percentiles are not (these are the highest paying occupations, generally), we infer the shape of wage distribution to be that of the highest income occupation for which the 75th percentile wage is available. In 2017 these are podiatrists. We then take the employment weighted average of these wage distributions to get an estimate of the overall wage distribution in the US in each year.

   With this wage distribution data in hand, we then look to see what percent of the population earned these multiples of the median wage in all years from 1997 through 2016. In 1997, by this measure, approximately 3% of the population were high-skill, and 17% medium skill. We run a linear regression line through these shares over the 21 years of data. The linear projection indicates that 5.1% and 24.% percent of the population will be high and low-skill respectively in 2050.

2. **Occupation Defined:** The idea behind this approach is to select a list of occupations that are plausibly high skilled, and to project growth in these occupations' share of output forward. Using BLS OES data, we calculate the share of the population, in each year back to 1997, in the following occupational categories: Researchers, Scientists (physicist, biologist, mathematicians, statistician, etc.), Engineers, Computer Scientists and Programmers, Post-secondary Teachers, Medical Doctors, Registered Nurses, Accountants, Lawyers, Executives, and Managers (excluding first-line supervisors). The occupations hand-sorted into this list of categories constituted 17% of employment in 2017.

   We are interested in the rate of change in the share of the population in these occupations over time, which we infer to be the growth rate of both high and medium skill occupations as a share of employment. To calculate this, we estimate a linear regression of the share of employees in these occupations over the last 21 years. Multiplying the estimated rate of employment share growth by year by 23 (for a projection 23 years in the future) and by our baseline high and medium skill employment share yields an estimate that 5.3% and 21.4% of the workforce will be providing high and medium-skill labor, respectively, in 2050.

3. **Education Defined:** The logic behind this approach is to select a list of university degrees that are plausibly high skilled, and to project growth in the share of the population being granted these degrees. We conjecture that the entering cohort high-skill labor supply tracks the following index: the share of the population granted STEM PhDs, plus 50% of the share granted medical degrees, plus 25% of the share granted STEM Masters degrees, MBAs or JDs, plus 12.5% of the share granted STEM undergraduate degrees. The sum of these populations being granted degrees in 2017 is 133,000 people. This should be contrasted with 176,000, which is 4% of the U.S. cohort entering the work force in our model in 2017.

   We assume that the high-skill share of the new labor force is a constant multiple of this measure over time. The high skill share is approximately 1.32 times the above index. We calculate the same fixed ratio of middle-skill new workers to new bachelors degrees granted. Projecting growth in these measures moving forward to 2050 yields 5.9% of workers entering the labor force being high skill in 2050, and 19.7% being medium skilled.

Table A22 summarizes the change in employment shares by skill group projected in 2050 under each sensitivity scenario. For the first two scenarios, employment shares after 2050 are fixed at their 2050 levels. For the final scenario, which is calibrated on the skill of those entering the labor force, employment shares continue to evolve for another half-century beyond that, at which point the workforce overall will the same skill provision mix as the cohort entering the labor force in 2050 and later.

Table A22: Population and Income Shares

Scenario/Year		High Skilled	Mid Skilled	Low Skilled
2017 Levels	population (%)	4.0	20.0	76.0
	labor income (%)	33.3	33.3	33.3
Baseline Automation Scenario, 2050	population (%)	4.0	20.0	76.0
	income (%)	41.9	32.5	25.5
Median-income denom. skill-groups, 2050	population (%)	5.1	24.1	70.8
	income (%)	47.7	32.0	20.3
NCES data estimated skill-groups, 2050	population (%)	4.6	23.5	71.9
	income (%)	45.9	32.9	21.2
BLS OCC estimated skill-groups, 2050	population (%)	5.5	26.5	68.0
	income (%)	48.6	33.4	18.0

To simulate each scenario in the model we make an additional assumption and adjustment. First, we assume that no one changes skill groups later in life. Rather, an increase in the high-skill share of the labor force is due to an increase in the share of the population 'born' into the high skill group. Members of the increasingly large high-skill group have their consumption paid for as children by the high-skill, and are bequeathed inheritances from high-skill grandparents, parents and spouses. This is for computational tractability, as allowing any skill group of individual to be born to any skill group of parent increases the amount of lifetime consumption, labor, leisure, and saving paths that need to be tracked by an order of magnitude. This modeling choice be rationalized as high-skill individuals paying for the induction of lower-class individuals into high society through non-distortionary taxation. Alternatively, it can literally represent a disporportionate increase in the fertility of the high skill, alongside perfect intergenerational immobility.

The adjustment we need to make is to the definition of automation in the model. As noted above, the effect of automation on inequality in the model is calibrated on the historical pattern of inequality in the U.S. Since 1980, the top 4% of earners have seen their share of income increase. In the baseline scenario, we calibrate automation's impact on the Cobb-Douglas share terms for the different types of labor to match this change exactly. However, if there was an increase in high-skill labor in the past, then there would have needed to be an even more rapid increase in the high-skilled's Cobb-Douglas share term to match the observed rate of productivity growth. We need to adjust for this factor to make the baseline and skill-change scenarios comparable. To make this adjustment, we use CBO data to construct empirical Lorenz curves from 1980 through 2017. For each year, we estimate the share of income paid to each skill group. Running a regression on this line generates a new estimate of the rate at which Cobb-Douglas coefficients have been changing in the U.S. over that interval.

Figures 10 and A16 show the change in US GDP and welfare, respectively relative to baseline, from automation under different labor supply scenarios. As can be sen growth in high skill labor supply alongside automation increases the benefits from automation on GDP and for almost all skill groups and birth cohorts. The low-skill benefit the most, as their skills become relatively abundant, leading automation to actually be good for them. Most skill cohorts of the middle and high skill benefit as well. While some of the currently living middle and high skill are made slightly worse off from the additional competition, middle and high skill workers of the Zoomer generation and later are generally made better off.

Figures A15 and A17 report the same outcomes but for a variation of the scenario where the entire world experiences the same high-skill labor supply growth. The beneficial consequences for US welfare and GDP growth are similar, but slightly attenuated.

Figures A18 and A19 report capital stocks, relative to no automation, for the same scenarios. As can be seen, capital stocks increase dramatically, explaining the long-run increase in high-skill welfare, despite facing more competition from high skill workers.

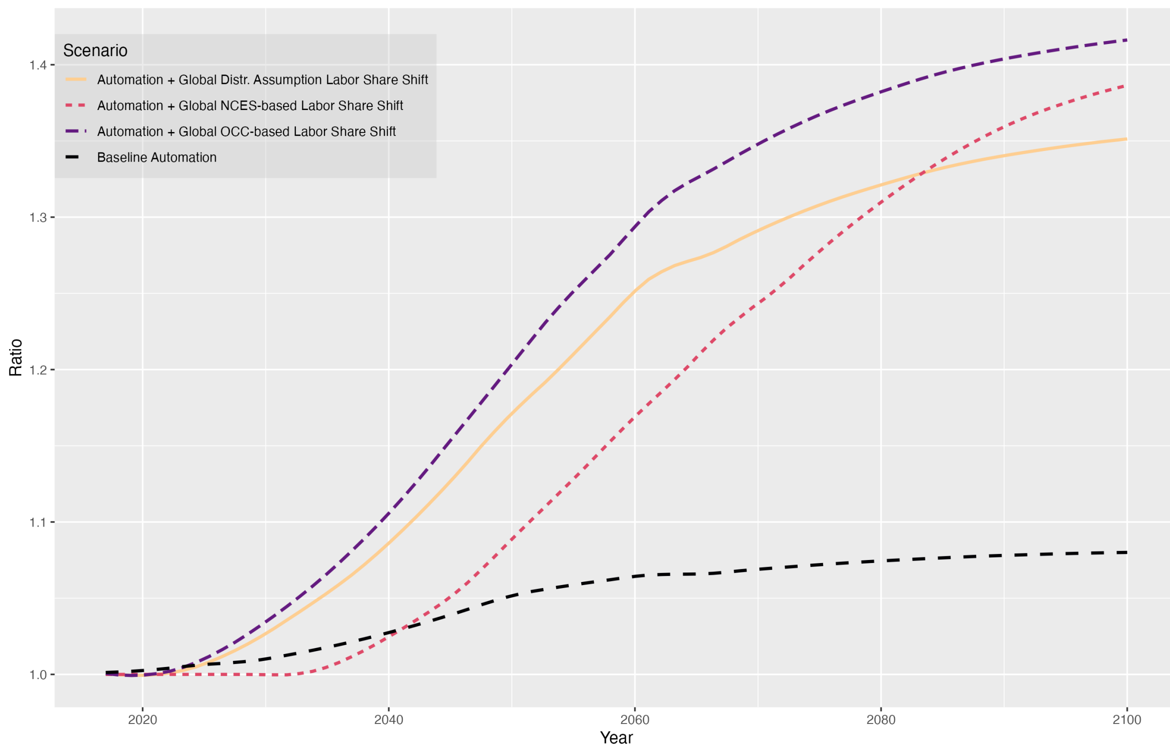

Figure A15: U.S. GDP Relative to No Automation, Global Scenarios

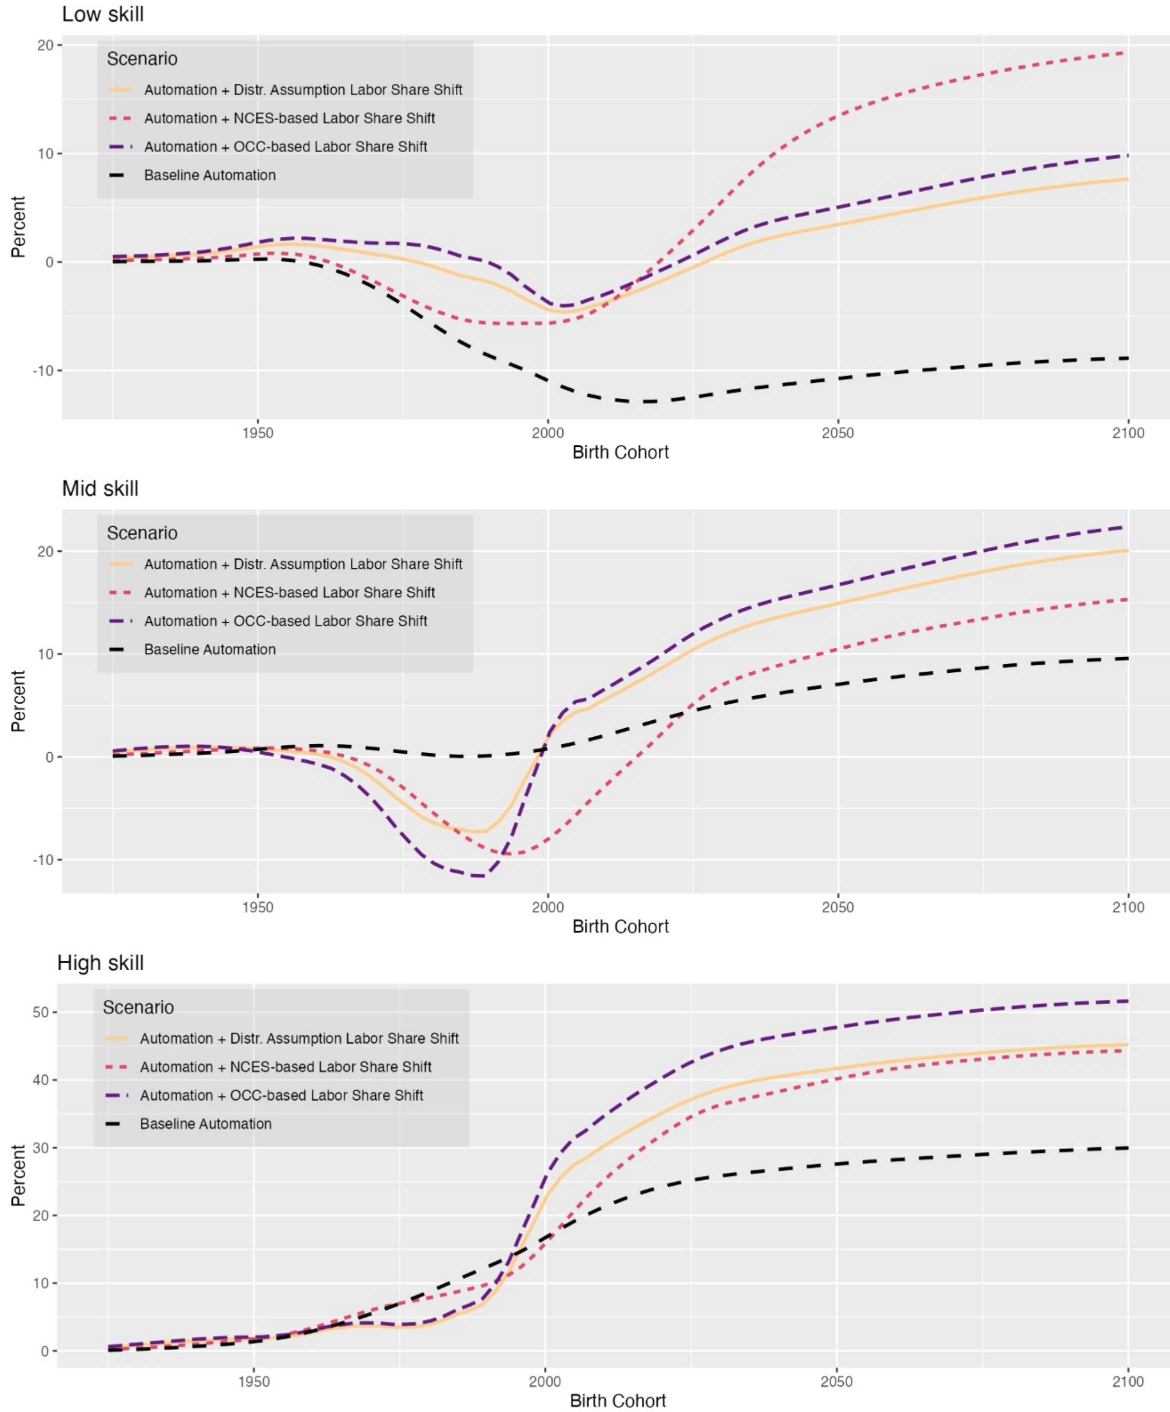

Figure A16: U.S. Welfare Relative to No Automation, US Human Capital Accumulation Scenarios

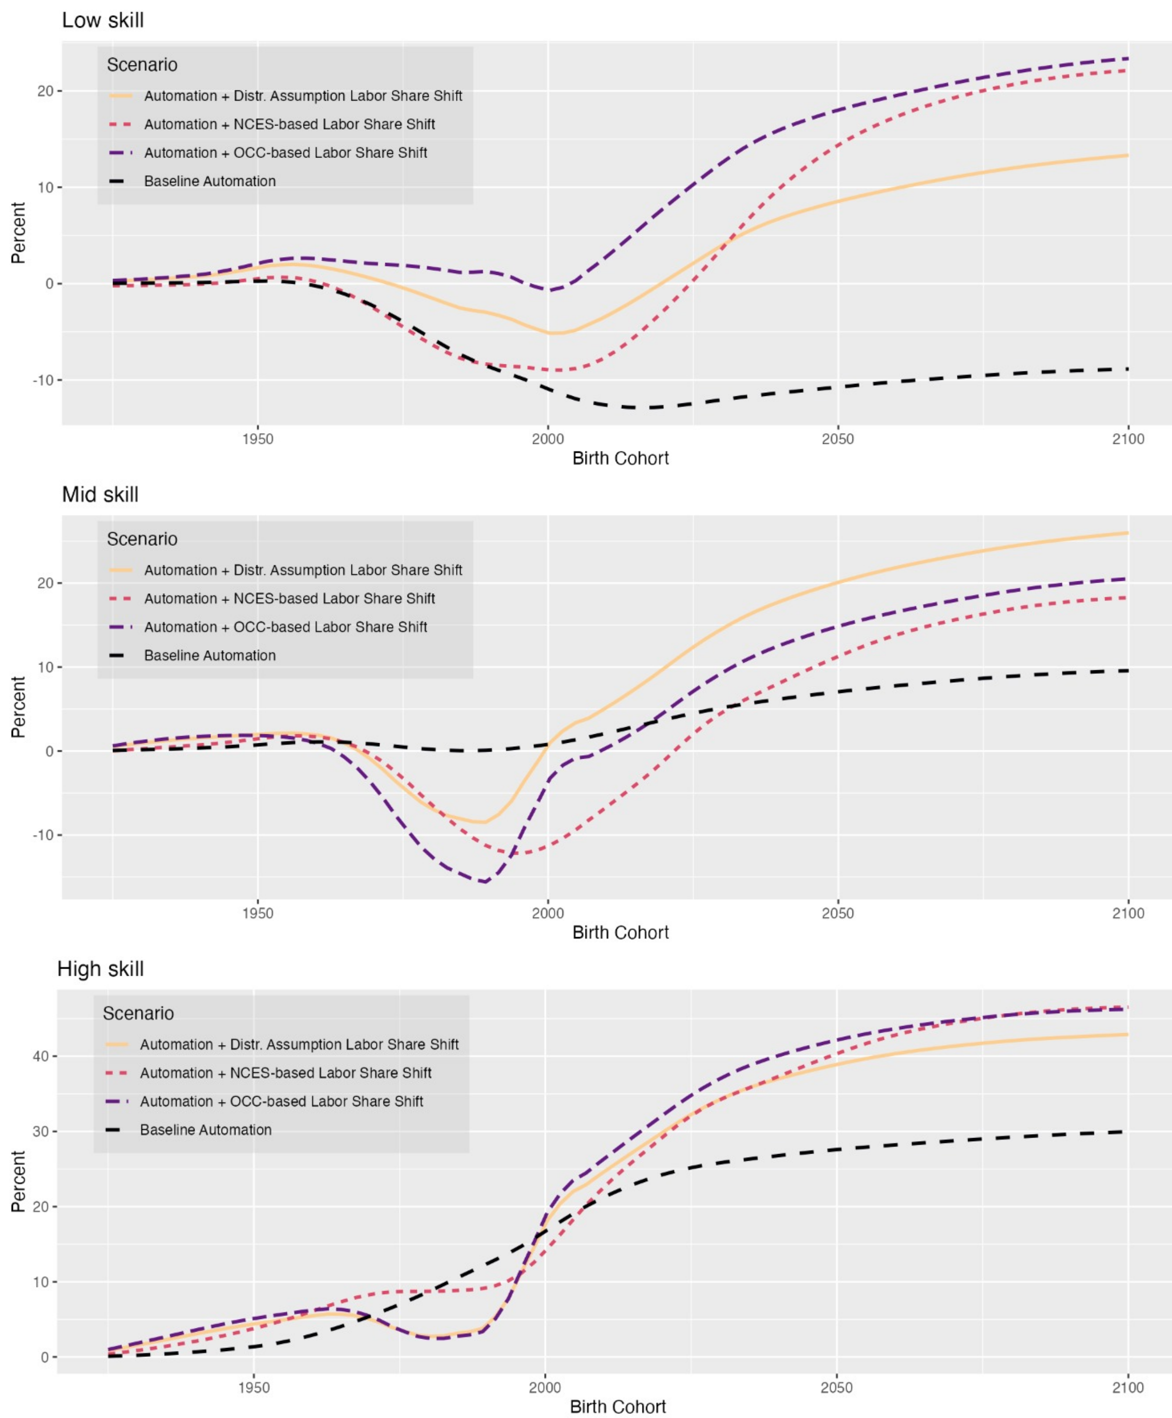

Figure A17: U.S. Welfare Relative to No Automation, Global Increases in the Relative Supply of Skilled Workers

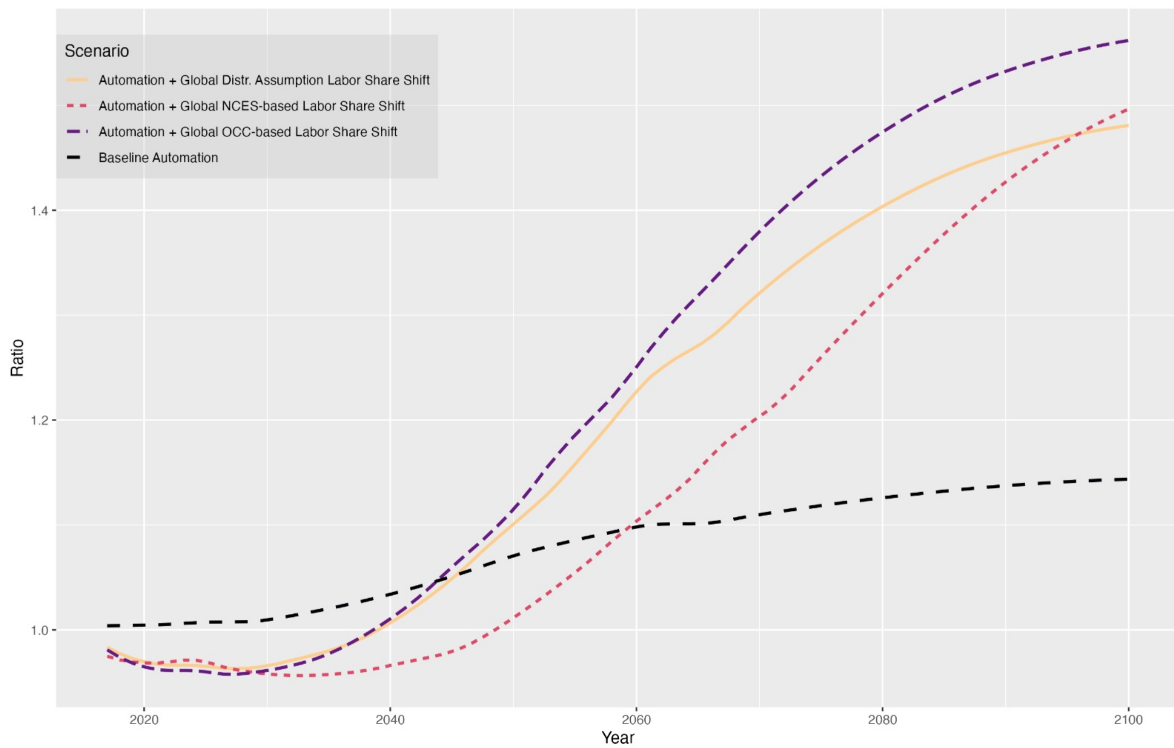

Figure A18: U.S. Capital Relative to No Automation, US Increases in the Relative Supply of Skilled Workers

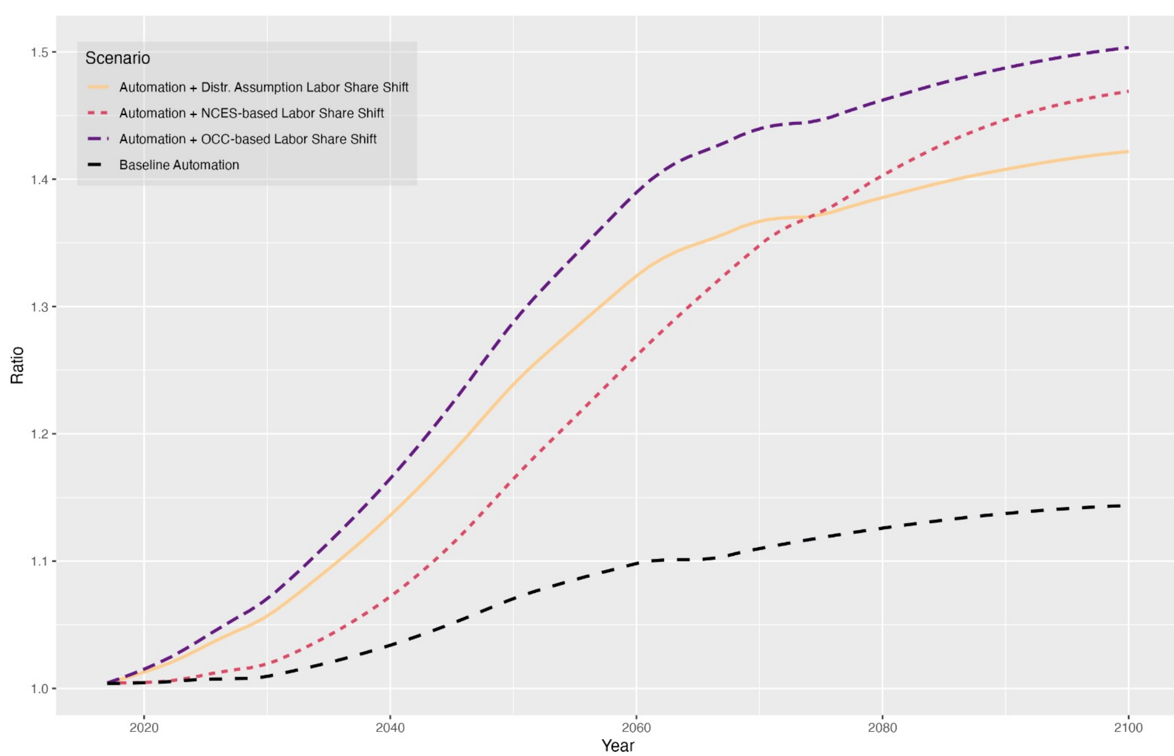

Figure A19: U.S. Capital Relative to No Automation, Global Human Capital Accumulation Scenarios

References

Daron Acemoglu and Pascual Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6): 1488–1542, 2018a.

Daron Acemoglu and Pascual Restrepo. Automation and new tasks: The implications of the task content of production for labor demand. *Journal of Economic Perspectives*, 33(2):3–30, 2018b.

Daron Acemoglu and Pascual Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6): 1488–1542, 2018c.

Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30, May 2019. doi: 10.1257/jep.33.2.3. URL https://www.aeaweb.org/articles?id=10.1257/jep.33.2.3.

Daron Acemoglu and Pascual Restrepo. Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, 128(6):2188–2244, 2020.

Daron Acemoglu, Claire Lelarge, and Pascual Restrepo. Competing with robots: Firm-level evidence from france. 110:383–88, 2020.

David Altig, Alan J Auerbach, Laurence J Kotlikoff, Kent A Smetters, and Jan Walliser. Simulating Fundamental Tax Reform in the United States. *American Economic Review*, pages 574–595, 2001.

David Altig, Alan J Auerbach, Laurence J Kotlikoff, Elias Ilin, and Victor Yifan Ye. Marginal net taxation of americans' labor supply. Working Paper 27164, National Bureau of Economic Research, May 2020. URL http://www.nber.org/papers/w27164.

Facundo Alvaredo, Anthony B. Atkinson, Thomas Blanchet, Lucas Chancel, Luis Bauluz, Matthew Fisher-Post, Ignacio Flores, Bertrand Garbinti, Jonathan Goupille-Lebret, Clara Martínez-Toledano, Marc Morgan, Theresa Neef, Thomas Piketty, Anne-Sophie Robilliard, Emmanuel Saez, Li Yang, and Gabriel Zucman. Distributional national accounts guidelines: Methods and concepts used in the world inequality database. 2020.

Alan J Auerbach and Laurence J Kotlikoff. National savings, economic welfare, and the structure of taxation. Working Paper 729, National Bureau of Economic Research, August 1981. URL http://www.nber.org/papers/w0729.

Alan J Auerbach and Laurence J Kotlikoff. National Savings, Economic Welfare, and the Structure of Taxation. In *Behavioral Simulation Methods in Tax Policy Analysis*, pages 459–498. University of Chicago Press, 1983.

Alan J Auerbach and Laurence J Kotlikoff. *Dynamic Fiscal Policy*, volume 11. Cambridge University Press Cambridge, 1987.

Alan J. Auerbach, Laurence J. Kotlikoff, Robert P. Hagemann, and Giuseppe Nicoletti. The economic dynamics of an ageing population. (62), 1989. doi: https://doi.org/https://doi.org/10.1787/054502801660. URL https://www.oecd-ilibrary.org/content/paper/054502801660.

Alan J Auerbach, Laurence J Kotlikoff, and Darryl R Koehler. US Inequality, Fiscal Progressivity, and Work Disincentives: An Intragenerational Accounting. Technical report, National Bureau of Economic Research, 2016.

David Autor, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen. The fall of the labor share and the rise of superstar firms. *The Quarterly Journal of Economics*, 135(2): 645–709, 2020.

David H. Autor. Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives*, 29(3):3–30, September 2015. doi: 10.1257/jep.29.3.3. URL https://www.aeaweb.org/articles?id=10.1257/jep.29.3.3.

Carlotta Balestra and Richard Tonkin. Inequalities in household wealth across oecd countries. 2018. doi: https://doi.org/https://doi.org/10.1787/7e1bf673-en. URL https://www.oecd-ilibrary.org/content/paper/7e1bf673-en.

Seth G Benzell and Erik Brynjolfsson. Digital abundance and scarce genius: Implications for wages, interest rates, and growth. Technical report, National Bureau of Economic Research, 2019.

Seth G Benzell and Guillermo Lagarda. Can russia survive economic sanctions? *Asian Economic Papers*, 16(3):78–120, 2017.

Seth G Benzell, Eugene Goryunov, Maria Kazakova, Laurence J Kotlikoff, Guillermo LaGarda, Kristina Nesterova, and Andrey Zubarev. Simulating Russia's and Other Large Economies' Challenging and Interconnected Transitions. *National Bureau of Economic Research Working Paper*, 2015a.

Seth G Benzell, Laurence J Kotlikoff, Guillermo LaGarda, and Jeffrey D Sachs. Robots are us:

Some economics of human replacement. Technical report, National Bureau of Economic Research, 2015b.

Seth G Benzell, Laurence J Kotlikoff, Guillermo Lagarda, and Jeffrey D Sachs. Robots Are Us: Some Economics of Human Replacement. *NBER Working Paper*, 2016.

Seth G. Benzell, Erik Brynjolfsson, Frank MacCrory, and George Westerman. Identifying the multiple skills in skill-biased technical change. *Working Paper*, 2019.

Seth G. Benzell, Laurence J. Kotlikoff, Guillermo Lagarda, and Victor Yifan Ye. Simulating u.s. business cash flow taxation. 2020. URL https://kotlikoff.net/wp-content/uploads/2020/10/Simulating_Business_Cash_Flow_Taxation-2.pdf.

Seth G Benzell, Laurence J Kotlikoff, Maria Kazakova, Guillermo LaGarda, Kristina Nesterova, Victor Yifan Ye, and Andrey Zubarev. The future of global economic power. Working Paper 30556, National Bureau of Economic Research, October 2022.

Elise Brezis, Paul Krugman, and Daniel Tsiddon. Leapfrogging in international competition: A theory of cycles in national technological leadership. *American Economic Review*, 83(5): 1211–19, 1993. URL https://EconPapers.repec.org/RePEc:aea:aecrev:v:83:y:1993:i:5:p:1211-19.

Jesse Bricker, Lisa Dettling, Alice Henriques, Joanne W. Hsu, Lindsay Jacobs, Kevin B. Moore, Sarah Pack, John Sabelhaus, Jeffrey Thompson, and Richard A. Windle. Changes in u.s. family finances from 2013 to 2016: Evidence from the survey of consumer finances. *Board of Governors of the Federal Reserve System*, 2017. URL https://www.federalreserve.gov/econres/scf_2016.htm.

Erik Brynjolfsson, Tom Mitchell, and Daniel Rock. What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings*, 108:43–47, 2018.

Bureau of Labor Statistics. Current employment statistics (national). March 2022.

Bureau Census. Historical income tables: Households. 2019. URL https://www.census.gov/data/tables/time-series/demo/income-poverty/historical-income-households.html.

Herbert Dawid. Chapter 25 agent-based models of innovation and technological change. volume 2 of *Handbook of Computational Economics*, pages 1235–1272. Elsevier, 2006. doi: https://doi.org/10.1016/S1574-0021(05)02025-3. URL https://www.sciencedirect.com/science/article/pii/

S1574002105020253.

Md. Deen Islam. The geography of automation. *mimeo*, 2021.

David J. Deming. The growing importance of social skills in the labor market. *Quarterly Journal of Economics*, 132(4):1593–1640, 2017.

Jan Philipp Dietrich, Christoph Schmitz, Hermann Lotze-Campen, Alexander Popp, and Christoph Müller. Forecasting technological change in agriculture—an endogenous implementation in a global land use model. *Technological Forecasting and Social Change*, 81:236–249, 2014.

Barry Eichengreen. Secular stagnation: the long view. *The American Economic Review*, 105(5): 66–70, 2015.

Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. Gpts are gpts: An early look at the labor market impact potential of large language models, 2023.

Hans Fehr, Sabine Jokisch, and Laurence Kotlikoff. The Developed World's Demographic Transition - The Roles of Capital Flows, Immigration, and Policy. (10096), November 2003.

Hans Fehr, Sabine Jokisch, Ashwin Kambhampati, and Laurence J Kotlikoff. Simulating the Elimination of the U.S. Corporate Income Tax. 2013a.

Hans Fehr, Sabine Jokisch, and Laurence Kotlikoff. The world's interconnected demographic/fiscal transition. *The Journal of the Economics of Ageing*, 2013b.

Carl Benedikt Frey and Michael A Osborne. The future of employment: how susceptible are jobs to computerisation. *Retrieved September*, 2013. URL http://arche.depotoi.re/autoblogs/wwwinternetactunet{_}8a3fe3331e0ad7327e18d9fe6ec3f0ad04dcea58/media/3722fa7d.The{_}Future{_}of{_}Employment.pdf.

Don Fullerton and Diane Lim Rogers. Lifetime Effects of Fundamental Tax Reform. *Economic Effects of Fundamental Tax Reform*, pages 321–354, 1996.

Maarten Goos, Alan Manning, and Anna Salomons. Explaining Job Polarization in Europe: The Roles of Technology, Globalization and Institutions. *CEP Discussion Papers*, 2011. URL http://ideas.repec.org/p/cep/cepdps/dp1026.html.

Nezih Guner, Remzi Kaygusuz, and Gustavo Ventura. Income taxation of u.s. households: Facts and parametric estimates. *Review of Economic Dynamics*, 17(4):559–581, 2014. URL https://EconPapers.repec.org/RePEc:red:issued:11-141.

Hans A. Holter, Dirk Krueger, and Serhiy Stepanchuk. How do tax progressivity and household heterogeneity affect laffer curves? *Quantitative Economics*, 10(4):1317–1356, 2019. doi: https://doi.org/10.3982/QE653. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/QE653.

ILOStat. Statistics on the working-age population and labour force. 2020. URL https://ilostat.ilo.org/topics/population-and-labour-force/.

IMF. Imf investment and capital stock dataset, 2019. 2019a.

IMF. Balance of payments and international investment position statistics. 2019b.

International Monetary Fund. Government Finance Statistics. *IMF*, 2017.

IRS. Individual income tax rates and tax shares, 2017. 2019. URL https://www.irs.gov/statistics/soi-tax-stats-individual-income-tax-rates-and-tax-shares.

Nir Jaimovich, Itay Saporta-Eksten, Henry Siu, and Yaniv Yedid-Levi. The macroeconomics of automation: Data, theory, and policy analysis. *Journal of Monetary Economics*, 2021.

Kenneth L. Judd. Redistributive taxation in a simple perfect foresight model. *Journal of Public Economics*, 28(1):59–83, 1985. ISSN 0047-2727. doi: https://doi.org/10.1016/0047-2727(85)90020-9. URL https://www.sciencedirect.com/science/article/pii/0047272785900209.

Loukas Karabarbounis and Brent Neiman. The global decline of the labor share. *The Quarterly journal of economics*, 129(1):61–103, 2014.

Lawrence F Katz and Robert A Margo. Technical change and the relative demand for skilled labor: The united states in historical perspective. In *Human capital in history: The American record*, pages 15–57. University of Chicago Press, 2014.

John Maynard Keynes. Economic possibilities for our grandchildren. In *Essays in persuasion; as reproduced in 2010 edition.*, pages 321–332. Springer, 1930.

Florian Knobloch, Hector Pollitt, Unnada Chewpreecha, Richard Lewney, Mark AJ Huijbregts, and Jean-Francois Mercure. Ftt: Heat—a simulation model for technological change in the european residential heating sector. *Energy Policy*, 153:112249, 2021.

Anton Korinek and Joseph E Stiglitz. *14. Artificial Intelligence and Its Implications for Income Distribution and Unemployment*. University of Chicago Press, 2019.

Per Krusell, Lee E Ohanian, José-Víctor Ríos-Rull, and Giovanni L Violante. Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica*, 68(5):1029–1053, 2000.

Karl Marx. *Capital: volume one.* Courier Dover Publications, 1867. 2019 printing in English translation.

Andrew McAfee and Erik Brynjolfsson. *Machine, platform, crowd: Harnessing our digital future.* WW Norton & Company, 2017.

Ulrich K Müller, James H Stock, and Mark W Watson. An econometric model of international long-run growth dynamics. Technical report, National Bureau of Economic Research, 2019.

NCES. Statistical compendium. Technical report, National Center for Education Statistics, 2019.

Jaeger Nelson and Kerk Phillips. Macroeconomic effects of reducing oasi benefits: A comparison of seven overlapping-generations models. *National Tax Journal*, 2019.

William D Nordhaus. Are we approaching an economic singularity? information technology and the future of economic growth. Technical report, National Bureau of Economic Research, 2015.

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

Jeffrey D Sachs, Seth G Benzell, and Guillermo LaGarda. Robots: Curse or blessing? a basic framework. Working Paper 21091, National Bureau of Economic Research, April 2015a. URL http://www.nber.org/papers/w21091.

Jeffrey D Sachs, Seth G Benzell, and Guillermo LaGarda. Robots: Curse or blessing? a basic framework. Technical report, National Bureau of Economic Research, 2015b.

Robert Tibshirani. Regression shinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B (Methodological)*, 58(1):267–288, 1996.

U.S. Bureau of Labor Statistics. Multifactor Productivity and Related Measures. URL http://www.bls.gov/mfp/mprdload.htm.

Wikipedia contributors. Imagenet — Wikipedia, the free encyclopedia, 2021. URL https://en.wikipedia.org/w/index.php?title=ImageNet&oldid=1021861077. [Online; accessed 9-May-2021].

World Bank. Fossil Fuel Rents, Natural Resource Rents, 2021. Data Retrieved from World Development Indicators, http://data.worldbank.org/indicator.

Joseph Zeira. Workers, Machines, and Economic Growth. *The Quarterly Journal of Economics*, 113(4):1091–1117, 1998a. URL http://ideas.repec.org/a/oup/qjecon/

v113y1998i4p1091-1117..html.

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

Hernando Zuleta. Factor saving innovations and factor income shares. *Review of Economic Dynamics*, 11(4):836–851, 2008.