Simulating the Global Effect of Transformative AI: Growth, Welfare, Economic Power, and Policy Responses

Professor Seth Gordon Benzell
Dr. Victor Yifan Ye
Digital Fellows, Stanford Digital Economy Lab

March 2024

STANFORD
DIGITAL
ECONOMY
LAB

Executive Summary

Transformative AI will revolutionize production technologies in virtually all industries, expanding the range of tasks that can be automated. Consequently, it has the potential to dramatically decrease the share of GDP paid to labor, and increase the returns to capital.

What if we succeed in increasing the rate of automation by 10 times or more? How much would quality of life improve for people alive today? What are the implications for wealth creation and inequality within and across countries, skill groups, and generations? And how can economic policy ameliorate downsides and boost the benefits of transformative AI?

In this white paper, we utilize a large-scale, global macrosimulation model we developed to explore these questions quantitatively. The model features seventeen regions, containing over 150 countries comprising 99 percent of global population and 98 percent of GDP. Simulated agents comprise three skill groups -- high, middle, and low -- that are more or less abundant depending on the region. Each agent's lifecycle is calibrated to comprehensive micro data, featuring labor force entry, child-rearing, retirement, idiosyncratic mortality, and region-specific, realistic tax and transfer policies. We investigate, using this model, the consequences of alternate technological growth scenarios (relative to a baseline of 'business as usual' growth). We further assess potential fiscal reforms.

We consider a ten-fold increase in the pace of technological change. By 2050, this technology will increase U.S. GDP by 71 percent. However, it leaves global GDP virtually unchanged. Countries with relatively expensive labor and cheap capital, such as the U.S., adopt the most advanced technologies immediately. Countries with expensive capital and cheap labor delay adoption for the opposite reason. Increased investment demand and dis-saving in technology-adopting regions leads to much higher real interest rates. Growth is reduced in non-adopting regions because of automation, as investors depart for more exciting shores. Within frontier-technology adopting regions, faster automation boosts income and welfare of high-skilled workers and those being born today in almost all scenarios. However, the extent to which currently living middle- and low-skill workers benefit depends on scenario-specific assumptions.

We evaluate UBI and Universal Basic Capital (UBC) programs designed to address inequality caused by transformative AI. The benefits of a UBI depend heavily on how it is financed. A UBI financed with corporate income taxes could be disastrous for a unilaterally implementing country, as AI-empowered businesses will simply depart, weakening the tax base. A $5,000 UBI financed with a progressive income tax or a Value Added Tax (VAT) would raise the welfare of the low-skilled considerably, with modest impact on GDP growth. A UBC program designed to boost the savings of the working class increases the welfare of low-skilled millennials and gen-xers, but has a negative effect on welfare and growth in the long-run because of deadweight loss from taxes and transfers.

Our research agenda of developing, applying and evaluating our global automation models is currently unfunded, and we are looking for grant, consulting and partnership opportunities.

Bios

Seth Benzell

sethgbenzell.com | @sbenzell

Seth G. Benzell is an Assistant Professor of Management Science at the Argyros School of Business and Economics at Chapman University. Before coming to Stanford and Chapman University, Seth was a postdoctoral associate at the MIT Initiative on the Digital Economy. Seth received his Ph.D. in Economics from Boston University in 2017.

Seth's work is in the economics of automation, digitization, and networks. He is also interested in public economics generally. His work has been published in AEJ: Applied Economics, PNAS, Vanderbilt Law Review, Sloan Management Review, and other peer-reviewed and non-peer-reviewed outlets. Seth has presented his research at the US Capitol, and as an expert for a US international public diplomacy mission. His research and op-eds have been featured in the New York Times, NPR Marketplace, The Hill, War on the Rocks, and other popular outlets. Current projects with other Stanford DEL members focus on measuring skill-biased technical change, taxation and regulation of digital platforms, measurement of network effects, and the macroeconomic implications of progress in artificial intelligence.

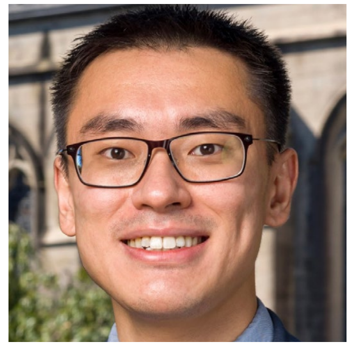

Victor Yifan Ye

yifanye.com

Victor Yifan Ye is a research scientist at Opendoor Technologies and a fellow at Boston University and Stanford DEL. His research focuses on computational economics and large-scale, micro-founded simulation methods, with applications in housing economics, public finance, macroeconomics, and the economics of artificial intelligence. Victor received his Ph.D in Economics from Boston University, and an M.S. in Statistical and Economic Modeling from Duke University, where he also received a B.S. in Economics and a B.A. in Philosophy.

Victor's publications have been covered by major media outlets including CNBC, The Wall Street Journal, Bloomberg, The Hill, The New York Times, Barron's, and Forbes. He has also been invited to present his work at the National Bureau of Economic Research, the National Tax Association, the Federal Reserve Bank of Atlanta, Boston, and Kansas City, the Gaidar Institute, the Congressional Budget Office, the Monetary Authority of Singapore, and the Hong Kong Monetary Authority.

Linkedin | NBER | Google Scholar

Introduction

The development of AI technologies and their economic impact are highly uncertain. Some, though, predict the imminent, ubiquitous deployment of powerful technologies that automate human labor in a variety of fields. Whatever the likelihood of this outcome, it is crucial to start preparing for this possibility immediately. Scenario analysis, even if the relative likelihood of scenarios remains unknown, is possible and wise.

To quantitatively assess the impact of advanced AI, we[1] developed a large-scale, multi-region computable general equilibrium, overlapping generations (OLG) model of the global economy focused on the international, distributional, and government fiscal impacts of new technologies. This model allows us to simulate different technological and policy scenarios, and see how they connect to political and economic goals.

For additional details, our Stanford Digital Economy Lab working paper, [Simulating Endogenous Global Automation,] is the most complete summary of our model.[2]

Model Elements

The model features 17 regions (e.g. the U.S., China, Russia, Western Europe, sub-Saharan Africa), each with realistically modeled demographics featuring individuals between the ages of 0 and 100 in each year. The model has over 450 demographic, fiscal, preference, and technology parameters, each of which is precisely calibrated using UN, IMF, and other macro data as well as microfindings on the marginal effect of new technologies (Acemoglu et al. 2020).

Each region has three distinct sectors: a household sector whose agents work, save, consume, and enjoy leisure, a government that taxes, transfers, and spends, and a business sector, which hires labor, rents capital, produces output that can be consumed or invested, and decide whether to use new (capital and high-skill biased) technologies or legacy technologies. Regions also receive fossil fuel rents that accrue to both private investors and governments. The seventeen regions are listed below.

The 17 Economic Regions of the Model

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 and the 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		

1- Professor Seth Gordon Benzell and Dr. Victor Yifan Ye of S-DEL, with the help, in particular, of Dr. Guillermo Lagarda and Professor Laurence Kotlikoff of Boston University
2 - You may also be interested in this VoxEU article explaining our research. The remainder of this white paper draws on both our working paper and the VoxEU article, occasionally verbatim.

Model Elements

How We Model AI and Automation

Technology can advance in many ways. It might mean the invention of new products, a way to overcome a previous obstacle to trade, or a new, more efficient way of creating an existing product. For the purpose of this study, we focus on two ways AI might change the macroeconomy. First, it can boost the overall economic output of the economy for the same amount of labor and capital inputs. Second, it can change the relative demand, and therefore shares of income, for capital versus labor, as well as demand for one skill type of labor versus another.

Our baseline scenario assumes steady productivity growth benefiting all regions of the world and additional catchup productivity growth for developing regions. The respective magnitudes of these effects are calibrated on UN and IMF economic and demographic projections, as well as the productivity growth estimates of Müller et al. (2019).

In our transformative AI scenarios, businesses gain access to more advanced production technologies over time. These new technologies are assumed to be capital intensive in the sense that, when utilized, businesses spend a higher share of costs on capital. This is how we model the advance of automation technology – firms getting the choice to produce using a more capital intensive technology. In aggregate, firms choosing to use the advanced technology increase the share of national income paid to capital, and decrease the share to labor.[3]

Businesses choose whether to adopt AI technologies. The extent to which businesses benefit from adopting the new technology depends on macroeconomic conditions and assumptions about how the technology works. Our estimates of the latter are based on returns to automation estimated in Acemoglu et al. (2020). Whenever capital is relatively expensive and labor relatively cheap, firms use legacy technologies. Whenever capital is cheap and labor expensive, firms particularly benefit from the new technology. The figure below shows when regions adopt the frontier automation technology in our baseline scenario:

Choice of technology by year and region

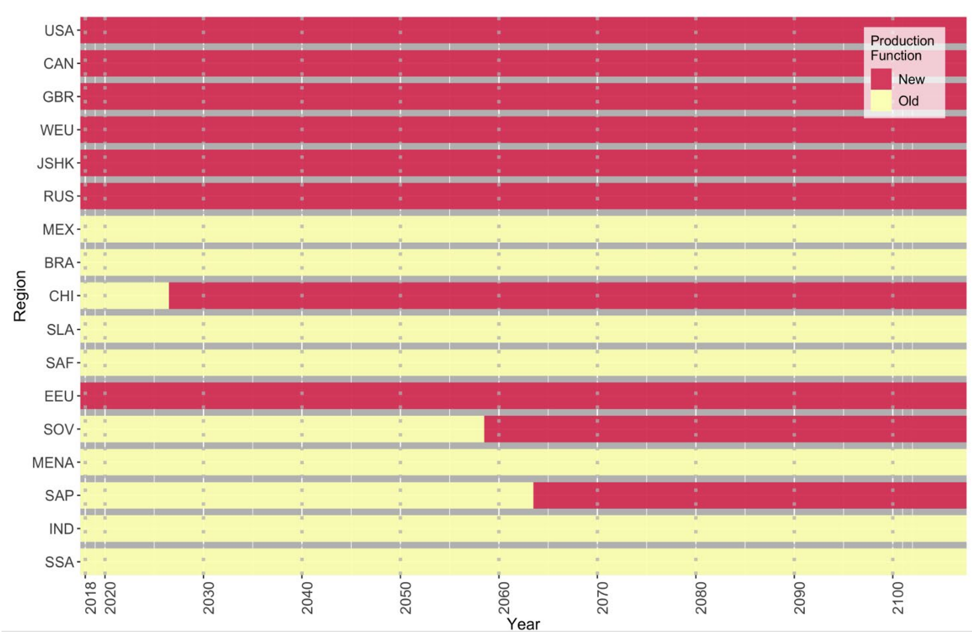

3-Decreasing labor shares of national income in the developed world has been a much noted phenomenon, and is often attributed in part to automation technology (Karabarbounis and Neiman, 2014).

Scenario Selection

Scenario Selection

We consider several scenarios concerning the aggregate economic impact of transformative AI

Technology Scenarios

10x Tech Change

• In our main technology scenario, we ask what would happen if transformative AI were to enable a much faster rate of automation. We select the magnitude of this shock to multiply the 'business as usual' decrease in the labor share of income by ten times over the next thirty years in the U.S.

• The figure below shows shares of national income for capital versus labor under the business as usual and 10x Tech Change scenarios. By 2060, Capital's share of income grows to 60 percent of GDP, up from today's rate of approximately 34 percent. We also present, in some figures, the results of the impact of a less dramatic 5x version of this scenario.

Shares of U.S. GDP paid to capital and different types of labor under the 10x Tech Change scenario

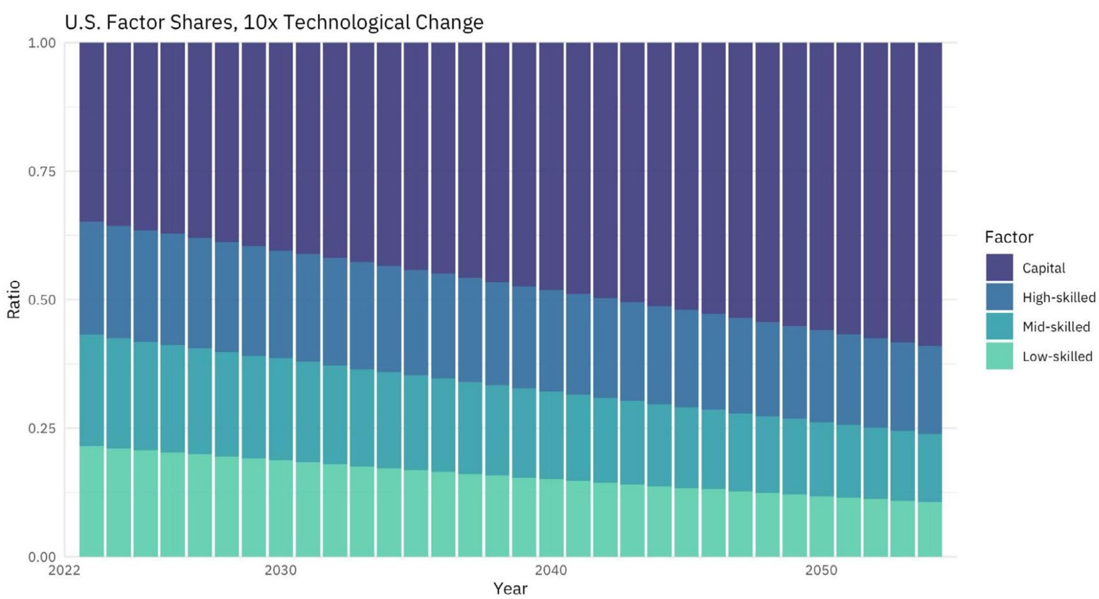

Skill Bias

In scenarios with this modifier, we additionally assume that the new automation technologies are strongly skill-biased. That is, beyond increasing capital's share of total income, adopting these technologies leads to a large increase in the high-skilled share of national income. This is a plausible scenario if one believes that "scarce architects" are needed to take full advantage of the new technologies, or if one believes that tasks susceptible to being automated by transformative AI are disproportionately performed in low-paying occupations.

Scenario Selection

Within a birth-cohort and region, individuals are in one of three skill groups:

- **High Skill** - The top highest-earning 4 percent of the population
- **Middle Skill** - The next 20 percent
- **Low Skill** - The bottom 76 percent

Under the + Skill Bias scenario, low- and mid-skilled workers lose ground to both automation and their high-skilled coworkers. By 2060, the lowest-income 76 percent of Americans see their share of national income decrease from 22 percent of GDP to only 4 percent.

U.S. Factor Shares Under "+ Skill Bias" Scenarios

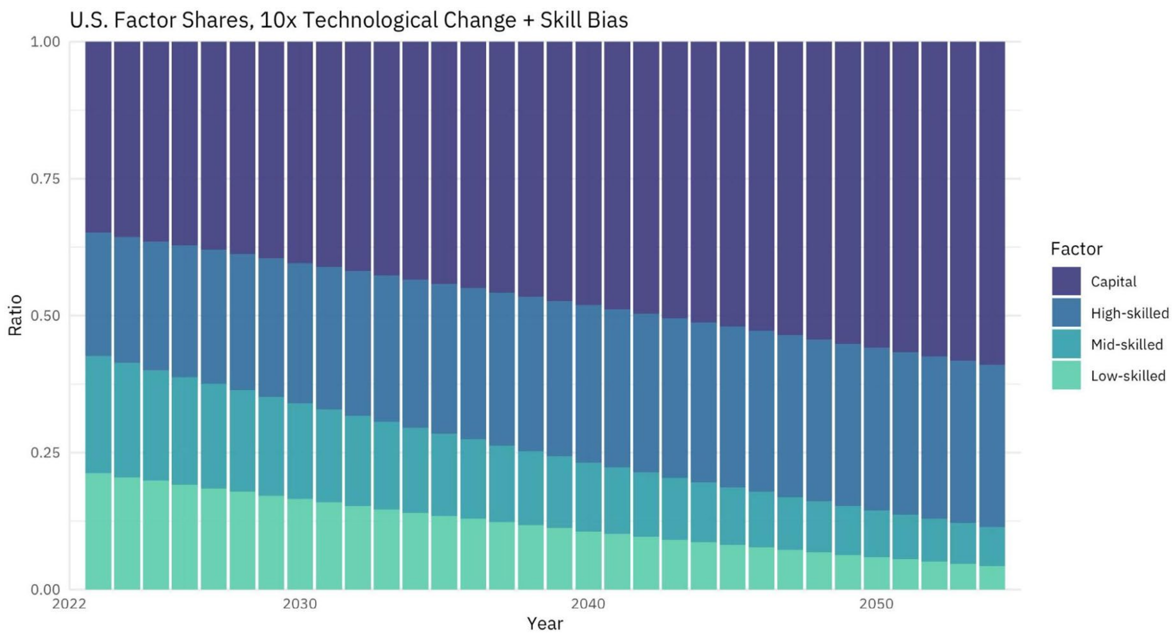

Better AI

In the scenarios described above, we assume that the arrival of automation technologies brings the same impact on productivity as measured in recent data. However, to the extent that "this time is different", the introduction of generative AI-based automation may entail a large boost to productivity beyond historical experience. In these scenarios, we assume that adopting the frontier AI technology yields an additional 0.7 percent total factor productivity (TFP) growth per year, on top of any innate advantage from adoption. This is enough to roughly double U.S. GDP, versus the 10x tech change scenario alone, by 2100.

4- See our related S-DEL Working Paper: Digital Abundance Meets Scarce Architects - Implications for Wages, Interest Rates, and Growth

Policy Scenarios

Policy Scenarios

We consider government fiscal responses to these scenarios. Outside of these new policies, we assume that governments follow a business-as-usual tax and spending policy as described in our working paper.

UBI of $5,000 per year (in 2022 U.S. Dollars)

- Paid equally to all adults annually
- Funded, alternately, through:
    - General Gov. Revenue: Increasing all personal taxes, maintaining current degree of progressivity
    - Income Tax: A new higher income tax rate for top earners
    - Consumption Tax: A flat, national value-added tax (e.g. a VAT) on all consumption

UBC Savings Boosting Policy[5]

In a future where production is much more dependent on capital and less so on labor, a Universal Basic Capital (UBC) policy could ensure that people of all ages and skill-groups have access to capital income.

In this scenario, we imagine an incentivized saving program designed to boost the aggregate savings for young low- and mid-skill workers under 35 years old by 10 percent of GDP by 2030 —approximately $3 trillion.

The savings boost is achieved with a budget-neutral tax incentive for young workers. The policy also includes a subsidy of $500 per year for young savers to approximate the benefit of exempting capital gains on incentivized savings.

To be precise: All low and mid-skill Americans (the bottom 96 percent of earners) between ages 21 and 35 receive a $500 per year subsidy. This number is roughly set to the annual tax advantage of an investment account that is fully tax-exempt up to ~$25k. This lasts until 2030, after which the policy is assumed to sunset. The funding for this subsidy comes out of general (federal, state, local combined) tax revenues. We calibrate the mix to be roughly 4:1 between income and consumption taxes. In addition, to make sure that young people actually save, we penalize them from withdrawing early. The way we represent this in the model is as a proportional tax on the consumption of young workers until age 35, reimbursed lump-sum and, consequently, is revenue neutral. We set the rates to induce an additional 10 percent of GDP’s worth of savings by 2030 - roughly $50,000 (2022 dollars) per young adult.

5- This policy is inspired, in part, by a number of UBC proposals of Nathan Gardels and others.

Technology Scenario Results

Technology Scenario Results

Impact on Global and U.S. Growth

Each of the technological change scenarios is massively positive for U.S. and Global GDP. For the U.S., the impact is more immediate. In the 10x Tech Change scenario, GDP rises 75 percent higher than the baseline level in 2050 and 117 percent higher in 2080. The “+Skill Bias” assumption cuts this growth by half, as scarce high-skilled workers increasingly hold back growth. The “+Better AI” assumption is even more favorable to the U.S., tripling GDP above baseline by 2090.

Global and U.S. GDP by Technology Scenario

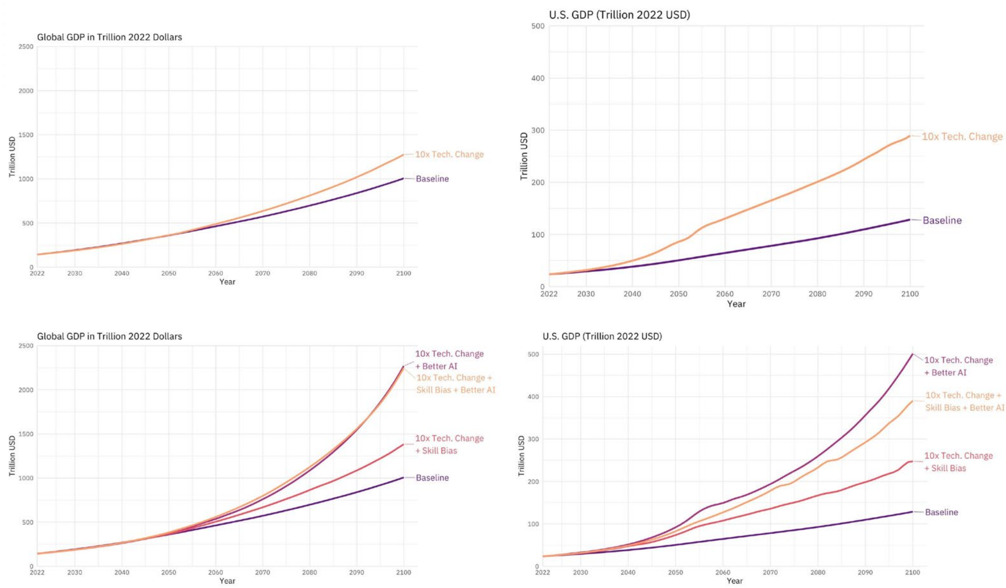

Impact on Global Rate of Return

It takes time for transformative AI to boost global growth. The main reason is scarce investment supply. Higher interest rates are one of the most dramatic implications of transformative AI. Under our ‘business as usual’ projections, a global ‘savings glut’ from the increasingly wealthy Asian and Middle Eastern middle class is enough to bring rates of return down to under 3 percent over the next three decades[6]

Under our transformative AI scenarios, however, two forces dramatically increase rates of return over the remainder of the century.[7] On the demand side, transformative AI is projected to substantially increase investment demand. On the supply side, people run down their savings in the short-term, desiring to smooth consumption in anticipation of future abundance created by AI. Therefore, under transformative AI scenarios, global capital formation is actually slower through mid-century due to the announcement of transformative AI. After mid-century, the overall productivity effect dominates, raising capital stocks and boosting the global economy.

6- For more on the global saving glut, see Barsk and Easton (2021) , Eichengreen (2015), and Bernanke (2015)
7- Previous research has also highlighted the connection between anticipated advances in AI and contemporary interest rates, as mediated by the two effects we consider. This has led Basil Halperin (joining S-DEL in 2024) to argue that AGI is not imminent.

Technology Scenario Results

Global rate of return on investment and global capital accumulation rate by year and technology scenario

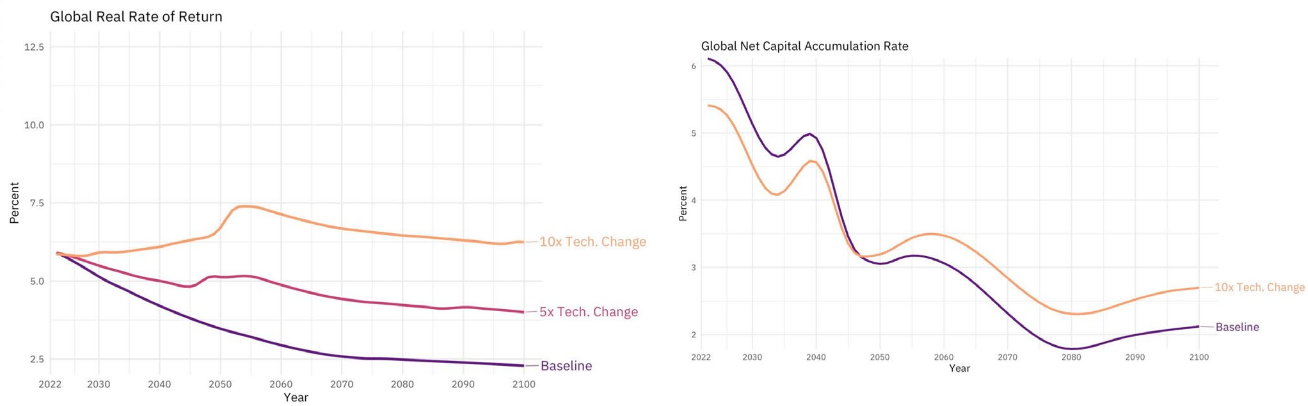

<u>Impact on Inequality Across Countries & The International Balance of Economic Power</u>

While regions adopting the newest AI technologies see massive returns, countries that don't keep up technologically are harmed by transformative AI.[8] This is due to the increased global cost of investment. China, for example, cannot adopt frontier technologies until 2046 due to relatively cheap labor and high interest rates. This effect has striking implications for the international balance of power. Over the next 30 years, transformative AI of the type we model would help the U.S. retain economic hegemony, in large part by constraining Chinese, Indian, and African catchup growth.

Real GDP growth rates by region, year, and technology scenario

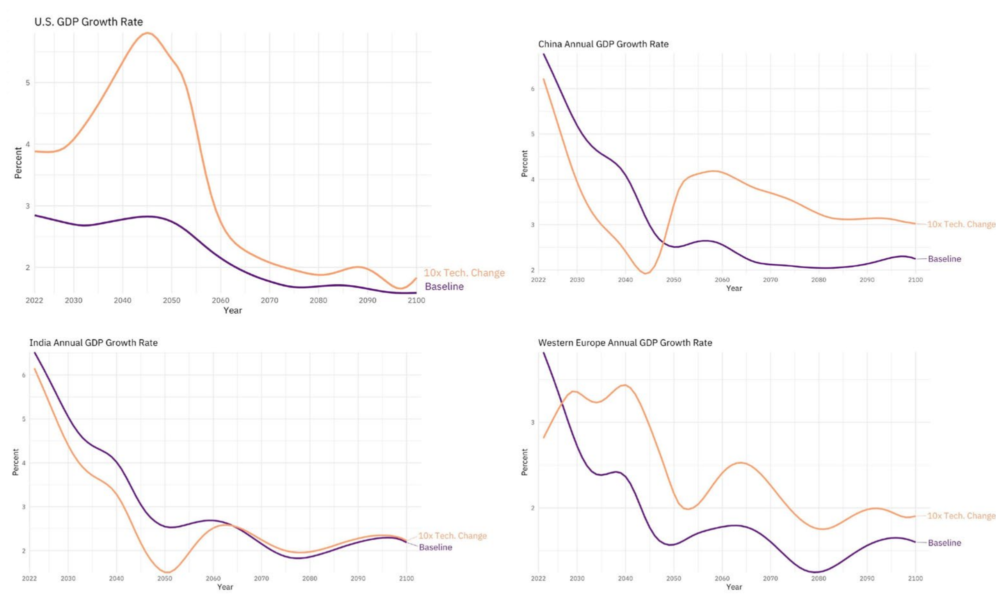

8- The theoretical possibility of automation to increase inter-country inequality because of these mechanisms was first proposed by (Zeira, 1998)

Technology Scenario Results

<u>Impact on Wages, Welfare and Inequality within the U.S: Within and Across Generations</u>

In the 10x Tech Change scenario, welfare gains from AI are roughly similar for all groups. A low-skilled individual born in 2020 is made 49 percent better off relative to baseline. A high-skilled individual born in the same year is 70 percent better off. However, this relatively balanced welfare gain is juxtaposed with a much more unequal increase to income. Between 2022 and 2060, the low-skilled (the bottom 76 percent of earners) average annual labor income is projected to only increase by $10,000 (in real 2022 dollars). On the other hand, the annual labor income of the high-skilled (the top 4 percent) increases by $350,000.

If transformative AI is associated with a large increase in labor income inequality, perhaps because AI is not able to automate the highest-compensated jobs (i.e. our "+ Skill Bias" scenarios) the inequality consequences are much starker. Under this assumption, all low-skilled individuals born before 2060 are made worse off, as well as mid-skilled individuals born before 2035.

Wage and welfare for a representative U.S. person in each skill group by year or birth-year and technology scenario

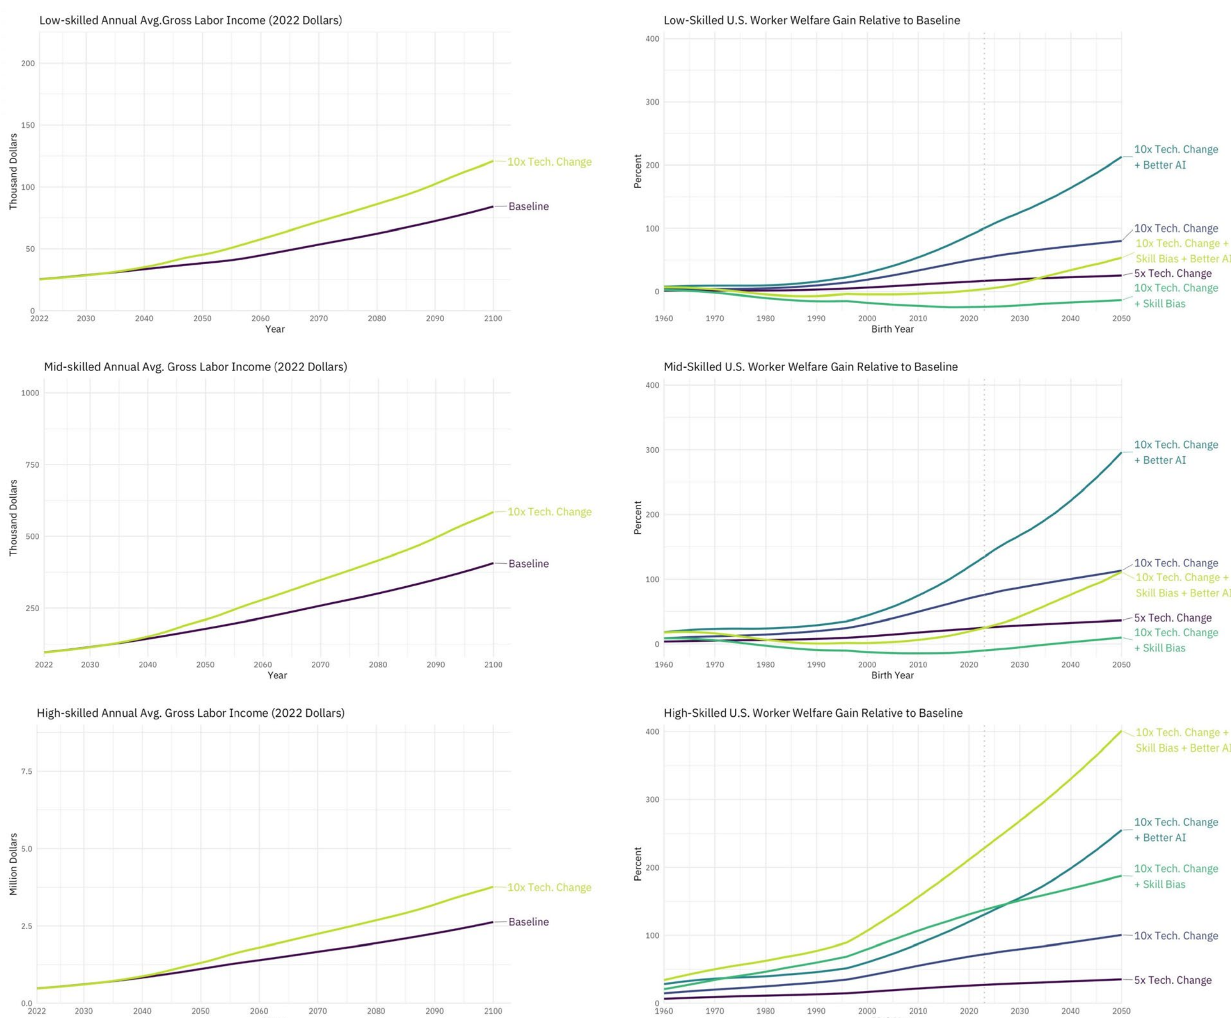

Policy Scenario Results

UBI Policies: Funding Type Matters!

UBI policies have been proposed to more broadly share the economic gains of rapid technological change. In this section, we consider a permanent, tax-exempt, revenue-neutral, $5,000 annual transfer to all U.S. persons in real 2022 dollars, financed in four different ways. The UBI is assumed to supplement, not replace, existing U.S. welfare transfers (e.g. food stamps, disability insurance), the fiscal impact of which grows with population and GDP.

The figures below show the welfare impact of the UBI under three different financing regimes: the current U.S. personal tax mix, a progressive income tax, and a VAT/consumption tax. Any of these policies would significantly raise the welfare of the bottom 76 percent of Americans. Low-skilled Americans favor the progressive income tax because the other taxes would be partially incident on them. The mid- and high-skilled are made worse off by the UBI. Regardless of the financing regime, the UBI's negative effect on GDP is insignificant. By 2050, the UBI policies lower GDP by approximately 2 to 3 percent.

Unplotted in the figures below is the effect of a corporate tax financed UBI. Because the capital-intensive businesses in our transformative AI simulations are globally mobile, growth rates are sensitive to corporate income taxation. If a country unilaterally raises corporate tax rates, this leads their tax base to decrease, making all groups worse off. Further, the corporate tax base is small relative to consumption and labor income, both of which amount to roughly a third of U.S. GDP. Hence, extraordinarily high rates of corporate taxes are required to fund a relatively modest UBI.

UBI Policy Impacts

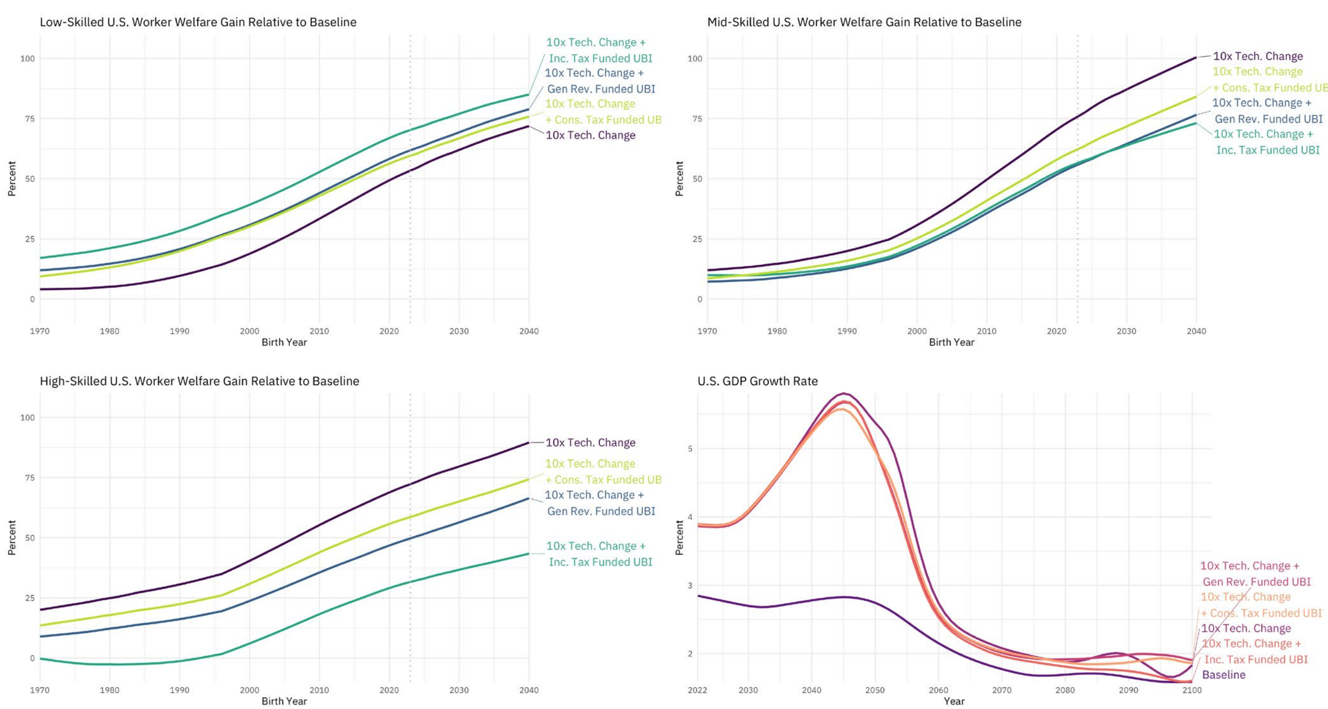

Policy Scenario Results

UBI Policy Impacts (continued)

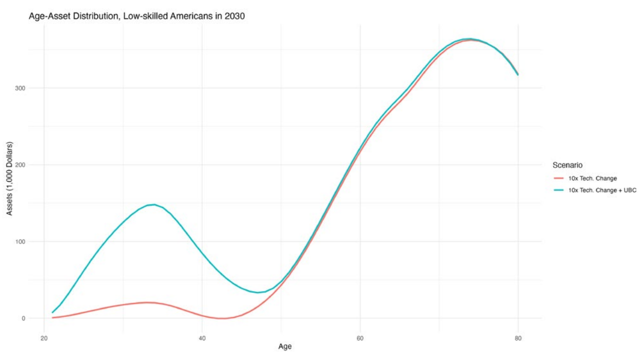

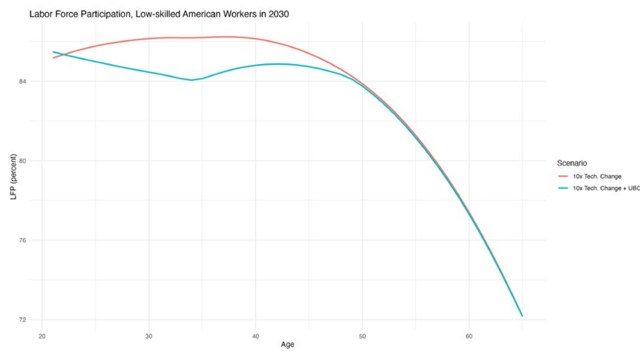

UBC Policy

A revenue-neutral UBC program designed to boost the savings of the mid- and low-skilled increases the welfare of low-skilled workers born between 1980 and 2009. However, it has a negative effect on welfare and long-run growth because of deadweight loss from the taxes and transfers. By 2050, the UBC lowers U.S. GDP by 0.5 percent, and by 2100 it lowers GDP by 5 percent.

One might be surprised by the negative effect on growth: shouldn't additional savings boost capital accumulation and output? The reason is that additional U.S. savings are a drop in the international bucket compared to massive wealth accumulation in Asia and the Middle East, having little impact on aggregate international capital formation. Growth is also reduced by a reduction in low-skill labor supply, primarily due to higher capital incomes in their middle age.

One might also be surprised by the relatively modest effect on low-skill welfare, and negative effects on the middle skilled. Partly this is because the UBC, while raising consumption overall, makes individuals save more when young in a way contrary to their preferences. If you think young people are irrationally impatient, then you should put less weight on this welfare analysis, in favor of a focus on the increases in assets, leisure and consumption. Another reason is that this policy is financed with general government taxes which are somewhat incident on low and mid-skilled Americans. A policy financed by a tax on only top percentile workers would favor the mid-skilled more.

A few tweaks would make the UBC more effective. First, while much of additional U.S. savings spill over to foreign countries, the opposite would be the case—with the U.S. benefiting more—if all countries were to adopt similar incentivized savings programs[9]. Second, countries experiencing autarky, due to international sanctions or by choice, would also benefit more from a UBC. Third, the policy combines incentivized savings with a *de facto* transfer to young workers. The transfer requires an increase in average net personal tax rates of about 4 percent, creating significant deadweight loss. A UBC with the subsidy component eliminated would be more beneficial for U.S. GDP.

9- International coordination on raising corporate tax rates would also make a corporate tax-funded UBI more attractive. Achieving a global corporate minimum tax has been a major diplomatic effort.

Conclusion

Conclusion

In this white paper, we simulate the impact of transformative AI on macroeconomic and fiscal outcomes. Specifically, we consider variations on a scenario where AI increases the pace of automation by 10x, increasing capital's share of GDP to 60 percent in 2060. This scenario strongly benefits the U.S., increasing GDP by 75 percent in 2050. The global effects are relatively muted as scarce capital bids up interest rates, leading some countries to fail to adopt frontier technologies and, consequently, fall behind. Inequality also increases within countries, although the magnitude depends heavily on the technology scenario. We also consider UBI and UBC policies, finding that a UBI of $5,000 a year would have a positive effect on the welfare of the low-skilled, at only moderate costs to growth.

We hope that these simulations of a possible economic future transformed by AI are helpful to you in thinking about such forces. Future simulations could integrate alternate assumptions about AI's impact on the economy, including increased life expectancy, changes to time preferences, or advantaging certain regions, occupations, types of capital, or age-cohorts.

We plan to consider more extreme AI scenarios, including allowing for technologies where capital's share of income increases to 90% or more. A challenge in simulating these scenarios is that they are incompatible with projected government fiscal policies and goals. In the short-term, certain non-automation adopting countries (e.g. Mexico, Sub-saharan Africa) are unable to afford planned welfare payments due to massive capital flight and reduction to GDP. The long-term complications are less ominous: global GDP growth in these scenarios is so rapid that governments are over-flush with tax revenues, requiring ancillary assumptions about handling excess fiscal revenue.

If you have questions about our simulations, want to learn more, or are interested in a collaboration, please reach out. We are excited to continue building, applying, and evaluating tools for projecting the global, long-run macroeconomic, and fiscal impacts of transformative technological change

References

Acemoglu, D, C Lelarge and P Restrepo (2020), "Competing with robots: Firm-level evidence from France", AEA Papers and Proceedings 110: 383–88.

Acemoglu, D and P Restrepo (2020), "Robots and jobs: Evidence from US labor markets", Journal of Political Economy 128(6): 2188–2244. Related VoxEU column here.

Barsky, R., & Easton, M. (2021). The global saving glut and the fall in US real interest rates: A 15-year retrospective. Economic Perspectives, 1, 2021.

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

Karabarbounis, Loukas, and Brent Neiman. "The global decline of the labor share." The Quarterly journal of economics 129, no. 1 (2014): 61-103.

Müller, U. K., Stock, J. H., & Watson, M. W. (2019). An econometric model of international long-run growth dynamics (No. w26593). National Bureau of Economic Research.

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