NBER WORKING PAPER SERIES

WILL AI GROWTH REACH THE DEVELOPING WORLD?

Daniel Björkegren
David Lagakos
Mitchell VanVuren

Working Paper 36035
http://www.nber.org/papers/w36035

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2026

For helpful comments we thank Jeff Berens, Oliver Hanney, Tarek Hassan, Chad Jones, and Joseph Levine. All potential errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research. Daniel Björkegren serves as a Nonresident Fellow at the Center for Global Development, has worked as a consultant for the World Bank, and has received research funding from AI companies.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w36035

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Daniel Björkegren, David Lagakos, and Mitchell VanVuren. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Will AI Growth Reach the Developing World?
Daniel Björkegren, David Lagakos, and Mitchell VanVuren
NBER Working Paper No. 36035
November 2026
JEL No. E27, O11, O41

ABSTRACT

This paper explores how economic growth driven by AI will impact the developing world. Our analysis is based on a growth model of a developed and developing economy where AI services are traded on the world market, and may substitute for or complement different labor inputs in production. Preferences allow for the possibility that demand for services provided by AI has an upper bound. The model predicts that developing countries may not benefit as much from AI-driven growth. First, they may lack complementary skills needed to fully utilize AI domestically. Second, they may be priced out of world markets due to their relatively lower productivity. We use our model to simulate how growth and welfare are affected by making workers more complementary to AI, subsidizing compute, speeding technology diffusion, and providing cash transfers. Numerical simulations of the model suggest that the potential for developing countries to be left behind – or having some workers lose in absolute terms – is substantial enough that researchers should study the impacts of AI on developing countries in real time.

Daniel Björkegren
Columbia University
dan@bjorkegren.com

Mitchell VanVuren
Vanderbilt University
mitchell.vanvuren@vanderbilt.edu

David Lagakos
Boston University
Department of Economics
and NBER
lagakos@bu.edu

1. Introduction

Artificial intelligence (AI) is improving at extraordinary speed. The price of electric lighting fell by a factor of 1,000 over the century following its introduction in the 1880s (Nordhaus, 1996); the price of GPT-3-level AI fell as much in just three years (Appenzeller, 2024). Earlier technological progress coincided with rising incomes in low-income countries but also widening gaps: between 1870 and 1990, the average per capita income of the world's poorest country grew by 60%, and in the richest country grew by 775% (Pritchett, 1997). Will economic growth that results from AI reach developing economies, or leave them behind?

The potential of AI to reshape labor markets and alter the course of economic growth has already become the focus of a large macroeconomics literature (see e.g. Aghion and Bunel, 2024; Acemoglu, 2025; Restrepo, 2026; Korinek, Jones, Sacher, Cotter and McCrory, 2026). Most of it has focused on advanced economies, such as the United States, which is sensible given that they are on the frontier of AI adoption. However, the more than half of the world's population that live in low- and middle-income countries are increasingly adopting AI themselves, in their workplaces and personal lives (Liu, Huang and Wang, 2025). And even those who do not themselves use AI will be impacted by it indirectly, as AI reshapes production, the products and prices in local and global markets, and scientific discovery.

Opinions range widely on whether AI will transform the developing world for the better or worse (or at all). AI could allow developing nations to better leverage their (largely) manual labor forces (Björkegren, 2026a; Kim, Tadepalli and Levine, 2026). It could also automate away their best export options, such as call centers (Mousa, 2026). Productivity increases from AI in advanced economies could lead to "winner-takes-all dynamics" that benefit only the economies that produce AI services while reducing economic opportunities elsewhere (Korinek and Stiglitz, 2021; Alonso et al., 2022). On the other hand, many low-income countries conduct little international trade, and have low domestic AI adoption: low-income countries account for only around 1 percent of the world's generative AI website traffic (Liu et al., 2025), which may suggest negligible impacts.

In this paper, we ask how AI might transform the economies of the developing world in the coming years, focusing on how they differ from, and interact with, advanced economies. Our analysis makes use of a growth model with two economies, representing the advanced world and the developing world. In the production of final goods, we allow some types of human capital to be more sub-

stitutable with AI (e.g. the ability to translate a specification from English into computer code) and some to be less substitutable (human capabilities that AI models may lack). We allow AI-intensive goods to have an arbitrary degree of substitutability with other goods. We also allow for satiation in preferences, where demand for either the AI-intensive or the non-AI-intensive good can reach a limit. For example, if AI completely automates legal services, one may still only need a lawyer's help once per year. We posit a world market for AI services, which allows the possibility that demand from the advanced economy affects equilibrium usage by the developing one. The two economies differ in their endowments of compute, their endowments of different types of human capital, and their overall productivity levels, which govern their relative income levels.

Using the model, we focus on several cases that are instructive for thinking about how AI could affect developing economies. To illuminate the simple forces in the model, we first shut down satiation in preferences and simulate the impact of continued growth in the productivity of AI. We show that the model admits optimistic scenarios, in which developing economies achieve sustained growth mirroring that of the advanced world, as well as pessimistic scenarios, in which developing economies are permanently left behind. The pessimistic scenarios occur when developing economies lack enough complementary inputs to make use of AI or, alternatively, when growth in the advanced economy is so rapid and AI's use in production is so broad that the advanced economy's expanding demand for AI services prevents developing economies from affording enough AI to achieve sustained growth from it. In these cases, income levels in the developing country are less affected by AI, despite the fact that AI productivity continues to increase, and the advanced and developing worlds diverge.

We next consider our model and allow for satiation in the AI-intensive sector. This captures the idea that AI may produce all the (say) legal services needed by society while a physical good such as food remains scarce. In this version of the model, the forces described above continue to be present but are opposed by the fact that satiation eventually drives down the price of the AI-intensive good. This diminishes the importance of the developing country's lack of skills and (potential) inability to afford AI, complicating the picture and leading to a wide variety of potential outcomes including both divergence and convergence. In the extreme case of world satiation in the AI-intensive good, growth in AI leads to substantially higher (average) incomes in both advanced and developing countries.

To assess the plausibility of these possibilities, we consider a range of values for substitution elasticities, degrees of satiation in preferences, and relative supplies of human capital that complements rather than substitutes for AI. Our approach is to draw from plausible distributions over these parameters and simulate the model's consequences over the next 25 years. We consider two paths for continued AI progress: one where the overall productivity of AI improves at a modest-but-rapid pace of 15% annually and one where it improves explosively at 50% annually. These values represent growth in the amount of output-relevant productive work that AI can perform, which has been much slower than improvements in performance on benchmarks or raw increases in usage (tokens) or data center capacity (Demirer et al., 2026; Chen and Stratton, 2026).[1] The goal of these exercises is not to make precise predictions about what will happen, but instead to describe a range of possibilities, illuminate the factors that determine the outcomes, and assess what the impacts of policies might be in different scenarios. The simulations also show which cases most warrant further study.

Many simulated scenarios produce divergence: in about half of our scenarios, the gap between the incomes of advanced and developing countries widens by at least 40 percent. These scenarios arise mostly from parameter values where developing economies lack complementary skills and these skills turn out to be particularly important for production; consequently, the developing country is less able to translate growth in AI to growth in output. Others arise when AI is sufficiently substitutable with labor, and growth in the advanced economy is sufficiently high, that the advanced country's demand for AI services prevents global prices from falling enough to sustain growth in the developing world.

We then turn to the distributional consequences of AI in developing economies. Developing economies have large low-education workforces (including subsistence farmers and microentrepreneurs), on which AI may have limited effects.

[1] Benchmark performance and token usage have increased at far faster rates. However, Chen and Stratton (2026) finds that an 8.43x increase in METR's time horizon benchmark plot corresponded to only a 1.74x increase in coding productivity. And both that study and Demirer et al. (2026) find that increases in coding productivity led to much lower increases in overall output. Brynjolfsson et al. (2026) finds that the average compensation required to forgo AI access for a month rose 27% between July 2025 and March 2026, also smaller than benchmark improvements. Korinek and McKelvey (2026) and Juniewicz and You (2026) document that global AI compute capacity has been growing at least 200% annually in the past few years (in terms of Nvidia H100 equivalents); such a trend may be continued if chips continue to become more efficient, but data center buildouts at that magnitude may hit physical constraints that prevent them from being sustained for the next 25 years.

Within the smaller share of workers in the knowledge-intensive occupations that are highly exposed to AI (Eloundou et al., 2024; Cerutti et al., 2025; Murugan et al., 2026), AI could raise inequality along the lines predicted in advanced economies. Workers whose tasks AI can perform see their wages decline in absolute terms even as aggregate incomes rise. Having skills complementary to AI does not always protect workers: if demand for their sector becomes satiated, workers whose skills are complementary to AI can also see wage declines (similar to Bessen, 2019). It is possible that affected cognitive workers include educated youth who have seen high unemployment and unrest in many developing country cities in recent years. Despite these distributional issues, across our simulated scenarios, our model predicts that wages rise on average, and for at least one worker group.

We simulate the impacts of four broad policies that have been discussed as ways to help the developing world to benefit more from AI driven growth. We model these in a stylized way and quantify them using parameter values we view as plausible, though still highly uncertain.

First, we simulate an increase in the supply of skills that complement AI. This can increase developing economy growth if inputs are gross complements in production, since then the returns to complementary human capital grow as the price of AI declines. There remain implementation questions: education may be a relevant lever, but it is not yet clear what skills will remain complementary to future AI and how to teach them at large scale. Additionally, this policy has smaller growth effects if AI can substitute for all types of labor, as in predictions of Artificial General Intelligence (AGI).

Second, we simulate a subsidy for AI adoption in developing economies. Subsidies might help offset the upward pressure on AI services by richer countries. Our simulation assesses the allocation cost of redirecting AI use. We find those allocation costs can be high: subsidies for AI use in the developing economy in our simulations can boost adoption but do not generate much catch up growth unless they are extremely large. In our model, subsidies inefficiently induce producers to replace human workers with AI in cases where humans would be more effective, and inefficiently shift AI inputs away from advanced economies, which use AI more efficiently.

Third, we simulate an increase in the rate at which general productivity diffuses to the developing country. Our model suggests that a large determinant of the growth of the developing country is how quickly it can adopt innovations from

the frontier economy (which may include business processes, management practices, and scientific innovations), and that speeding it up could increase developing economy growth and in some scenarios lead to convergence. It is an open question how diffusion could be effectively achieved since accelerating diffusion is less a concrete policy than a policy objective.

Finally, we consider large-scale transfers from advanced to developing economies – a global adaptation of universal basic income policies (O'Keefe et al., 2020; Yelizarova, 2025). In the simplest version of our model, transfers mechanically have no effect on production or AI use, but can be used to share consumption gains obtained by the advanced economy. Our simulations suggest that the magnitude of transfers necessary to offset AI-induced divergence in consumption levels between advanced and developing economies could be large, and may exceed the GDP of developing economies. In our model, transfers funded by ownership shares in the AI sector (Patel et al., 2026; Bostrom, 2025) may scale with the magnitude of AI's impact on the economy. We additionally consider targeting transfers to all residents of developing economies (universal basic income), as well as to specific groups. There may be a tension between insuring against disruption and alleviating pre-existing poverty, if the groups that experience wage losses due to AI are educated people like lawyers and software developers, rather than lower income people like subsistence farmers that are the more traditional target of aid.

We conclude that the stakes are sufficiently large that researchers should increase efforts to model and measure the impacts of AI in the developing world. We consider two directions for this agenda. First, we describe a measurement agenda, suggesting what could be measured to reduce uncertainty. While there are substantial measurement efforts in advanced economies, they remain limited in the developing world. We should monitor how AI capabilities evolve to complement or substitute the human capabilities in developing countries, whether demand for AI-produced output becomes satiated, and how quickly productivity gains transfer to developing countries. Second, we describe features omitted from our simple model that warrant additional exploration. Particularly notable is our limited treatment of international trade; future work should model richer interactions as comparative advantages are potentially reshaped by AI. Additional modeling may also clarify the key forces to monitor.

Related literature. There is an emerging theoretical literature considering how AI may alter growth in a single (advanced) country. Aghion et al. (2019) and Trammell and Korinek (2026) model the growth impacts of automating production and R&D while Korinek and Suh (2024) and Korinek et al. (2026) provide calibrated estimates of different scenarios. Rockall et al. (2025) consider how AI could shrink or expand within-country inequality, and make the point that high income workers who are more exposed to AI may benefit due to complementarity or capital returns. Acemoglu and Restrepo (2018) and Ide and Talamàs (2025) both develop detailed microfoundations for the process of AI-automation and study the implications for the structure of production.

Some papers consider how AI may affect the relationships between (predominantly wealthy) countries. Athey et al. (2026) consider AI as an imported factor, competition in AI, and resulting impacts on different countries while Barhoumi et al. (2026) consider qualitative scenarios (without a formal model).

Some studies consider AI impacts on low- or middle-income countries specifically. Korinek and Stiglitz (2021) theoretically consider automating physical labor not just cognitive labor, as well as the impact on demand for natural resources, which are often produced in LMICs. Alonso et al. (2022) consider how automation can widen income gaps between high- and low-income countries, focusing in particular on substitutability with labor. Several papers modify quantitative international models in order to make predictions for outcomes in a wide swath of countries (Benzell et al., 2021; Cerutti et al., 2025; Che et al., 2026). Ide and Talamàs (2026) consider specialization in knowledge and problem-solving tasks between advanced and developing countries while Liu (2024) considers the possibility that AI prevents developing countries from developing by reducing demand for skilled services. Imam and Temple (2026) take a development-accounting approach and argue that cross-country diffusion of AI-related productivity gains depends on the particular nature of its influence on costs.

Diffusion of productive technologies to developing countries may be slowed due to differences in inputs (Basu and Weil, 1998; Zeira, 1998; Acemoglu and Zilibotti, 2001; Comin and Mestieri, 2018; Acemoglu et al., 2026b). AI may be less appropriate for low capital economies, or could make frontier technologies more appropriate by providing previously missing skills. In that sense, countries may be able to import missing skills from AI rather than foreign humans (Antràs et al., 2006). Porzio et al. (2022) provide evidence that accumulation of human capital

increased the transition out of agriculture and drove structural change.

Most empirical estimates focus on wealthy countries. Many empirical estimates measure how exposed different occupations are to AI (Eloundou et al., 2024), and then aggregate up (Acemoglu, 2025; Aghion and Bunel, 2024). Other work studies how future productivity increases from AI can be estimated using returns to stock markets (Blumenfeld et al., 2026; Eisfeldt et al., 2023), or are reflected in interest rates (Chow et al., 2025; Andrews and Farboodi, 2025; Björkegren, 2026b). Surveys of AI use are now being collected regularly in advanced economies. These data show, for example, that roughly one in five Americans now uses generative AI almost every day (Das et al., 2026; Verasight, 2026), and that its adoption for work tasks is already widespread (Bick, Blandin, Deming and Schumacher, 2026).

There is less empirical evidence in low- and middle-income economies, and the effects there may differ. AI adoption rates are lower (Chatterji et al., 2025; Appel et al., 2025; Daepp and Slaughter, 2026; Iscenko et al., 2026), and workers' jobs are less exposed to AI on average (Cerutti et al., 2025; Murugan et al., 2026). In exposed occupations, job postings have declined less since the release of ChatGPT in developing economies than in advanced economies (Sakai et al., 2026). Evidence does not suggest a universal relationship between human capital and causal gains from AI (Brynjolfsson et al., 2025b; Otis et al., 2026). There is also some empirical evidence from previous waves of innovation; for example, de Souza and Li (2023) finds that robots imported into Brazil decreased employment and wages of low-skilled workers between 1997 and 2014.

2. A Simple Growth Model with AI

We consider two economies, advanced (R for "rich") and developing (P for "poor"), representing the world's most and least developed areas. Each economy has two sectors. One of them, which we index using A, may use labor inputs and AI in production. The other, indexed by N, uses only labor. We think of the A sector as the services and manufactured goods whose production process could be materially improved by using AI, and the N sector as activities where AI is not valuable in the production process. We are agnostic about which specific sectors in reality map into our model's broad sectors, since debates are still very much ongoing as to which sectors have the most potential for AI use in production.

Preferences. Consumers allocate income between goods produced by sectors A and N. Their preferences allow for the possibility that consumption in a sector eventually becomes satiated. Utility is given by

$$U = \left[\mu\left(C_A^{\frac{\varepsilon-1}{\varepsilon}} - \tfrac{\varepsilon-1}{\varepsilon} b_A^{\frac{1}{\varepsilon}} C_A\right) + (1-\mu)\left(C_N^{\frac{\varepsilon-1}{\varepsilon}} - \tfrac{\varepsilon-1}{\varepsilon} b_N^{\frac{1}{\varepsilon}} C_N\right)\right]^{\frac{\varepsilon}{\varepsilon-1}}$$

where C_A and C_N are consumption of the AI-intensive and non-AI-intensive goods. Parameter ε governs the elasticity of substitution between the two goods and is restricted to be positive. Parameters b_A and b_N determine the extent to which the goods exhibit satiation in preferences, and are also restricted to be positive. Our baseline analysis assumes $b_A = b_N = 0$, so that preferences reduce to standard CES with elasticity of substitution ε. We also consider cases where $b_j > 0$, so that marginal utility from consumption of good j declines to zero at $C_{i,j} = 1/b_j$. This allows us to consider the possibility that AI increases production of the AI-compatible sector beyond the point that consumers value. The utility function applies to all countries and worker types in the model, and realized utility levels will differ by country and type in equilibrium.

Labor endowments. There are three types of labor inputs: human capital that is complementary to AI (H_C), human capital that is substitutable with AI (H_S), and low-skilled labor (H_L). Let $\mathcal{H}$ be the set of human capital types: $\{C, S, L\}$. Labor inputs can move freely across sectors within an economy but not across economies. We assume that the developing world is endowed with a relatively larger share of L workers, and relatively smaller shares of C and S (the two categories more 'exposed' to AI; Cerutti et al. (2025); Murugan et al. (2026)). In reality, it is uncertain which skills will ultimately be complementary or substitutable with AI, so we keep labor inputs abstract and consider a range of assumptions about relative stocks of H_C and H_S.

Production. The production in country i of good j, for $j \in \{A, N\}$, is given by:

$$Q_{i,j} = Z_i \left(\omega_{C,j} X_{C,i,j}^{\frac{\sigma-1}{\sigma}} + \omega_{S,j} X_{S,i,j}^{\frac{\sigma-1}{\sigma}} + \omega_{L,j} X_{L,i,j}^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}},$$

where Z_i is country i's overall productivity level, which may capture the quality of infrastructure, institutions, and market access, for example. The terms $X_{C,i,j}$,

$X_{S,i,j}$, and $X_{L,i,j}$ represent "equipped labor services" of complementary human capital, substitutable human capital, and low-skilled labor. The $\omega_{\ell,j}$ terms are the weights on each equipped labor term and satisfy $\sum_{\ell \in \mathcal{H}} \omega_{\ell,j} = 1$ for both j, and σ is the elasticity of substitution across equipped labor inputs. When σ is larger than one it means that the equipped labor services are gross substitutes to one another, expenditure shifts towards the sector where costs are falling the fastest (as in Ngai and Pissarides, 2007), and no individual input is strictly necessary for production. In contrast, σ values less than one mean they are gross complements, so that expenditure shifts towards sectors where prices fall slowest (i.e. a Baumol effect), and a positive level of all labor services is necessary for production.[2]

The production function for equipped labor services in the A sector of type ℓ, for $\ell \in \{C, S, L\}$, combines labor inputs and AI services:

$$X_{\ell,i,A} = \left(\alpha_\ell \left(H_{\ell,i,A} \right)^{\frac{\eta_\ell - 1}{\eta_\ell}} + (1 - \alpha_\ell) \left(\text{AI}_{\ell,i,A} \right)^{\frac{\eta_\ell - 1}{\eta_\ell}} \right)^{\frac{\eta_\ell}{\eta_\ell - 1}},$$

where $H_{\ell,i,A}$ and $\text{AI}_{\ell,i,A}$ are the human capital and AI inputs used to create the equipped labor inputs of type ℓ in country i in the AI-intensive good. The parameter α_ℓ represents the relative importance of human capital inputs of type ℓ relative to AI inputs, and the elasticity of substitution η_ℓ governs the extent of substitution possibilities between human capital inputs and AI.

We assume $\eta_S > \eta_L > 1 > \eta_C$, meaning that AI is most substitutable for H_S, less substitutable for H_L (but still a gross substitute) and complementary to H_C. This captures the possibility that when AI prices fall, producers attempt to substitute most away from H_S, followed by H_L, and may substitute toward H_C. While recent progress in AI that can complete cognitive tasks affects η_S and η_C, progress in robotics could also influence η_L, which governs substitution with the large amount of labor that does much manual work in low-income countries.[3] Since human capital stocks are fixed, in equilibrium the only adjustment is through the relative wages of each of these labor groups and the intensity of AI use.

[2] We abstract from physical capital in production of the final goods for simplicity; the only capital type object in the model is stock of computing power to produce AI.

[3] Robotics may hit diffusion frictions; for example, in many countries where agricultural plot sizes are small, even existing mechanization such as tractors has yet to see full adoption (Foster and Rosenzweig, 2022; Caunedo and Keller, 2023).

In the AI-incompatible sector, production does not use AI, so

$$X_{\ell,i,N} = H_{\ell,i,N}, \tag{1}$$

where $H_{\ell,i,N}$ is the input of human capital of type ℓ in country i into the non-AI-intensive good. The non-AI-intensive good is thus produced using a simple CES production function of the H_C, H_S, and H_L inputs.

AI services. AI services are produced linearly from compute, with productivity Z_{AI}:

$$\text{AI}_i = Z_{\text{AI}} X_{i,AI}, \tag{2}$$

where the variable $X_{i,AI}$ is called compute, the capital representing economy i's stock of GPUs and data centers. Z_{AI} is the productivity of computing AI services, which is shared between both countries, and captures the capabilities of the services rendered by AI for a given number of GPUs.

For simplicity, we assume that the advanced economy has an endowment of compute, $X_{R,AI} > 0$, whereas the developing economy does not: $X_{P,AI} = 0$. In reality this is not exactly right, but the vast majority of the world's computing power is currently in advanced economies.

AI services are traded in a global market, and their price p_{AI} adjusts to equate world demand and supply. The proceeds from sales of AI services are paid to the advanced economy, since they are endowed with the compute stock.

The main force of interest we consider is how the quantity of AI services grows year over year. This can result from building more compute ($X_{R,AI}$) or increasing the productivity of compute (Z_{AI}). In our model, these two channels are isomorphic; we assume that each grows exogenously at rate γ_Z and γ_X, respectively, so that the total quantity of AI services growing exogenously at rate $\gamma = \gamma_Z \gamma_X$ each year. As the production function for equipped labor services makes clear, this growth should be considered to be growth in the amount of productive work that can be done by world AI supply measured specifically in terms of the labor that it substitutes or complements in a CES manner.

Recent work on the incorporation of AI systems into production streams strongly suggests that growth in is impressive but substantially lower than growth in benchmark abilities or compute capacity (Demirer et al., 2026; Chen and Stratton, 2026). In many of our simulations we will hold fixed the TFP terms in the rest of the

economy, Z_R and Z_P, for simplicity. This is not because we do not think any other forces will cause growth around the world, but to focus attention on the growth impacts of growth in the capabilities of AI.[4]

Differences between the R and P economies. Economies differ in their productivity Z_i, endowments of the three labor types (H_C, H_S, H_L), and endowments of compute, where $X_{R,AI} > 0$ and $X_{P,AI} = 0$. Technological parameters are common across countries.

Equilibrium. A representative firm operates the production technology in perfect competition, taking wages, goods prices, and AI prices as given. Workers receive their marginal revenue products, and wages adjust to clear labor markets in each country.

The model is frictionless and static. We focus on how access to AI changes aggregate economic outcomes and the relative earnings of different labor types. Notably, our model abstracts away from transition concerns, namely unemployment. While there are certainly substantial concerns along these lines, we focus on long-term growth analysis and, in particular, whether developing economies will be able to substantially benefit from AI proliferation.

Despite being relatively simple, the equilibrium of this economy does not, in general, admit a clean closed-form expression. Still, as we show below, it provides insight into the effects of the proliferation of AI in two special cases of the model that are analytically tractable.

Equilibrium objects of interest. For our results that follow, it will be useful to define a few quantities that come up repeatedly. Let A be the numeraire good so that $p_A = 1$. Then the *factor cost shares* are

$$\lambda_{\ell,i} = \frac{w_{\ell,i} H_{\ell,i}}{Y_i}, \qquad \lambda_{\text{AI},i} = \frac{p_{\text{AI},i} \text{AI}_i}{Y_i}$$

[4]For simplicity, we assume that growth in the stock of compute (e.g. data centers) is captured in this exogenous growth rate γ_X. In reality, compute construction is capital expenditure which responds endogenously to both the return to compute as well as the interest rate. We abstract from this endogeneity and, in our simulations that follow, simulate for multiple growth rates. Situations where AI leads to a high return to capital expenditure and thus an investment boom in compute construction are captured in our high growth-rate simulations.

where Y_i is GDP in economy i. These shares reflect the total expenditure on labor type ℓ relative to GDP, for $\ell \in \{C, S, L\}$, and the total expenditure on AI inputs in economy i relative to its GDP.

The *AI cost share in equipped labor* of type ℓ in country i for good A, $\theta_{\ell,i}$, is:

$$\theta_{\ell,i} \equiv \frac{p_{\text{AI},i}\text{AI}_{\ell,i}}{P_{\ell,i}X_{\ell,i}} = \frac{1}{1 + \left(\frac{\alpha_\ell}{1-\alpha_\ell}\right)^{\eta_\ell}\left(\frac{w_{\ell,i}}{p_{\text{AI},i}}\right)^{1-\eta_\ell}}$$

where P_ℓ is the familiar CES-ideal price of the bundled good X_ℓ:

$$P_{\ell,i} = \left[\alpha_\ell^{\eta_\ell} w_{\ell,i}^{1-\eta_\ell} + (1-\alpha_\ell)^{\eta_\ell} p_{\text{AI},i}^{1-\eta_\ell}\right]^{\frac{1}{1-\eta_\ell}}$$

and the expression for $\theta_{\ell,i}$ follows directly from the cost minimization problem. Thus $\theta_{\ell,i}$ describes the total cost share of AI in producing the equipped labor bundle $X_{\ell,i}$ as a function of the wage of type-ℓ labor in country i and the price of AI inputs there.

The *net demand elasticity for labor* ℓ, $D_{\ell,i}$:

$$D_{\ell,i} = (1 - \theta_{\ell,i})\sigma + \theta_{\ell,i}\eta_\ell$$

$D_{\ell,i}$ describes how responsive demand for a labor type H_C, H_S, or H_L is with respect to its own price. When $w_{\ell,i}$ increases, two adjustments happen in equilibrium. First, the firms substitute *across* types of equipped labor and away from the now more expensive one. Second, the firms substitute *within* the $X_{\ell,i}$ bundle away from ℓ and towards AI. The relative importance of these two channels is dictated by the AI cost share in that type of equipped labor. When the equipped labor input type is dominated by AI, within-type substitution matters more than across-type substitution. This is similar to the logic underlying skill-biased technical change in Krusell, Ohanian, Ríos-Rull and Violante (2000).

Characterizing scenarios. We pause briefly to highlight the range of scenarios that the framework can accommodate. A macroeconomic literature on AI in production has emphasized the central role that "weak links" that are not amenable to automation play in production (Jones and Tonetti, 2026). For us, this is captured in the case of $\sigma < 1$. In this case, the complementary skill/labor type C captures the humans' ability to perform tasks that are difficult to automate and fundamentally

necessary for production and complementary to AI-performed tasks.

On the opposite end of the spectrum, many have also discussed the potential of AGI capable of performing a wide range of tasks. In our model, the combination of parameters $\mu = 1$ and $\sigma > 1$ corresponds to an automation-dominant scenario where all goods that humans value can be produced with AI, and its production does not require specialized human labor. In the limit, this case leads to a labor share of GDP that converges toward zero, as in the models of Korinek et al. (2026) and Restrepo (2026).

Fewer previous studies have considered how satiation in preferences can limit the growth impacts of AI. This is distinct from the weak-links possibility described above, and instead comes in through preferences. In our model, if $b_A > 0$, consumers may reach an upper bound on the amount of AI-intensive goods they demand. As we demonstrate in the sections that follow, demand satiation can interact with aspects of production, including the extent to which AI can substitute for labor of different types.

3. Will AI Leave Developing Countries Behind?

We begin with the case in which consumers do not become satiated in either good ($b_A = b_N = 0$). We show that the model predicts two reasons why increasing AI capabilities may lead to less growth for developing countries than for richer countries. First, developing countries may lack the skills needed to complement AI. Second, more productive countries may absorb AI supply and prevent its price from falling enough to be widely used in developing countries.

3.1. Complementary skills limit developing-country gains

Consider the case where $H_S = 0$ and $\omega_{S,A} = 0$ so that only H_C, H_L, and AI are used in production. We also consider the case where $1 > \sigma > \eta_C$ so that complementary labor is actually more complementary to AI than low-skilled labor. This case may correspond to the short- and medium-run where the skills to leverage AI developments will be crucial in production.

We are interested in what happens when the quantity and productivity of compute grow (recall that they are isomorphic in the model). The two countries are linked through the equilibrium price of AI, which must (weakly) fall as the compute grows. They are also linked via balanced trade in which the developing economy ships N and A goods to the advanced economy in exchange for AI services.

We therefore study how an individual country responds to a reduction in the price of AI which the following remarks clarify.

Aggregate Labor-Income Gains Depend on AI Expenditure: A small decline in the price of AI saves money on existing AI purchases. In our competitive model, these savings translate into higher labor income. Formally, the total increase in labor income is equal to the income share of AI, as in the logic of Hulten (1978):

$$\lambda_C \frac{d \ln w_C}{-d \ln p_{\text{AI}}} + \lambda_L \frac{d \ln w_L}{-d \ln p_{\text{AI}}} = \lambda_{\text{AI}}.$$

How and why does development matter? The development-relevant primitives in the model are the stocks of high- and low-skill workers, H_C and H_L. Production is constant returns to scale so the relative skill endowment $\frac{H_C}{H_L}$ is the relevant quantity. We can thus study how the impact of improved AI productivity *changes* between the advanced and the developing economy by examining how outcomes change with respect to $\frac{H_C}{H_L}$.

Proposition 1. *Suppose $\eta_L > 1 > \sigma > \eta_C$, $\mu = 1$, and $b_A = 0$. Then we have*

$$\frac{d\lambda_{\text{AI}}}{d \ln(H_C/H_L)} = \Xi\Big(\eta_L\theta_L - \eta_C\theta_C - \sigma(\theta_L - \theta_C)\Big) > 0$$

where λ_{AI} is the economy's cost share in AI and $\Xi = \frac{\lambda_C\lambda_L}{\lambda_L D_C + \lambda_C D_L}$.

In words, the proposition says that the economy's expenditure share on AI inputs is increasing in the stock of complementary human capital relative to low-skilled labor. The reason this matters is that the economy's gains from rising AI productivity will be directly proportional to λ_{AI}, as posited above. Since advanced economies have higher stocks of complementary human capital, they thus gain more than the developing economy from AI productivity growth.

It is worth noting the "sufficient statistic" flavor of this proposition. The cost share of AI, λ_{AI}, is an endogenous but measurable statistic indicating a country's marginal gains to improvements in AI technology or increases in compute. In practice, the observation that AI services are used much less extensively in developing economies (Liu et al., 2025; Chatterji et al., 2025; Appel et al., 2025; Daepp and Slaughter, 2026) implies through the lens of our model that they will grow less as a result of increasing AI capabilities.

3.2. More productive countries absorb AI supply

We now consider the case of greater substitutability between AI and labor, in which $H_C = 0$, $\omega_{C,A} = 0$, so that the only labor inputs are human capital that is substitutable with AI and low-skilled labor. We also assume that $\eta_S > \eta_L > \sigma$. This may correspond more closely to longer run, speculative scenarios around Artificial General Intelligence (AGI) where there are few labor tasks that AI cannot accomplish. We consider the limit case, exploring what happens as the productivity of compute in producing AI services grows infinitely.

This leads us to our main proposition for the high substitutability case, which posits divergence forever between the advanced and developing economies.

Proposition 2. *Suppose $\eta_S, \eta_L > 1$, $\mu = 1$, and $b_A = 0$. If AI productivity Z_{AI} grows at a constant rate γ then the following limits hold:*

$$\lim_{t\to\infty} \gamma_{S,R} = \frac{\gamma}{\eta_S}, \qquad \lim_{t\to\infty} \gamma_{L,R} = \frac{\gamma}{\eta_L}$$
$$\lim_{t\to\infty} \gamma_{S,P} = 0, \qquad \lim_{t\to\infty} \gamma_{L,P} = 0$$

where $\gamma_{\ell,i}$ is the growth rate of the wage for labor type ℓ in country i. Further, we have

$$\lim_{t\to\infty} w_{S,P} = \bar{w}_{S,P}, \qquad \lim_{t\to\infty} w_{L,P} = \bar{w}_{L,P}$$

where $\bar{w}_{S,P}, \bar{w}_{L,P}$ are positive, finite constants.

In words, this result says that, in the long run, wage growth for both skilled and unskilled laborers in the advanced country approaches a balanced growth path with a strictly positive growth rate; as long as AI productivity growth continues, workers in these countries get to enjoy ever-growing wealth. In contrast, wages in the developing country adjust in the short-run (not necessarily positively) and eventually stop growing, reaching some constant finite level. The story is one of divergence, where workers in the developing country are left behind and, over time, income gaps continue to grow without bound.

What is the intuition behind this result? It is an outcome of constant returns to scale production and the initial difference in productivity levels. When AI is a gross substitute for other inputs in the economy, there is a threshold price at and below which a country's demand for AI services becomes infinite as (under substitutability) production can be scaled endlessly, and their use of AI is limited only

by the quantity of compute available. Unsurprisingly, this threshold is increasing in productivity. Thus the advanced economy's threshold is higher than the developing country's and, as prices are falling, is hit first. Once at the threshold, the advanced country *purchases all additional available AI services insatiably*. This endless pool of demand prevents the global price of AI from falling any further and, while the quantity of AI services used by the advanced country continues to grow (and thus wages continue to grow), the developing economy stagnates.

3.3. Simulations, no satiation

Though it is inherently speculative, we perform here some simulations to examine whether the outcomes laid out above are plausible or if they require extreme parameter values. Our goal is not to precisely predict the future but rather to get a sense of whether certain scenarios can likely be ruled out or in and provide some insight into which parameters are particularly important in determining future trajectories. We consider a range of parameter values:

The value of σ: We conduct our simulations with two cases for the elasticity of substitution across skill-types: substitutes with $\sigma = 2$, and complements with $\sigma = 0.75$. Whether this parameter is above or below unity is the key difference between our cases above and determines whether AI-complementary skills are ultimately *required* for production. Both cases are plausible, based on evidence from substitution across education groups, and from the task literature. To the extent that these skill categories represent education levels, the traditional skills literature finds evidence that these groups operate as substitutes (Katz and Murphy, 1992; Bils et al., 2024). But our categories also differentiate between skills used to perform AI-complementing or -substituting tasks. The task-based literature more consistently finds evidence of complementarities (the "weak links" of Jones and Tonetti, 2026) though most studies pre-date modern AI. To the extent that our division of skills proxies for suitability for performing tasks, these estimates may better represent our (abstract) notion of substitutability between groups.

Other parameters: We consider the case of a "typical" advanced country and a "typical" developing country. The TFP difference $\frac{Z_R}{Z_P}$ is set to generate a factor seven GDP gap between the advanced and developing country and the skill endowment $\frac{H_C+H_S}{H_L}$ is set to differ by a factor of seven, which roughly corresponds to differences between the US and India. The remaining technological parameters and the extent to which the total skill endowment is split between AI-complementary

and -substitutable skills (which we denote $\chi = \frac{H_C}{H_S}$) possess genuine uncertainty. Our approach is to assign (fairly generous) prior distributions to these parameters and then simulate model outcomes under a variety of draws from the parameter distribution in order to gauge how this uncertainty translates to uncertainty in model outcomes. Appendix A provides details as well as the precise distributions used.

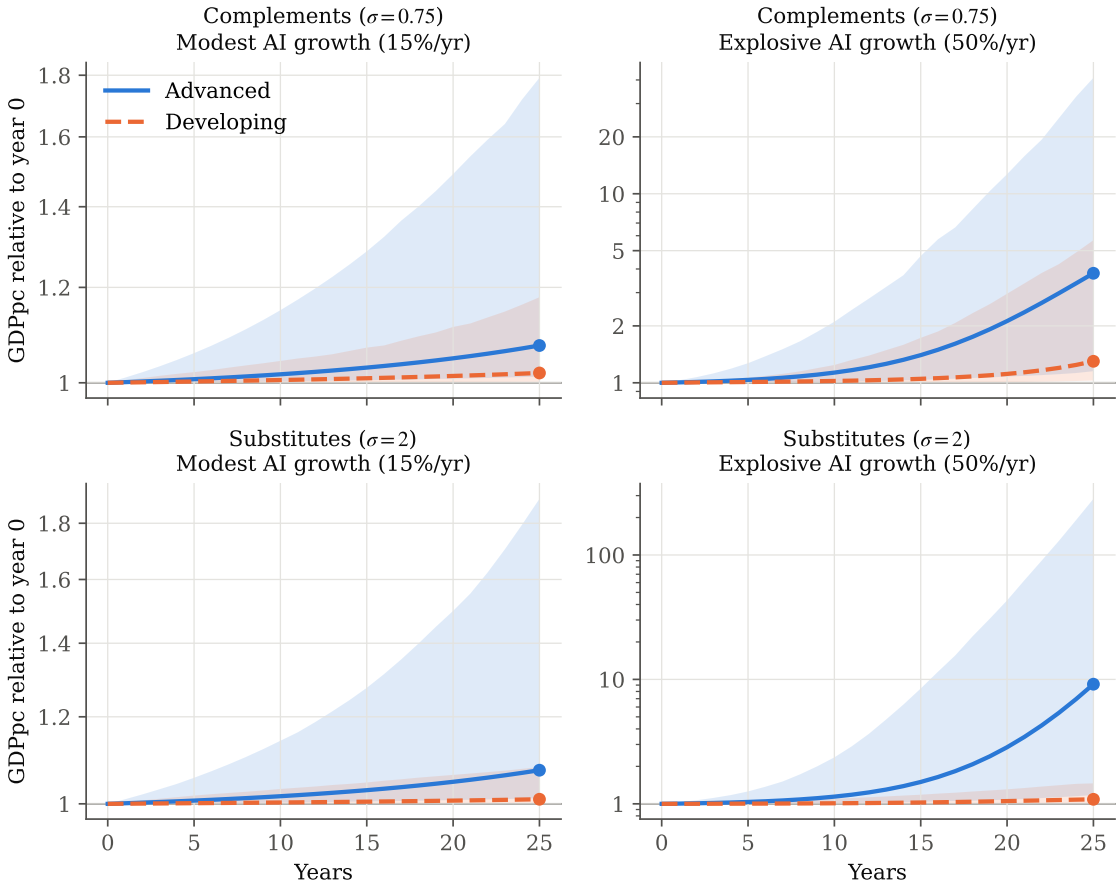
Results: With the parameterization laid out above, we simulate outcomes for 25 years under two scenarios. The first we call the "modest growth" scenario where productivity in the AI sector grows by 15 percent per year.[5] The second is the "explosive growth" scenario where productivity grows by 50 percent per year. In each case, we define the "representative" draw as the single parameter draw whose path remains closest to the median outcome in each period.[6] We plot the time path of this representative draw, while the shaded areas represent the bounds that contain 95 percent of scenarios.

Figure 1 plots the time path of GDP for the advanced and developing countries, relative to their initial GDP level, under all four scenarios we consider. From the figure, it is abundantly clear that the concerns raised by Propositions 1 and 2 above are not edge cases that only occur under particular parameterizations. The median outcome is a small increase in developing country income, accompanied by a much larger increase in advanced country income, in all four cases considered. Under modest growth, the typical path increases the gap between advanced and developing regardless of the value of σ. Under explosive growth, however, σ has a substantial impact. Under the parameterization $\sigma = 2$ and explosive growth, 100 percent of parameter draws result in the "Divergence, Forever" result of Proposition 2 where the gap between advanced and developing economies grows without bound. Even in the case of complements, the gap between advanced and developing economies grows to 2.4 times its initial level in the typical case. Overall, in essentially all cases, growth in AI increases developing economy income, but increases advanced economy income more.

[5]We assume that this comes from growth in productivity Z_{AI} so that $\gamma_{\text{AI}} = 0.15$ and $\gamma_{Comp} = 0.0$ but any combination of growth in the productivity or quantity of compute that generates 15 percent annual growth in the production of AI services gives identical results.

[6]We do this via a simple least-squares computation.

Figure 1: GDP Growth: Advanced vs Developing (No Satiation)



This figure plots GDP relative to its $t = 0$ value over 25 years in both the advanced and developing countries under the assumptions of modest vs explosive growth (left vs right columns) and gross complementarity or substitutability of AI (top vs bottom rows). Each path corresponds to the "representative" draw as described in the text, while the shaded areas contain 95 percent of all simulated scenarios. The parameters here feature no satiation in demand for AI-produced goods and services.

3.4. General results with satiation

Standard economic theory assumes that individuals benefit from increasing consumption of each good without bound, and the "Divergence, Forever" result of Proposition 2 above rests on the assumption that the advanced country's appetite for AI services never runs out. As we discussed above, it seems possible that while AI will be extremely effective for producing certain outputs, there may be some goods whose productivity is not greatly improved with AI (such as land-intensive goods like agriculture or housing). It is also possible that AI productivity growth could be so prolific that AI-produced goods are no longer scarce and society is in essence satiated with respect to these outputs. For example, AI may allow us to produce an infinite quantity of legal services, but I may only encounter a handful of situations in which I need to pursue or defend against a lawsuit in a given year.

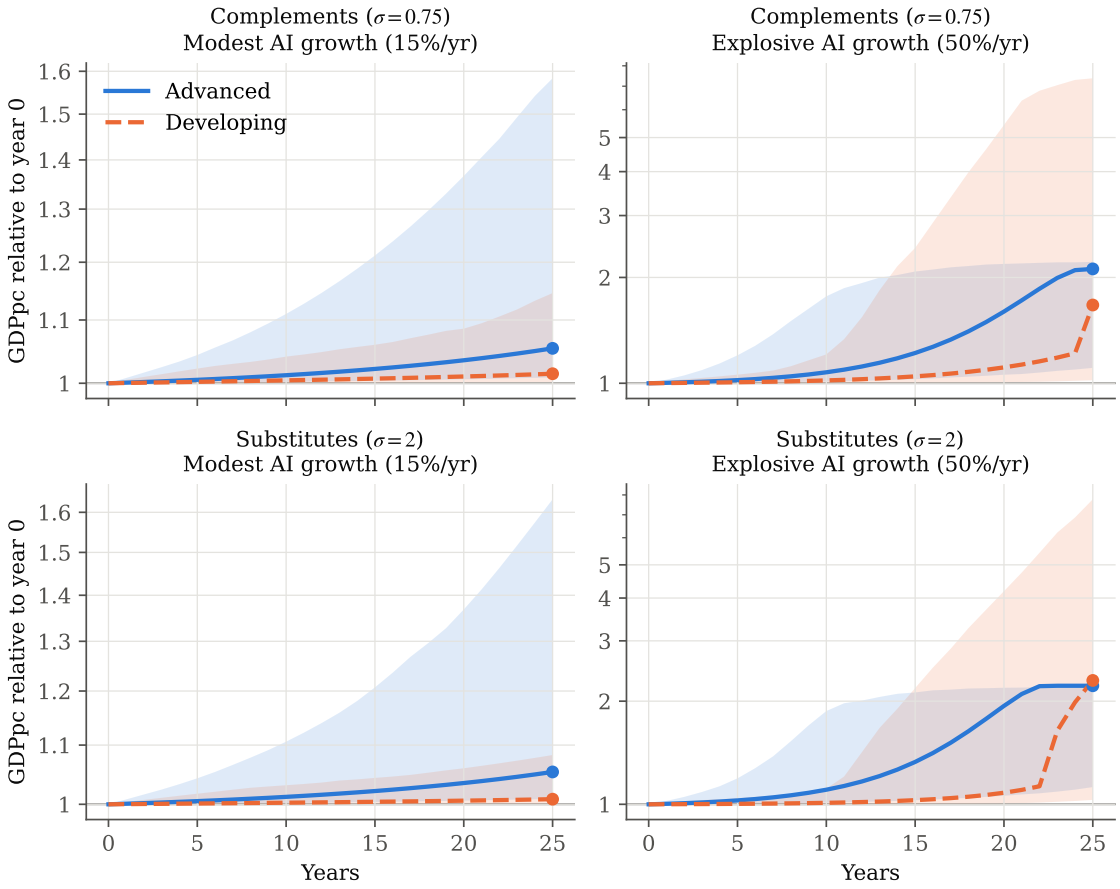
We capture this possibility by turning on two features of the model that have been dormant so far. First, we allow for $\mu < 1$ so that consumers demand some goods that AI cannot produce. Second, we choose $b_A > 0$ so that there is satiation in AI-produced goods and services. Under these assumptions, one can intuit what happens when AI generates prolific growth. The advanced country's consumption of sector A goods eventually reaches $\frac{1}{b_A}$, at which point it is satiated and it no longer demands (additional) AI services. The global price of AI then continues to fall and the developing country is able to continue its growth, breaking past the limit of Proposition 2.

3.5. Simulations with satiation

Figure 2 extends the simulation exercise to the model with satiation. We set the new parameters illustratively, without an empirical calibration. First, we set b_A so that the satiated level of demand in AI-compatible services is 1.5 times the advanced country's production at $t = 0$. Second, we set μ, the parameter governing the preference weight between the two sectors, so that the advanced country features a 30 percent employment share in sector N (at $t = 0$). Finally, we choose $\varepsilon = 0.5$ so that the two sectors are complements, though we show in Appendix Figure B1 that our results change very little under an assumption of $\varepsilon = 1.5$.[7]

[7]One important thing to note about this exercise is that the relative price of sectors A and N changes along the growth path. In the extreme case of full satiation, the price of the satiated sector falls exactly to zero. Consequently, the results for real output here depend heavily on what is chosen to be the "base year". Consistent with practice, we mitigate this issue by chain-weighting the price index.

Figure 2: GDP Growth: Advanced vs Developing (with Satiation)



This figure plots GDP relative to its $t = 0$ value over 25 years in both the advanced and developing countries under the assumptions of modest vs explosive growth (left vs right columns) and gross complementarity or substitutability of AI (top vs bottom rows). Each path corresponds to the "representative" draw as described in the text, while the shaded areas contain 95 percent of all simulated scenarios. The parameters here feature satiation in demand for AI-produced goods and services at 1.5 times the advanced country's $t = 0$ consumption level.

For modest levels of AI growth, the results under satiability look very similar to those without satiability. The median outcome continues to be small amounts of growth in the developing economy accompanied by divergence, regardless of the value of σ, with some particularly aggressive outcomes suggesting this divergence is substantial. As before, there are essentially no parameter combinations that generate any amount of convergence between advanced and developing economies.

The differences between the two cases arise in the explosive-growth exercises, particularly in the case of within-sector substitutability because, as discussed above, the mechanism driving divergence no longer holds. Eventually, demand for AI-produced output (and thus AI services) in the advanced country is satiated, and global prices can fall further, allowing the developing country to also fully automate sector A and catch up. The case of within-sector complementarity is more complicated. In this case, production cannot be fully automated even in sector A. As a result, whether the satiation point even binds depends on the precise combination of parameters drawn. If it does, then a gentler version of the logic above for the case of substitutes begins to operate, and the developing country continues to automate while the advanced country does not, helping to close the gap. On the other hand, $\sigma = 0.75$ is the case where labor H_C, which the developing country lacks, is necessary, which pushes towards divergence for the same reasons as in the baseline case of Figure 1. Which channel dominates, and thus whether convergence or divergence prevails, depends purely on parameter draws and thus, as reflected in the bounds in the figure, is highly uncertain.

These results suggest that the possibility of satiation could exert a force towards convergence. However, it is worth noting that even in this case, growth in AI generates substantial distributional issues *within* developing countries, and these may be more extreme in the case of satiation. We turn to this in Section 4 below.

3.6. Innovation and productivity catch-up could overturn divergence

An important reason that our model generates divergence is that developed countries are generally more *productive*: they are able to produce more output from the same set of inputs. This general productivity, captured in TFP Z_i, may result from using the best and latest production techniques, scientific knowledge, and institutions that improve coordination (individuals may trust each other more and can rely on courts to settle disputes). Here we consider an extension that allows for AI to affect general productivity in both the developed and developing country.

This reflects two channels.

First, AI may speed scientific discovery and discovering which processes are effective, which can improve the productivity of developed countries that lie on the frontier (Cockburn et al., 2018; Hao et al., 2026; Hill and Stein, 2026). Second, AI itself may influence the rate at which technology diffuses to developing economies, which are within the technological frontier and typically operate as adopters rather than pioneers (Keller, 2004). Although recent technologies have diffused faster than previous generations of technologies (Comin and Hobijn, 2004, 2010), they still take time. Kalyani et al. (2025) estimates that it will take 50 years for jobs related to the inventions created between 1976 and 2014 to diffuse fully across the US, let alone across the developing world. It is possible that AI may speed diffusion and shrink productivity gaps either by directly aiding in adoption (e.g. an LLM can provide guidance to a technician about how to repair a complex machine) or by leading to technological breakthroughs that are particularly shareable (e.g. the discovery of a new medication that, once the formula is known, can be manufactured easily).

In a similar vein, expanding AI capabilities may allow institutional "technology" to diffuse across countries more easily — productivity may be held back by individuals, governments, and firms not employing best practices (Bloom et al., 2013; Verhoogen, 2023). While earlier generations of information technology increased access to information about best practice, actions delegated to AI systems could simply implement it. Alternatively, developing economies face substantial frictions, including contracting and delegation (Akcigit et al., 2021); machines that think may be able to coordinate even in contexts where human trust and the quality of institutions for human coordination are low (Hadfield and Koh, 2025; Shahidi et al., 2025).

Within the model we introduce two new forces to consider potential frontier productivity gains and diffusion spillovers. First, we capture the possibility that AI proliferation improves frontier TFP in the developed country by allowing Z_R to grow with Z_{AI} according to $\frac{Z_{R,t+1}}{Z_{R,t}} = \left(\frac{Z_{\text{AI},t+1}}{Z_{\text{AI},t}}\right)^{\kappa_g}$, where κ_g represents the elasticity of TFP growth to AI productivity growth. Our previous results hold $\kappa_g = 0$ so that frontier TFP is constant over time; we now consider cases where it increases ($\kappa_g > 0$). Second, we allow general productivity to diffuse to the developing country, according to $\frac{Z_{P,t+1}}{Z_{P,t}} = G_{t+1}^{\ln\left(\left(\frac{Z_{\text{AI},t+1}}{Z_{\text{AI},t}}\right)^{\kappa_d}\right)}$, where $G_t = \frac{Z_{R,t}}{Z_{P,t}}$ represents the technological

gap between the advanced and developing country. κ_d governs how strongly AI progress accelerates the diffusion of technology. If $\kappa_d = 0$, developing country TFP remains constant, so that the gap widens if frontier productivity grows. If $\kappa_d > 0$, faster AI productivity growth speeds the diffusion of technology from the developed to developing country. A larger κ_d implies faster diffusion for a given rate of AI progress.

Whether productivity gaps narrow depends on the balance of frontier advancements and diffusion. Under these assumptions, it turns out that the long-run productivity gap between countries $\frac{Z_{R,t}}{Z_{P,t}}$ converges to e^{κ_g/κ_d}. In other words, it depends *only* on the values of κ_d and κ_g.

3.7. Simulations with productivity catch-up, no satiation

Appendix Figure B2 displays the results of our simulation exercise in this extended model, under the baseline and three counterfactuals with different parameters κ_g and κ_d. We include the case of $(\kappa_g, \kappa_d) = (0, 0)$ (i.e. the baseline case) for comparison. We then consider the case where κ_g is set so that the advanced economy experiences one percent per year TFP growth in the case of explosive AI productivity growth, the case where κ_d is set so that the developing economy experiences diffusion that closes one percent of the gap to the advanced economy per year, and the case where both κ_g and κ_d are increased.

The figure conveys that in our model, growth in the developing economy depends mostly on κ_d: diffusion can have a large impact on developing economy growth. In contrast, κ_g alone has little impact and in the most basic case where κ_g is positive and κ_d is zero, growth in the developing country is actually marginally *lower* than in the baseline case: the larger productivity gap further crowds the developing country out of using AI (increasing the impact of Proposition 2). The baseline productivity gap between the advanced and developing economy is so large that the later gap does not depend much on frontier growth, at least for non-extreme assumptions for κ_g. In the case where $\kappa_d > 0$ (so that the gains in frontier productivity from κ_g will diffuse from the advanced to the developing economy), increasing κ_g alone does not much change the developing economy. Other forces omitted from our model could make κ_g more important for developing countries; in particular if developing countries benefit from trading final goods with advanced economies following comparative advantage. We return to this later.

4. How Will AI Affect Workers Within Developing Countries?

In addition to affecting overall growth, AI may have differential effects on different groups of workers within developing countries. As before, we first describe results assuming that consumers do not become satiated.

4.1. Complementarity

We begin by restricting attention to the case of complementary skills, where $H_S = 0$ and $\omega_{S,A} = 0$ (as in Section 3.1). Under this restriction, we have intuitive formulas for individual wage changes.

Gains are split between C and L workers based on elasticities: When the price of AI declines, the total increase in country income λ_{AI} is split between high- and low-skill labor according to 1) elasticities of substitution and 2) AI cost shares:

$$\frac{d\ln w_C}{-d\ln p_{\text{AI}}} = \underbrace{\frac{1}{(\lambda_C+\lambda_L)}}_{\text{Normalization}} \left(\lambda_{\text{AI}} + \frac{\Xi}{\lambda_C} \left(\underbrace{\eta_L\theta_L - \eta_C\theta_C}_{\text{Within Nests}} - \underbrace{\sigma(\theta_L-\theta_C)}_{\text{Across Nests}}\right)\right)$$

$$\frac{d\ln w_L}{-d\ln p_{\text{AI}}} = \frac{1}{(\lambda_C+\lambda_L)} \left(\lambda_{\text{AI}} - \frac{\Xi}{\lambda_L} \left(\eta_L\theta_L - \eta_C\theta_C - \sigma(\theta_L-\theta_C)\right)\right)$$

Workers of each skill-type gain according to the (properly normalized) aggregate gains (λ_{AI}), subject to a reallocation term whose sign follows $\eta_L\theta_L - \eta_C\theta_C - \sigma(\theta_L - \theta_C)$, as Ξ is a strictly positive general equilibrium feedback term.

The reallocation term (which, as signed, is larger when the high-skill workers benefit more) balances two channels. The first represents the differential rates of within-nest substitution. If AI is a better substitute for H_L, then a decline in price will decrease demand for H_L relative to H_C, leading to a smaller wage gain (subject to differences in cost shares not reversing this logic). The second term represents substitution across nests. If AI is used more intensively in H_L (as it would be if it is a better substitute), then a decline in the price of AI leads to a larger reduction in the price of X_L, causing aggregate substitution away from X_C and increasing relative demand for H_L and thus wages. Some simple rearranging shows that this term is always positive — those who can leverage AI gain relatively more than those who cannot. This is essentially the within-country equivalent of Proposition 1.

4.2. Substitutability

We next consider the substitutability case, where $H_C = \omega_{C,A} = 0$ (as in Section 3.2) so that AI can more easily substitute for both labor types. The supplies of H_S and H_L are fixed, while wages and AI use adjust in equilibrium. The wage responses follow

$$\frac{d\ln w_S}{-d\ln p_{\text{AI}}} = \underbrace{\frac{1}{(\lambda_S+\lambda_L)}}_{\text{Normalization}} \qquad \left(\lambda_{\text{AI}} + \frac{\Xi}{\lambda_S}\big(\underbrace{\eta_L\theta_L - \eta_S\theta_S}_{\text{Within Nests}} - \underbrace{\sigma(\theta_L-\theta_S)}_{\text{Across Nests}}\big)\right)$$

$$\frac{d\ln w_L}{-d\ln p_{\text{AI}}} = \frac{1}{(\lambda_S+\lambda_L)} \qquad \left(\lambda_{\text{AI}} - \frac{\Xi}{\lambda_L}\big(\eta_L\theta_L - \eta_S\theta_S - \sigma(\theta_L-\theta_S)\big)\right),$$

where now $\Xi = \frac{\lambda_S\lambda_L}{\lambda_L D_S + \lambda_S D_L}$.

The same two forces determine how gains are distributed between workers. Within each bundle, better AI encourages substitution away from labor, which reduces the labor required per unit of equipped services. Demand shifts towards the bundle whose price falls more, which offsets some of this displacement. The reallocation term, $(\eta_L - \sigma)\,\theta_L - (\eta_S - \sigma)\,\theta_S = D_L - D_S$, captures the balance of these forces. Unlike in the complementarity case, its sign is ambiguous: S workers may lose or gain relative to L workers, and their wages can fall even as aggregate income rises.

These equilibrium wage responses may differ across economies with different skill endowments. Our next proposition asks whether, if cheaper AI reduces S wages, that decline is larger when S is scarcer.

Proposition 3. *Suppose $H_C = \omega_{C,A} = 0$, $\eta_S > \eta_L > \sigma > 1$, $\mu = 1$, and $b_A = 0$. Then if S wages are falling due to cheaper AI, $\frac{d\ln w_S}{-d\ln p_{\text{AI}}} < 0$, we have*

$$\frac{d}{-d\ln(H_S/H_L)}\left(\frac{d\ln w_S}{-d\ln p_{\text{AI}}}\right) < 0$$

In words: the decline in the S wage is *larger* when H_S-labor is *more* scarce. This result suggests a reason that AI-substitutable workers may be particularly vulnerable in developing economies. Although they represent a smaller share of employment, that scarcity raises their wage premium relative to other workers, all else equal. In our model, this premium amplifies their exposure to displacement by cheaper AI. This result depends on labor endowments being fixed. Education

and retraining could change those supplies and therefore the wage response, but are unlikely to change them in the short run.[8]

4.3. Simulations for the distribution, no satiation

We return to the simulation exercise of Section 3.3 (details can be found in Appendix A), this time examining the changes in wages for each type of labor *within* the developing country rather than relative growth between the advanced and developing country. Figure 3 displays the change in relative wages for all three types (L, S, and C) in all four of our baseline scenarios.

The figure communicates two key results. First, consistent with much of the literature, certain groups (in our case, S who is most substitutable with AI) are at risk of experiencing substantial wage losses. In all scenarios, a substantial portion of the parameterizations result in declining wages for S-type and even sometimes L-type wages (though the median outcomes suggest roughly stagnant, rather than falling, wages). The second is the degree of uncertainty in these impacts — the range of outcomes contained in our 95 percent bounds are an order of magnitude larger than the outcomes along the "median" path, highlighting the importance of parameters that are essentially unknown as well as the need for further research.

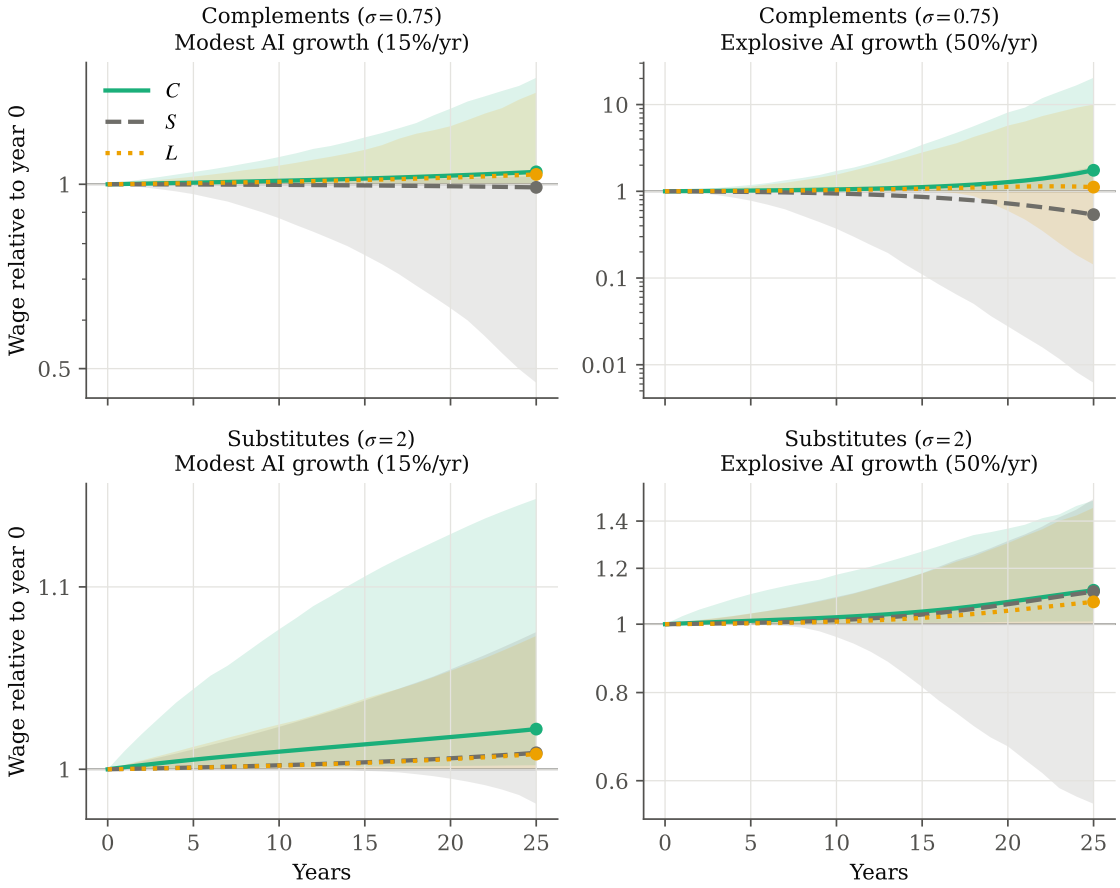
4.4. Results with satiation

As before, we extend our exercise to include the case of satiation. What makes this case interesting is the possibility that some labor types are *only* useful in the AI sector. For example, suppose that H_C cannot be used in sector N at all (i.e. $\omega_{C,N} = 0$) which, in words, would be the situation that occurs when a (say) lawyer's skills are only useful for producing legal services. In this case, proliferation in AI and the possibility that society only needs a finite quantity of legal services may eventually make the lawyer's skills completely irrelevant (on the margin), *even* if they are complementary and not substitutable with AI. This is summarized in the following result:

Proposition 4. *Suppose that* $\mu < 1$, $b_A > 0$, $\eta_S > \eta_L > 1 > \eta_C$, *and* $\omega_{C,N} = 0$. *Additionally, we need* $\frac{1}{b_A} < Z_i \, \omega_{C,A}^{\sigma/(\sigma-1)} \, \alpha_C^{\eta_C/(\eta_C-1)} H_C$, *which ensures that the satiation*

[8]Our simple model uses a common TFP term across labor bundles within each country. Cross-country differences in skill productivity could offset the scarcity premium in the S sector, and weaken this mechanism. Such offsetting differences have been documented for different education levels (Rossi, 2022), but it is not clear that this maps cleanly to a notion of AI-complementary and -substitutable skills.

Figure 3: Wage Growth in the Developing Economy (without Satiation)



This figure plots wages relative to their $t = 0$ value over 25 years for all three labor types in the developing country under our assumptions of modest vs explosive growth (left vs right column) and gross complementarity or substitutability of AI (top vs bottom rows). Each path corresponds to the "median" draw as described in the text, while the shaded areas contain 95 percent of all simulated scenarios. The parameters here feature no satiation in demand for AI-produced goods and services.

point can be technologically reached. Then we have

$$\lim_{z_{\text{AI}} \to \infty} w_C / p_N = 0$$

In essence, as AI productivity grows, the wage of H_C labor goes to zero (relative to the price of the non-AI good). Because labor is not mobile between skill categories, this results in widespread low earnings for certain skill types. It can push wages towards zero which would be akin to perpetual unemployment within our model. The prolific growth of AI allows us to make all the A goods that we need and, eventually, there is no additional demand for this sector. As a result, those who are well-suited exclusively for this sector lose.

4.5. Simulations for the distribution, with satiation

Figure 4 shows the results of our within country simulations in the case of satiation, using the same parameters as Section 3.5.[9] Under modest AI growth, the results are nearly identical to those in Figure 3, since satiation never binds. In contrast, the results change dramatically in the case of explosive growth, and the degree of uncertainty in outcomes is much higher. While this is true of all labor types, it is particularly stark for the wages of AI-complementary labor H_C. Many scenarios exhibit a tenfold or more increase in earnings while a substantial number of scenarios now indicate zero or close-to-zero wage gains. Consistent with Proposition 4, a small number of scenarios (about 2 percent) even feature a *fall* of tenfold or more in H_C earnings. This uncertainty stems from when and whether satiation in the AI-compatible sector occurs within the developing country. If it never occurs (because either the developing country's lack of complementary skills prevents it or because the advanced country absorbs all the AI services), then H_C labor remains extremely (and increasingly) valuable. However, if it does occur, then the value of H_C labor rapidly plummets to near nothing.[10] This result highlights the need for research into *who* stands to benefit and lose.

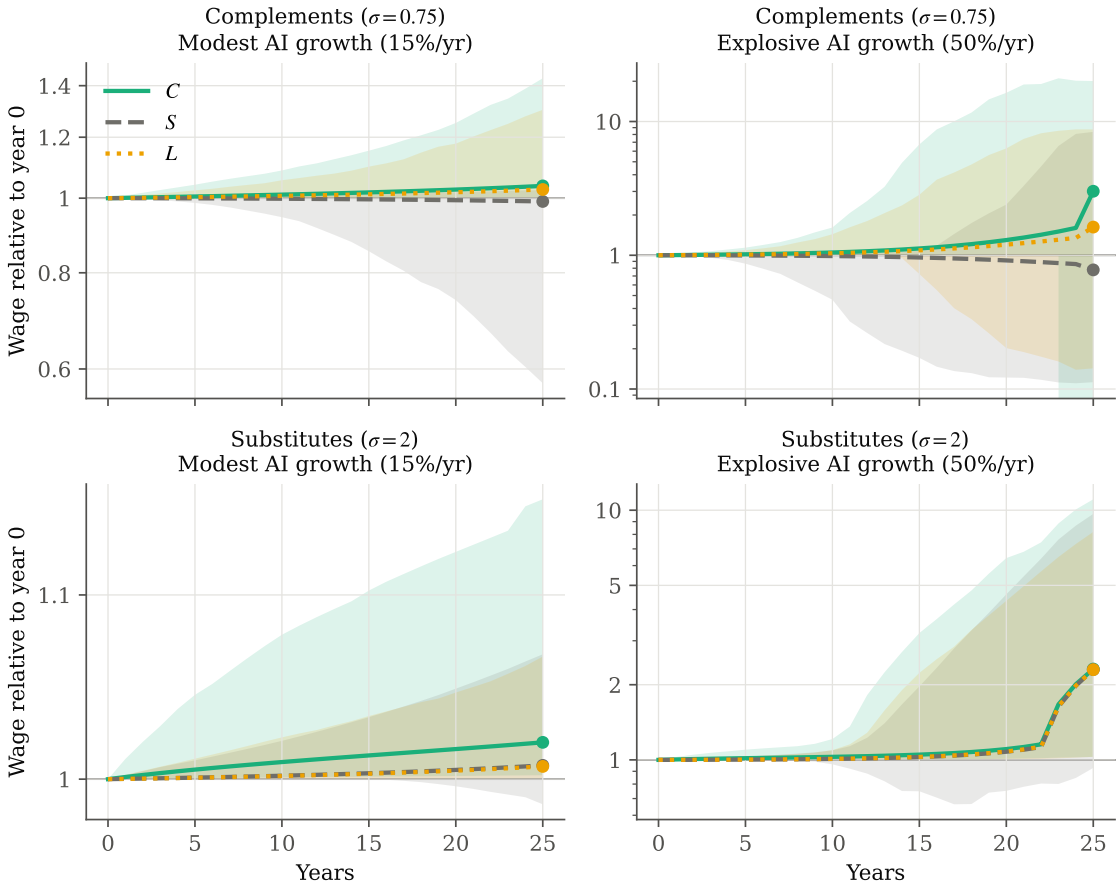
5. The Potential Impact of Policies and Greater Diffusion

In response to the potentially large stakes for developing countries, policymakers are debating various actions to take. We simulate four broad policies, modeled

[9] As before, we use $\varepsilon = 0.5$. Appendix Figure B3 shows the case of $\varepsilon = 1.5$ which is very similar.

[10] Relatedly, Bessen (2019) documents a pattern of rising and then falling employment in historical industries attributed to satiation.

Figure 4: Wage Growth in the Developing Economy (with Satiation)



This figure plots wages relative to their $t = 0$ value over 25 years for all three labor types in the developing country under our assumptions of modest vs explosive growth (left vs right column) and gross complementarity or substitutability of AI (top vs bottom rows). Each path corresponds to the "median" draw as described in the text, while the shaded areas contain 95 percent of all simulated scenarios. The parameters here feature satiation in demand for AI-produced goods and services at 1.5 times the advanced country's t = 0 consumption level.

in an intentionally stylized way. First, a policy that makes labor in the developing country more complementary to AI, filling the developing country's lack of complementary inputs. Second, a subsidy on AI usage in the developing economy, which might confront low AI adoption. Third, increasing the diffusion of innovations from the advanced to developing economy. Finally, cash transfers from the advanced to to developing economy. We consider how these policies affect production and consumption.

5.1. Making labor complementary to AI

Our first policy imagines it is possible to increase the stock of AI-complementary skills in the developing world. We model the complementarity policy as transitioning developing country workers from H_L to H_C until the stock of H_C matches that of the developed country.

We view this as a speculative exercise, for two reasons. First, the capabilities of AI are rapidly evolving, and it is hard to pinpoint exactly what skills will be most AI-complementary. One candidate is skills in how to use AI; however, these have not been stable. AI training designed in 2023 may have emphasized that LLMs are not competent at math and that mathematicians will have a stable job; advances in capabilities have disproved the first and called the second into question. It may be that the abilities more complementary to AI are not skills in using AI but tacit knowledge (such as that possessed by senior managers or political operators), or deep technical knowledge in some domains that have yet to be automated (such as PhD-level biology ability). Second, even if the appropriate skills are identified, more research will be needed to determine how to instill them.

Figure 5 displays output growth under the baseline scenario as well as this policy and the various other policies we will discuss.[11] For compactness, each column shows the output that would result under the policy at $t = 25$, with a line representing the tail end of the transition path. There is nothing surprising about the earlier portion of the transition path.

In the case where the different inputs are gross complements (i.e., "weak links") represented in the top row of our various figures, making more workers complementary to AI generates growth. Under modest AI growth, the policy results in a

[11] In addition to having an impact on how a country grows, increasing the stock of skills also has a *level* effect as it simply increases productivity. Here we focus on the impact of growth and normalize the post-policy point by GDP in year 0 *with the stock of skills increased*, so that this level effect disappears from the figure.

small increase in growth in the median scenario, but a substantial increase in the "tail" of parameters that involve substantial growth. Growth effects are larger in the case of explosive growth. This is intuitive, as increasing H_C directly addresses the issues underlying Proposition 1.

In the case of substitutes, where AI can produce the necessary outputs alone, the increase in complementary skills has a smaller effect on output. This is the case where H_C labor is not ultimately necessary for production, and the threat of divergence instead arises from AI supply being absorbed by the advanced country. While it is theoretically possible that an increase in the stock of complementary skills pushes the developing country's AI productivity past the threshold needed to overcome this channel, in our simulated economies (which, recall, feature income differences akin to those between the US and India) this does not occur even once. The output response to increasing complementary skills thus depends on whether those skills remain essential to production as AI capabilities expand.

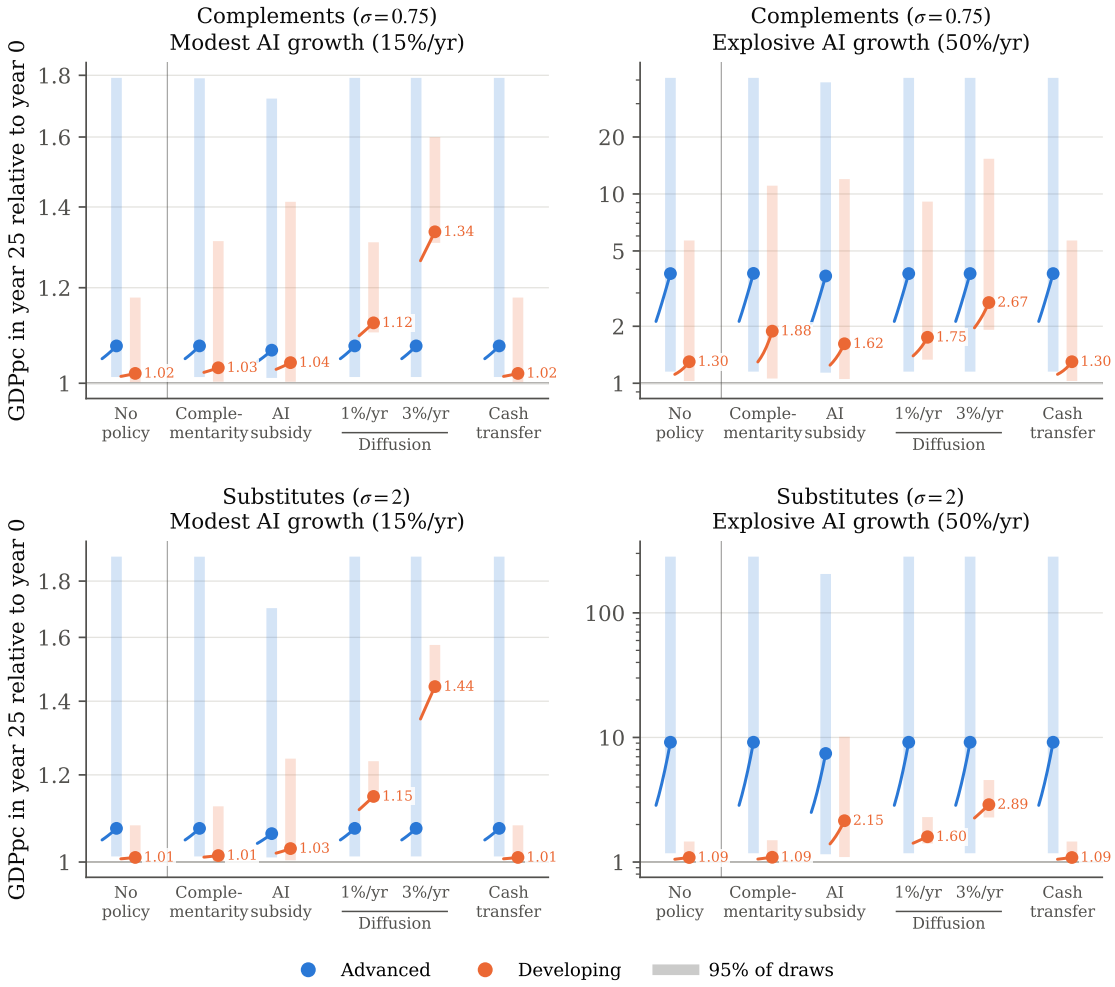
5.2. AI usage subsidy

The second policy we consider is an AI usage subsidy in the developing country, which could lower local prices enough to spur adoption. Some proposals call for subsidizing imported AI services, which could be financed by governments, global funds, or providers (Dennis et al., 2025).

In our model, we implement this by assuming the developing country is charged $(1-\tau)p_{\text{AI}}$ per unit of AI services while the advanced country continues to pay p_{AI}. We model the subsidy as being paid for via a flat tax on the returns to X_{AI}, so that the subsidy essentially takes the form of AI services providers "eating the losses" of selling to developing countries at less than the market-clearing price. We choose the subsidy rate to be equal to (one minus) the ratio of TFPs in the developed and developing country Z_P/Z_R which can be easily shown to be the smallest possible subsidy under which the Divergence, Forever result in Proposition 2 no longer holds. As we will find, this will turn out to be a massive subsidy.

The third set of points in each panel of Figure 5 shows output growth under the AI subsidy. The subsidy policy has a modest effect in the case of complements, but the impact in the case of substitutes (particularly with explosive growth) is substantial. This is unsurprising as the subsidy policy directly combats the "priced out" effect underlying Proposition 2 by allowing the developing economy to continue to afford AI services, even as the advanced economy grows indefinitely.

Figure 5: Impact of Policies and Diffusion on Year 25 GDP



This figure displays the impact of various policies and our diffusion counterfactuals on the $t = 25$ GDP (relative to the $t = 0$ level) for both the advanced and developing economies in the representative simulation. Each panel displays the results under a different set of assumptions regarding modest vs explosive growth (left vs right column) and gross complementarity or substitutability of AI (top vs bottom rows). These simulations have not accounted for the cost or feasibility of different policies so should not be interpreted as a cost/benefit analysis. The shaded bars contain 95 percent of simulations. The parameters here feature no satiation in demand for AI-produced goods and services.

However, the distortions induced by the subsidy can be large. In the case of $\sigma = 2$ and explosive AI growth (where the subsidy has a larger effect on growth), the developing country only gains 11 cents of value per dollar spent on the subsidy in the representative path. In other words, they would be indifferent between receiving the entire subsidy or receiving a cash transfer equal to 11% of the value of the subsidy. Further, in the long run, the implicit size of the subsidy is extremely large, representing 34% of all world AI use (meaning that 34% of world AI is being given for free) which, at the non-subsidized price, would be worth nearly 600% of the developing country's GDP. The reason for this substantial inefficiency, and thus cost, can be understood through the lens of allocative efficiency. The developed country has higher TFP and thus can more productively make use of AI services. By moving these services to the developing country (where TFP is lower), their value declines. Additionally, within the developing country, the subsidy induces the production side to employ AI in place of humans in cases where humans would be more productive.

Omitted forces Our model omits mechanisms such as social learning, adoption spillovers, and borrowing constraints, which could increase the growth effects of a subsidy. Our exercise is best understood as quantifying possible allocation costs of redirecting AI use.

Local compute While we model the case where foreign AI service providers are subsidized, another policy under consideration is for governments or global funds to subsidize the building of local compute infrastructure within developing countries (UN, 2025). Within our model, this would not stop the developing country from being priced out of AI, since compute is traded on a world market. If the advanced economy uses AI services more productively, its users would be willing to outbid developing country users to use any local infrastructure. There may be other unmodeled reasons for building compute infrastructure locally, such as latency; data security; or worries that access to compute may cease to be freely traded, or used as a tool of geopolitical leverage.

5.3. Productivity diffusion

Our productivity results in Section 3.6 suggest that a large determinant of the growth of the developing country is how quickly it can adopt frontier innovations. If AI speeds the advance of the frontier, there are more productivity gains to dif-

fuse. We simulate the effects of faster diffusion on growth. Diffusion is a policy objective more than a directly implementable policy, since it is not clear what actions a policymaker could take.

The two sets of points under the "Diffusion" bracket in Figure 5 show the impact of increased diffusion on GDP in the developing country. For simplicity, we fix the general productivity of the frontier economy ($\kappa_g = 0$), and consider how diffusion might shrink the productivity gap. We plot the impact for three different levels of κ_d, corresponding to cases where the productivity gap between the developed and developing country shrinks each year by 0% ("No policy"), 1%, and 3%.

Diffusion substantially improves growth on average, in all cases. When AI productivity growth is modest, even 1% diffusion allows the developing economy to grow faster than the advanced economy. Under 3% diffusion, its growth is substantially faster. This is the case where the proliferation of AI represents only a modest change to the production structure of the economy (growth is modest, so the share of production input expenditure on AI stays below 10 percent in both developed and developing economies) but the expansion in capabilities enables an entire new source of growth for developing economies via technological and institutional diffusion.

Note that the technologies diffused may represent not only those related to AI (such as models, and processes for tapping AI in decisions) but also "normal" innovations downstream of AI (such as new drug discoveries, seeds, or technology enabled by new mathematical findings). In practice, more work would be needed to understand how concrete policies affect diffusion.

5.4. Cash transfers

One policy commonly recommended to distribute gains from AI automation is cash transfers. Though our model abstracts from money entirely, we can implement a broad version of this policy via direct transfers of goods and services. Enacting this globally would require a transfer from developed to developing countries (Korinek and Stiglitz, 2021; O'Keefe et al., 2020; Yelizarova, 2025).

We simulate two possible endowment designs that have been proposed. The first, a *broad endowment*, annually transfers a fixed portion of the developed country's GDP to the developing country. This could be implemented by the advanced country committing to these ongoing transfers as it becomes wealthier, through government aid or private donations. Alternately, an equivalent implementation

in our model would be if the developing country were granted an ownership share of the advanced country's overall output (Patel et al., 2026). Others have proposed more narrowly exposing the developing economy to upside in the growth of the AI industry (Schaal, 2026). We thus also consider a second policy, an *AI industry endowment*, which annually transfers a fixed portion of the AI sector's output to the developing country.[12]

We also consider targeting the cash transfers to different groups. We first consider *universal basic income* which sends equal payments to all residents of the developing economy, and then consider policies that target only S or L workers.

No impact on production in our model. In the simplest version of our model, the production and the consumption sides of each economy can be fully separated, if one restricts attention to the case of insatiable demand where preferences are homothetic. Because of that separability, in our model cash transfers have no impact on any of the outcomes displayed in Figure 1 and thus no impact on the underlying divergence in *production*. This is apparent in Figure 5, which shows that output growth is identical to the baseline in the presence of cash transfers. This is the case for both endowments and all types of targeting.[13] Some empirical evidence finds that, contrary to our model, cash transfers have a modest effect on income (Crosta et al., 2024).

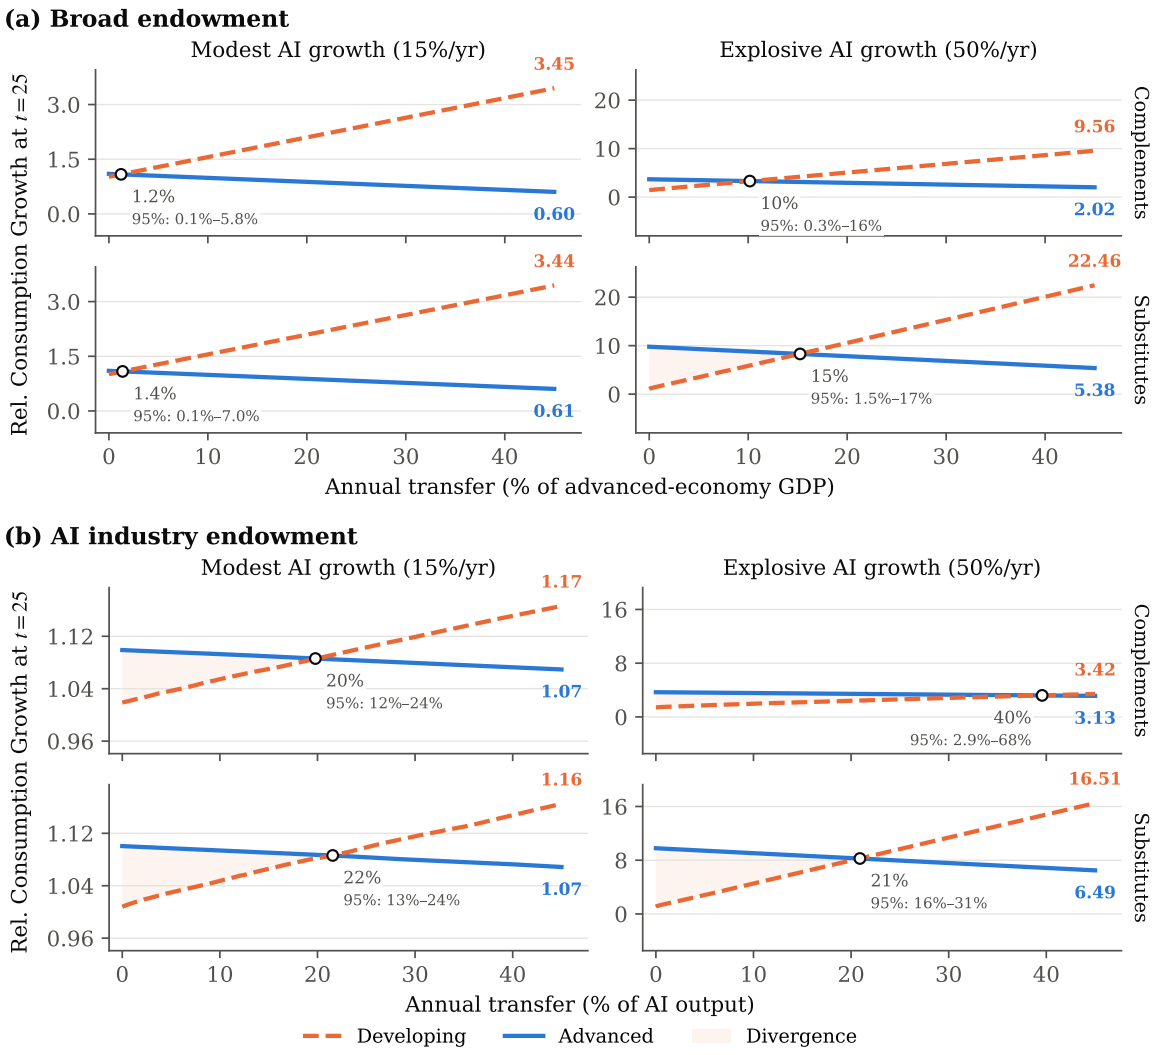
Transfers may need to be large to prevent further divergence in consumption. Even if they do not impact production in our model, cash transfers do impact consumption and thus welfare. Figure 6 summarizes the welfare impact of universal basic income funded by both types of endowment by displaying the level of consumption at $t = 25$ relative to $t = 0$ for both the advanced and developing economy (i.e. a y-value of 2 represents a case where the country's consumption doubles after twenty-five years relative to the initial period) versus the size of transfers.[14] Panel

[12] Although our model abstracts away from corporate governance, distributing ownership of the AI sector could also more widely distribute representation for decisions made by the AI industry (Bostrom, 2025).

[13] Ownership only changes who claims income; it does not change production. In particular, this means that in the case of Proposition 2 (Divergence, Forever), the developing country eventually stops using any additional AI services because it would rather sell them to the developed country at the market price than use them itself.

[14] The simulation uses the same parameters as the simulation in Figure 1 with the exception that we allow the populations of the developed and developing economies to differ in size. The difference in populations becomes a first-order consideration for transfers — halving the advanced

Figure 6: Cash Transfers May Need to be Large to Prevent Further Divergence



This figure displays the combined impact of AI growth and cash transfers on $t = 25$ consumption relative to $t = 0$ consumption for both advanced and developing economies in the representative simulation. The upper panel displays the impact of the broad endowment while the lower panel displays the AI endowment. Each sub-panel displays the results under a different set of assumptions regarding modest vs explosive growth (left vs right column) and gross complementarity or substitutability of AI (top vs bottom rows). The parameters here feature no satiation in demand for AI-produced goods and services.

(a) displays the impact of the broad endowment while panel (b) displays the impact of the AI industry endowment. We also label the transfer level at which the two lines in each panel intersect, representing the size of the transfer that would be needed to prevent any further divergence in consumption between the advanced and developing economies (in the typical simulation).

The magnitudes of the transfers necessary to even modestly offset divergence can be large, particularly in the case of explosive growth. In many of our simulations, transfers large enough to prevent consumption gaps from rising would require 1-15% of the entire advanced economy's output, or 20-40% of the AI industry's. This is much larger than historic levels of US development aid, which in 2024 amounted to 0.22% of GDP (OECD, 2025), and the entire projected AI philanthropic spend, of 0.1-0.3% of US GDP ($37-100B annually, as computed by Ransohoff (2026)). If the entirety of those sums were devoted to cash transfers they would be sufficient to prevent gaps from expanding in only a small number of parameterizations of our model (mostly under modest AI growth).

Such transfers would also be large relative to developing country production. Even transfers representing 5% of advanced country GDP would not be enough to restore convergence in most cases, but would represent 62% (complements) or 208% (substitutes) of the developing country's production annually in the median case. The transfers generated by the AI endowment are of a similar, though slightly more modest, magnitude of 14% and 148% for the case of complements and substitutes respectively.[15] Sovereign wealth funds offer precedents for such endowments, but disbursements from existing funds are much smaller. In 2025, Alaska rebated an amount equivalent to 0.8% of GDP to residents, plus an amount equivalent to 4% to government services (Alaska, 2026a,b; FRED, 2026). Other governments use sovereign wealth funds to finance government expenditure, including Norway (9% of GDP in 2025), Kiribati (13-15%), and Timor-Leste (70%) (Finance, 2026; IMF, 2026; Banco Central de Timor-Leste, 2025). Other evidence finds that cash transfers have proven to be effective even at surprisingly large scales: Egger et al. (2022) evaluates transfers corresponding to 15% of local GDP. How-

economy also halves the size of the transfer that a developing country resident receives. We assume the developing economy population is 1.43 times the population of the advanced economy, reflecting the relative proportion of the world's population living in countries with PPP GDP per capita above and below 20,000 dollars. We choose this cutoff specifically so that China is counted as an advanced economy, reflecting its participation in the AI industry.

[15] In the case of modest growth, cash transfers represent 26/27 percent of developing country production while the endowment generates transfers of only 2/2 percent (complements/substitutes).

ever, in explosive AI growth scenarios, transfers at the scales needed to reduce divergence could substantially reorganize the power structures both between and within countries.

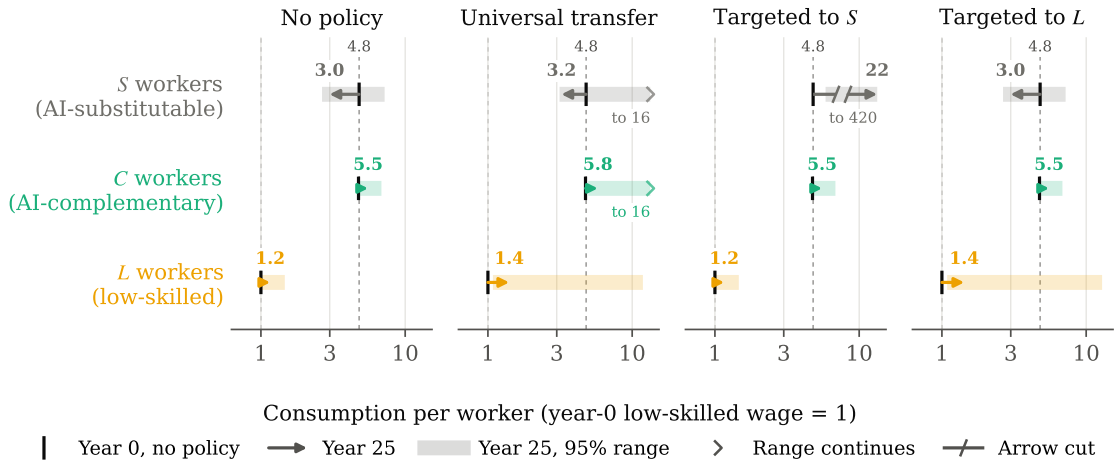
Broad vs AI endowment. The size of broad-based transfers necessary to prevent further widening of the consumption gap varies substantially with the speed of AI growth, as shown by the different locations of the intersections of the blue and orange lines in Figure 6(a), ranging from 1.2-15%. In contrast, roughly the same AI industry endowment (20-40%) prevents further widening of the consumption gap between the countries, across scenarios. AI-linked transfers end up paying out more in scenarios where the AI sector earns more, which in our model occurs when AI growth is more extreme and generates the most divergence. In practice, the payouts would depend on which firms are represented in the endowment and the degree to which those firms capture this growth. Transfers financed by a broader endowment depend on returns elsewhere in the economy as well.

Targeting. We have so far modeled universal transfers that give an equal amount to each resident of the developing country, but policymakers commonly focus transfers on specific groups (Hanna and Olken, 2018). Here there may exist trade-offs between targeting AI harms and pre-existing poverty.

In advanced economies, policymakers are considering transferring cash to compensate workers who face downside risks from AI. In our model, the workers most exposed to risks of earnings losses due to AI are skilled (those with AI-substitutable skills in all cases (S), and even those with AI-complementary skills in the case of satiation (C)–as highlighted in Figures 3 and 4). Affected jobs may include for example call center workers, lawyers, and software developers, which are relatively wealthy in developing countries. They would represent very different profiles than the typical targets for cash transfers in developing countries, which tend to be low-skilled individuals (L), which may include subsistence farmers or day laborers, who are substantially poorer.

Figure 7 demonstrates this tension by displaying the change in consumption for each labor group in the case of no cash transfers (baseline), as well as transfers of 0.5% of advanced economy GDP targeted towards different groups in the developing economy: universally, towards S workers, and towards L workers. For clarity, we limit ourselves to a single parameterization (explosive growth, $\sigma = 2$).

Figure 7: Compensating Those Harmed by AI Could Mean Helping Better-off Workers



This figure displays the impact of cash transfers on the equilibrium consumption of each type of developing economy workers, net of transfers. The values displayed come from the case of explosive growth (γ grows at 50% per year) with gross substitutes ($\sigma = 2$) and a 0.5% broad endowment transfer. The total amount of transfers is the same in each panel, only the recipients differ.

The first column presents consumption for each labor type in the initial year (base of arrow) and year 25 (tip of arrow), without any policy. S and C initially had the same, higher consumption level. In this parameter draw, S sector worker wages fall by a substantial amount but still consume more than L workers.

Compared to the no policy simulation, universal transfers do indeed make workers better off, but they do not substantially offset the losses that the most exposed S workers suffer as a result of being displaced by AI (second column). S workers are still worse off. Targeting the same amount of transfers exclusively towards these workers (S, third column) is effective at mitigating that harm, particularly because there are so few of these workers. However, as the levels show, this policy involves targeting transfers towards a rich segment of the population. Even if they were left alone, the skilled-but-displaced workers would be wealthier than the subsistence farmers. Targeting transfers exclusively towards low-skilled workers (L, final column) would increase the consumption of that group but not mitigate the AI disruption faced by S workers.

6. Discussion

Our framework suggests that advances in AI could profoundly impact the relationship between developing and developed economies, even if developing economies make limited use of the technology. In this section we consider what quantities can be measured to better anticipate these impacts. We also lay out questions omitted from our framework.

6.1. Measurement

Evidence on substitution, demand, and productivity growth can distinguish our scenarios, while responses to changes in AI access, prices, and complementary inputs can diagnose what constrains particular economies. Note that low AI usage alone (Liu et al., 2025; Chatterji et al., 2025; Appel et al., 2025; Daepp and Slaughter, 2026; Iscenko et al., 2026) is not enough to distinguish between whether complementary inputs are missing, prices are high relative to domestic productivity, or demand is satiated. Table 1 summarizes a measurement agenda, which we develop below.

Distinguishing scenarios Several measures can help distinguish between scenarios. We expect that the first two will be visible first in advanced economies, and so efforts can piggyback on existing measurement systems being built in those economies. However, additional measures are needed to understand how those capabilities will affect developing economies.

AI capabilities and supply. One can monitor AI model capabilities based on benchmark performance, and monitor how those translate into performance on real world tasks, and substitute for or complement human labor.

Satiation. It could be that AI dramatically improves the output of the AI sector, but consumers are becoming satiated in those goods. One indication of such satiation would be falling price-sensitivity in sectors that are exposed to AI as demand becomes saturated.

Productivity gaps. In addition to monitoring how aggregate TFP evolves using firm surveys, one can monitor what decisions are delegated to AI, and assess whether resulting practices improve across countries.

Table 1: Evidence to Diagnose Scenarios and Constraints

Question	**Relevant Model Objects**	**Informative Evidence**
Panel A: Distinguishing scenarios		
Can AI substitute for all human skills?	Substitution between AI and labor σ,$\boldsymbol{\eta}$	Demand for AI as AI productivity increases
How quickly will AI supply grow?	Compute X_{AI} and AI capabilities Z_{AI}	Growth in compute, and capabilities per unit of compute.
Will demand for AI-produced output keep expanding?	Demand composition and substitution, μ, ε; satiation, b_A, b_N	Expenditure responses to real price changes, quality improvements, and new products.
Will productivity gaps narrow?	Productivity innovation κ_g and diffusion κ_d; also $\sigma, \varepsilon, H_C, H_S$	Adoption and implementation of improved practices; productivity growth in both frontier and developing country.
Panel B: Diagnosing constraints		
Which complementary inputs constrain AI use?	Skill endowments (H_C, H_S) and their interaction with AI; other constraints represented by Z_i	Varying AI access with candidate complements, such as training or localization, and measuring how adoption and productivity change.
To what extent are high AI prices constraining the developing world?	World AI price relative to local wages, relative productivity, factor prices, and AI demand	Local usage and output responses to AI prices, combined with advanced-country demand, and how AI supply is allocated across countries.

Diagnosing constraints Other measures can help diagnose the factors that most constrain the benefits of AI in each location.

What complementary inputs constrain AI use? In some scenarios, developing economies benefit less from AI because they lack complementary factors. It may be possible to identify which factors are most constraining by measuring differences in adoption and productivity gains.

Our model considers complementarity between AI and human capital, which could be observed if some groups (like college graduates) have higher usage or productivity gains from adopting AI than others. Evidence is mixed as to what types of workers benefit most from previous AI systems: Brynjolfsson et al. (2025b) finds GPT-3 assistance helps less experienced call center workers in the Philippines; Otis et al. (2026) finds that a GPT-4-based entrepreneur assistant helps higher profit businesses more than low profit businesses in Kenya. Another indication is whether productivity with AI greatly improves with training: Humlum and Vestergaard (2026) find that workplaces that encourage AI use have higher adoption in Denmark. Benchmarks with locally relevant tasks (McPeak et al., 2024) can also illuminate whether limitations arise from AI systems themselves or from issues with deployment.[16]

Lack of other complementary factors could also limit productive domestic use (World Bank, 2026), including training data in local languages (Bandarkar et al., 2024; Adebara et al., 2025; Hupkes and Bogoychev, 2025); digitization of economic activity (Jensen and Weigel, 2026); and access to smartphones, computers, reliable internet access, and electricity (World Bank, 2026). The importance of these gaps could be diagnosed by assessing whether gaps in these factors coincide with gaps in use; for example, Daepp and Slaughter (2026) finds less usage in local languages relative to major languages.

In some cases it may be possible to design AI systems that are more complementary to developing economies, for example by building interfaces that are more accessible for populations with less connectivity and less access to computers, or tailored to local needs (like predicting the weather outcomes most relevant to local agricultural decisions as in Aitken et al. (2026)).

[16]For physical automation, it may be possible to differentiate between specific mechanization goods that complement or substitute for human labor based on a textual description (de Souza and Li, 2023), but current AI systems have such broad capabilities that it is difficult to do so.

AI prices. Alternately, developing countries may have productive uses of AI that are uneconomical at the equilibrium world price. In this case, we may observe that their use of AI is highly price sensitive, and increases substantially with natural price reductions or temporary subsidies. Assessing the response to temporary subsidies may be a useful diagnostic even if one does not wish to subsidize AI in the long run. When the market for AI services is global, expanding the supply of compute infrastructure anywhere in the world may help reduce prices.

Table 2: Cost Illustrations for API Access

Country	Daily Output per Person ($)	Cost of One Billion Input Tokens ($)	Days of Output Required
United States	360	10,000	28
Vietnam	20	10,000	500
Malawi	3	10,000	3,333

Notes: First column represents GDP per capita in 2025 as reported by the World Bank, divided by 250 working days. Token cost assumes $10 per million input tokens, the current API price for the frontier model GPT-6 Astra at OpenAI (output tokens are more expensive, at $50 per million).

Are developing economies getting priced out of AI inputs already? Some evidence suggests the answer could be yes. In emerging markets, many people lack smartphones that could access AI services (Björkegren, 2025), but prices of smartphones have risen as global computer memory prices have responded to increased demand for AI workloads (Shanahan, 2026). That is expected to reduce shipments of the devices to sub-Saharan Africa. For AI model services, Table 2 provides some illustrative calculations. Imagine a potential creator uses AI tokens as an input. Suppose that this creator needs one billion tokens to create their app. Consider the cost relative to GDP per capita in the United States ($360 per day), Vietnam ($20 per day), and Malawi ($3 per day). It would take 28 days of production for an average worker in the US to pay for these tokens of the frontier model GPT-6 Astra, which cost around $10 per million input tokens (output tokens are more expensive, so consider our calculation a lower bound). The same billion tokens would take around 500 days to earn in Vietnam, and around 3,333 days in Malawi. These simple calculations are meant to illustrate a point: if there is a single world price of a token, then the labor input required to make use of that token will be far higher in developing countries, since wages in dollars at exchange rates are so much lower

there. This gap could widen if wages grow faster in advanced countries than in developing nations.

Distribution Economies may also wish to monitor losses to different groups. If targeting pre-existing poverty requires identifying people with low levels of consumption (in our model, L), targeting people harmed by AI may require identifying those whose outcomes *change* (in our model, S).[17] It may be possible to monitor changes in outcomes such as earnings and employment (Ravanilla, 2026), or in precursors such as productivity on different tasks and job postings (Humlum and Vestergaard, 2026; Brynjolfsson et al., 2025a).

6.2. Other considerations

Our framework is necessarily simple. The impacts of AI on developing countries could be affected by many other forces, which can be modeled in future work.

Trade. In our model, the developing economy imports only AI services, in exchange for the numeraire good. In particular, there is no trade based on comparative advantage. We consider scenarios where the developing economy benefits from general productivity spillovers from the advanced economy, which can allow it to utilize better production technologies (e.g. designs for more energy efficient machines), produce better products (e.g. factories may learn a design for a new medicine) and implement better business practices (e.g. improve capital allocation). In reality it is also possible that the developing economy benefits from productivity improvements in the advanced economy via trade. This could be because other goods produced in the advanced economy improve in quality. For example, developing countries often import medications from advanced economies and proliferation in AI could change the quality of these medications. It could also simply occur in quantity terms where high productivity advanced economies are willing to trade substantial resources in exchange for goods in which developing countries have comparative advantage.

We also assume away any form of labor mobility. Allowing workers to move between advanced and developing economies would have substantial impacts on the implicit welfare results, as workers move between economies with different productivity levels.

[17] If demand for AI intensive output becomes satiated, complementary workers (C) could also see wage declines, as shown in Figure 4.

Adjustment, learning, and endogenous skills. We model workers as having a fixed type (C, S, or L). However, increased AI capabilities may change the incentive to acquire skills, as well as the cost (if e.g. AI revolutionizes the education sector). To the extent that these forces allow developing economies to organically acquire AI-complementary skills, they may close the gap between our baseline scenarios and our more optimistic scenarios in which the supply of complementary skills is increased. Proliferation and the associated labor market shifts could also eliminate entry level jobs through which skills were historically acquired (Brynjolfsson et al., 2025a; Garicano and Rayo, 2026). Though it is not clear if these would differentially impact skills that are substitutable or complementary with AI, it is at least possible that this would make the shortage-of-skills scenario more dramatic.

Compute. We model compute and AI services as being in fixed supply each period, growing exogenously, and fully owned by the advanced country. Thus payments to compute represent pure profits and fully accrue to the rich country. While the value of these payments has minimal impact on our simulations of convergence and divergence, exploring deviations from these assumptions could be worthwhile.[18] Most notably, compute in reality requires investment. Second, if value from portions of the compute construction supply chain is captured within developing countries (such as the production of the raw materials necessary for GPU manufacturing), this may allow them to indirectly capture some of the income allocated to compute construction which would work against divergence.

Our model treats AI as a single production technology whose quantity and quality are fungible (consistent with a large macroeconomic literature on capital more broadly). However, different AI models differ substantially in their capabilities and requirements, and distinguishing between these dimensions may be important for developing economies. For instance if the countries producing AI restrict access to the most capable systems, this may be an additional force pushing towards divergence. On the other hand, it could be that basic non-frontier models are sufficient to cover many economic use cases. The market competition for compute resources between advanced and developing countries could be lessened either if those basic models are so efficient that they produce high value per unit of compute, or if they can be run on different compute resources (such as on

[18] In our simulations, similar results hold even if we restrict attention only to convergence in labor income.

mobile phones or vintage GPUs that are not useful for frontier computation). That could help mitigate our "Divergence, Forever" scenario.

Energy. AI systems demand energy, which could raise prices (Bogmans et al., 2025). If energy prices rise globally, that could affect broader uses of energy in developing countries, such as transportation and heavy industry. However, access to cheap energy within a country may not be sufficient to drive local data center investment (Lokshin, 2026), so there may be different impacts on tradable and nontradable forms of energy.

Political systems, conflict, and regulation. Even if AI theoretically improves productivity, if it benefits some workers at the expense of others, that could lead to conflict that could undermine economic activity (Acemoglu et al., 2026a). It is possible that conflict and instability are more likely in weaker states. Weaker states might also be prone to being destabilized by emerging military technologies like drones. Alternately, technologies like facial recognition could enable mass surveillance (Beraja et al., 2023a,b) and the ability to monitor the content of electronic communication at scale could increase centralized authority.

On the other hand, advanced countries that have politically powerful workforces exposed to AI may encounter backlash that constrains the technology. If such constraints emerge less in developing countries, that could allow for regulatory arbitrage. For example, if fully robotic surgery is disallowed in advanced economies it may be allowed in countries with small medical systems, enabling new forms of medical tourism.

7. Conclusion

It is too soon to know exactly how advances in AI will affect the world's diverse economies. However, our results suggest that the stakes for the developing world could be enormous: depending on the underlying forces, AI could increase or reduce divergence, and it could cause domestic inequality of a type not seen previously. AI has spurred frenetic economic activity around the globe. Our results suggest that researchers should model and measure how AI may affect the developing world, before the dust settles.

References

Acemoglu, Daron, "The Simple Macroeconomics of AI," *Economic Policy*, 2025, *40* (121), 13–58.

_ , A. Arda Gitmez, and Mehdi Shadmehr, "Automation and Repression," June 2026.

_ and Fabrizio Zilibotti, "Productivity Differences," *The Quarterly Journal of Economics*, May 2001, *116* (2), 563–606.

_ and Pascual Restrepo, "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review*, June 2018, *108* (6), 1488–1542.

_ , Ufuk Akcigit, and Simon Johnson, "Technology and Economic Development," January 2026.

Adebara, Ife, Hawau Olamide Toyin, Nahom Tesfu Ghebremichael, Abdelrahim Elmadany, and Muhammad Abdul-Mageed, "Where Are We? Evaluating LLM Performance on African Languages," June 2025. arXiv:2502.19582 [cs.CL].

Aghion, Philippe and Simon Bunel, "AI and Growth: Where Do We Stand?," 2024.

_ , Benjamin F. Jones, and Charles I. Jones, "Artificial Intelligence and Economic Growth," in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds., *The Economics of Artificial Intelligence: An Agenda*, University of Chicago Press, May 2019, p. 0.

Aitken, Colin, Michael K. Tippett, Pedram Hassanzadeh, Katherine Kowal, Rendani Mbuvha, John H. Marsham, Shruti Nath, Ousmane Ndiaye, Douglas J. Parker, Caroline M. Wainwright, Michael Kremer, and William R. Boos, "Can we create a 'race to the top' for weather forecasts to inform smallholder farmer decisions?," October 2026. arXiv:2610.00782 [physics.ao-ph].

Akcigit, Ufuk, Harun Alp, and Michael Peters, "Lack of Selection and Limits to Delegation: Firm Dynamics in Developing Countries," *American Economic Review*, January 2021, *111* (1), 231–75.

Alaska, "The Fund," 2026.

_ , "Permanent Fund Dividend - Summary of Dividend Applications & Payments," 2026.

Alonso, Cristian, Andrew Berg, Siddharth Kothari, Chris Papageorgiou, and Sidra Rehman, "Will the AI Revolution Cause a Great Divergence?," *Journal of Monetary Economics*, 2022, *127*, 18–37.

Andrews, Isaiah and Maryam Farboodi, "Do Markets Believe in Transformative AI?," September 2025.

Antràs, Pol, Luis Garicano, and Esteban Rossi-Hansberg, "Offshoring in a Knowledge Economy," *The Quarterly Journal of Economics*, February 2006, *121* (1), 31–77.

Appel, Ruth, Peter McCrory, Alex Tamkin, Miles McCain, Tyler Neylon, and Michael Stern, "Anthropic Economic Index report: Uneven geographic and enterprise AI adoption," Technical Report, Anthropic 2025.

Appenzeller, Guido, "Welcome to LLMflation – LLM Inference Cost Is Going Down Fast," 2024.

Athey, Susan, Fiona Scott Morton, and Catherine E. Tucker, "Artificial Intelligence, Competition, and Welfare," in "The Economics of Transformative AI," University of Chicago Press, September 2026, pp. 95–118. Section: The Economics of Transformative AI.

Banco Central de Timor-Leste, "Petroleum Fund of Timor-Leste," Technical Report LXXVII December 2025.

Bandarkar, Lucas, Davis Liang, Benjamin Muller, Mikel Artetxe, Satya Narayan Shukla, Donald Husa, Naman Goyal, Abhinandan Krishnan, Luke Zettlemoyer, and Madian Khabsa, "The Belebele Benchmark: a Parallel Reading Comprehension Dataset in 122 Language Variants," in "Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)" 2024, pp. 749–775. arXiv:2308.16884 [cs.CL].

Barhoumi, Karim, Fabia A. De Carvalho, Michael Gorbanyov, Yosuke Kido, David Koll, Dragana Ostojic, Baoping Shang, Natalia Tamirisa, Sally Toms,

Era Dabla-Norris, Anh Nguyen, and Yunhui Zhao, "Global Economic and Financial Implications of Artificial Intelligence: Lessons from a Scenario Planning Exercise," *IMF Notes*, April 2026, *2026* (002), 1.

Basu, Susanto and David N. Weil, "Appropriate Technology and Growth," *The Quarterly Journal of Economics*, November 1998, *113* (4), 1025–1054.

Benzell, Seth G., Laurence J. Kotlikoff, Guillermo LaGarda, and Victor Yifan Ye, "Simulating Endogenous Global Automation," September 2021.

Beraja, Martin, Andrew Kao, David Y Yang, and Noam Yuchtman, "AI-tocracy," *The Quarterly Journal of Economics*, August 2023, *138* (3), 1349–1402.

_ , _ , David Y. Yang, and Noam Yuchtman, "Exporting the Surveillance State via Trade in AI," September 2023.

Bessen, James, "Automation and jobs: when technology boosts employment," *Economic Policy*, October 2019, *34* (100), 589–626.

Bick, Alexander, Adam Blandin, David J. Deming, and Tyler R. Schumacher, "What Work Does Generative AI Do?," Working Paper 35677, National Bureau of Economic Research August 2026.

Bils, Mark, Bariş Kaymak, and Kai-Jie Wu, "Labor Substitutability Among Schooling Groups," *American Economic Journal: Macroeconomics*, 2024, *16* (4), 1–34.

Björkegren, Daniel, "SMS Charges May Be Inhibiting the Use of AI for 1 Billion People," 2025.

_ , "The intelligence is plenty but the workers are few," April 2026.

_ , "Market Beliefs about Open vs. Closed AI," January 2026. arXiv:2512.14969 [econ.GN].

Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts, "Does Management Matter? Evidence from India," *The Quarterly Journal of Economics*, February 2013, *128* (1), 1–51.

Blumenfeld, Alex, Jonathon Hazell, Chen Lian, and Andreas Schaab, "The Macroeconomic Effect of AI: Sizing the Software Engineering Channel," September 2026.

Bogmans, Christian, Patricia Gomez-Gonzalez, Giovanni Melina, Andrea Pescatori, and Sneha Thube, "Power Hungry: How AI Will Drive Energy Demand," Technical Report WP/25/81, IMF 2025.

Bostrom, Nick, "Open Global Investment as a Governance Model for AGI," 2025.

Brynjolfsson, Erik, Avinash Collis, Felix Eggers, Sophia Kazinnik, and David Nguyen, "What is Generative AI Worth?," April 2026.

_ , Bharat Chandar, and Ruyu Chen, "Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence," *Stanford Digital Economy Lab. Published August*, 2025.

_ , Danielle Li, and Lindsey Raymond, "Generative AI at Work," *The Quarterly Journal of Economics*, May 2025, *140* (2), 889–942.

Caunedo, Julieta and Elisa Keller, "Capital Embodied Structural Change," *Unpublished manuscript*, 2023.

Cerutti, Eugenio, Antonio Garcia Pascual, Yosuke Kido, Longji Li, Giovanni Melina, Marina M Tavares, and Philippe Wingender, "The Global Impact of AI: Mind the Gap," Technical Report WP/25/76, IMF 2025.

Chatterji, Aaron, Thomas Cunningham, David J. Deming, Zoe Hitzig, Christopher Ong, Carl Yan Shan, and Kevin Wadman, "How People Use ChatGPT," September 2025.

Che, Natasha, Weining Xin, and Taichi Yoshida, "AI and Economic Divergence in Asia," *IMF Working Papers*, August 2026, *2026* (166), 1.

Chen, Fiona and James Stratton, "Artificial Intelligence in the Firm: Bottlenecks in Software Production," 2026.

Chow, Trevor, Basil Halperin, and J Zachary Mazlish, "Transformative AI, existential risk, and real interest rates," *Working Paper*, 2025.

Cockburn, Iain M., Rebecca Henderson, and Scott Stern, "The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis," in "The Economics of Artificial Intelligence: An Agenda," University of Chicago Press, January 2018, pp. 115–146.

Comin, D. and B. Hobijn, "Cross-country technology adoption: making the theories face the facts," *Journal of Monetary Economics*, January 2004, *51* (1), 39–83.

Comin, Diego and Bart Hobijn, "An Exploration of Technology Diffusion," *American Economic Review*, December 2010, *100* (5), 2031–2059.

_ **and Martı Mestieri**, "If Technology Has Arrived Everywhere, Why Has Income Diverged?," *American Economic Journal: Macroeconomics*, 2018.

Crosta, Tommaso, Dean Karlan, Finley Ong, Julius Rüschenpöhler, and Christopher R. Udry, "Unconditional Cash Transfers: A Bayesian Meta-Analysis of Randomized Evaluations in Low and Middle Income Countries," August 2024.

Daepp, Madeleine I. G. and Isaac Slaughter, "How Early Adopters Used Generative AI Worldwide: Variation by Country Income and Language," May 2026. arXiv:2605.30685 [cs.CY].

Das, Amreeta, Caroline Falkman Olsson, and Yafah Edelman, "Near-Daily AI Use Among US Adults Has More Than Doubled in Six Months," Epoch AI Data Insight 2026. https://epoch.ai/data-insights/near-daily-ai-use-doubled.

de Souza, Gustavo and Haishi Li, "Robots, Tools, and Jobs: Evidence from Brazilian Labor Markets," Technical Report, Federal Reserve Bank of Chicago 2023.

Demirer, Mert, Leon Musolff, and Liyuan Yang, "Writing Code vs. Shipping Code: Productivity Effects Across Generations of AI Coding Tools," May 2026.

Dennis, Claire, Sam Manning, Stephen Clare, Boxi Wu, Okechukwu Jake Effoduh, Chinasa T Okolo, Lennart Heim, and Katya Klinova, "Options and Motivations for International AI Benefit Sharing," 2025.

Egger, Dennis, Johannes Haushofer, Edward Miguel, Paul Niehaus, and Michael Walker, "General equilibrium effects of cash transfers: experimental evidence from Kenya," *Econometrica*, 2022, *90* (6), 2603–2643.

Eisfeldt, Andrea L., Gregor Schubert, and Miao Ben Zhang, "Generative AI and Firm Values," May 2023.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, "GPTs are GPTs: Labor Market Impact Potential of LLMs," *Science*, June 2024, *384* (6702), 1306–1308.

Foster, Andrew D and Mark Rosenzweig, "Are There Too Many Farms in the World? Labor-Market Transaction Costs, Machine Capacities and Optimal Farm Size," *Journal of Political Economy*, 2022.

FRED, "Gross Domestic Product: All Industry Total in Alaska," September 2026.

Garicano, Luis and Luis Rayo, "Training in the Age of AI: A Theory of Career Viability," *Dropbox*, 2026.

Hadfield, Gillian K. and Andrew Koh, "An Economy of AI Agents," in "The Economics of Transformative AI," University of Chicago Press, December 2025, pp. 119–138.

Hanna, Rema and Benjamin A. Olken, "Universal Basic Incomes versus Targeted Transfers: Anti-Poverty Programs in Developing Countries," *Journal of Economic Perspectives*, November 2018, *32* (4), 201–226.

Hao, Qianyue, Fengli Xu, Yong Li, and James Evans, "Artificial intelligence tools expand scientists' impact but contract science's focus," *Nature*, January 2026, *649* (8099), 1237–1243.

Hill, Ryan R. and Carolyn Stein, "How Artificial Intelligence Shapes Science: Evidence from AlphaFold," April 2026.

Hulten, Charles R., "Growth Accounting with Intermediate Inputs," *The Review of Economic Studies*, October 1978, *45* (3), 511–518.

Humlum, Anders and Emilie Vestergaard, "Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI," 2026.

Hupkes, Dieuwke and Nikolay Bogoychev, "MultiLoKo: a multilingual local knowledge benchmark for LLMs spanning 31 languages," April 2025.

Ide, Enrique and Eduard Talamàs, "Artificial Intelligence in the Knowledge Economy," *Journal of Political Economy*, December 2025, *133* (12), 3762–3800.

_ **and** _ , "The Impact of AI on Global Knowledge Work," *Journal of Monetary Economics*, January 2026, *157*, 103876.

Imam, Patrick and Jonathan Temple, "Relative Development and the Intelligence Divide: Human Capital, Technology Diffusion, and AI," *IMF Working Papers*, September 2026, *2026* (190), 1.

IMF, "Kiribati: 2026 Article IV Consultation-Press Release; Staff Report; and Statement by the Executive Director for Kiribati," *IMF Staff Country Reports*, May 2026, *2026* (099). ISBN: 9798229045421.

Iscenko, Zanna, Scott Strand, Yiyuan Chen, Guillaume Aimard, Mihai Codreanu, Vivek Sampathkumar, Alex Imas, Julian Jacobs, Evalyne Muiruri, Juan Mateos-Garcia, Jia Jen Ng, Samirah Javed, Josh Martin, Omar Ajmeri, Denis Calin, Andrew Kim, Fabien Curto Millet, and James Manyika, "Google's AI & Economy ATLAS v1.0: Mapping Gemini Usage in the Economy," Technical Report, Google 2026.

Jensen, Anders D. and Jonathan L. Weigel, "No Taxation without Administration: Bringing the State Back into the Public Finance of Developing Countries," *Journal of Economic Literature*, March 2026, *64* (1), 246–80.

Jones, Charles I. and Christopher Tonetti, "Past Automation and Future AI: How Weak Links Tame the Growth Explosion," 2026. Unpublished Manuscript, Stanford University.

Juniewicz, Isabel and Josh You, "New global AI computing capacity is doubling every 9 months," 2026.

Kalyani, Aakash, Nicholas Bloom, Marcela Carvalho, Tarek Hassan, Josh Lerner, and Ahmed Tahoun, "The Diffusion of New Technologies," *The Quarterly Journal of Economics*, May 2025, *140* (2), 1299–1365.

Katz, Lawrence F. and Kevin M. Murphy, "Changes in Relative Wages, 1963–1987: Supply and Demand Factors," *Quarterly Journal of Economics*, 1992, *107* (1), 35–78.

Keller, Wolfgang, "International technology diffusion," *Journal of economic literature*, 2004, *42* (3), 752–782.

Kim, Oliver, Karthik Tadepalli, and Joseph Levine, "What Will AI Do To Global Inequality?," 2026. https://www.global-developments.org/p/what-will-ai-do-to-global-inequality.

Korinek, Anton and Donghyun Suh, "Scenarios for the Transition to AGI," March 2024.

_ **and Joseph E. Stiglitz**, "Artificial Intelligence, Globalization, and Strategies for Economic Development," 2021. NBER Working Paper 28453.

_ **and Patrick McKelvey**, "Measuring the AI Economy," June 2026.

_ **, Charles I Jones, Szymon Sacher, Tess Cotter, and Peter McCrory**, "Economic Scenarios for Transformative AI," Technical Report 2026-02, Anthropic Institute 2026.

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

Liu, Yan, *Generative AI: Catalyst for Growth or Harbinger of Premature De-Professionalization?*, Washington, DC: World Bank, September 2024.

_ **, Jingyun Huang, and He Wang**, "Who on Earth Is Using Generative AI? Global Trends and Shifts in 2025," Technical Report, World Bank, Washington, DC October 2025.

Lokshin, Michael, *Cheap Energy Might Not Be Enough: A Trade Model of AI Compute Services*, Washington, DC: World Bank, July 2026.

McPeak, Grady, Anja Sautmann, Ohia George, Adham Hallal, Eduardo Arancón Simal, Aaron L. Schwartz, Jason Abaluck, Nirmal Ravi, and Robert Pless, "An LLM's Medical Testing Recommendations in a Nigerian Clinic: Potential and Limits of Prompt Engineering for Clinical Decision Support," in "2024 IEEE 12th International Conference on Healthcare Informatics (ICHI)" June 2024, pp. 586–591. ISSN: 2575-2634.

Mousa, Deena, "AI Could Keep Poor Countries Poor," 2026. https://newsletter.deenamousa.com/p/ai-could-keep-poor-countries-poor.

Murugan, Arul, Tomás Aguirre, Abhishek Nagaraj, and Rishi Bommasani, "The Jagged Global Economy: Frontier AI Unevenly Exposes National Economies," June 2026. arXiv:2607.05404 [cs.CY].

Ngai, L Rachel and Christopher A Pissarides, "Structural change in a multisector model of growth," *American economic review*, 2007, *97* (1), 429–443.

Nordhaus, William D., "Do Real-Output and Real-Wage Measures Capture Reality? The History of Lighting Suggests Not," in "The Economics of New Goods," University of Chicago Press, January 1996, pp. 27–70.

OECD, "Preliminary ODA levels in 2025," Technical Report 2025.

of Finance, Ministry, "The Revised National Budget: Preserving Norway's strong foundation in a turbulent world," May 2026.

O'Keefe, Cullen, Peter Cihon, Ben Garfinkel, Carrick Flynn, Jade Leung, and Allan Dafoe, "The Windfall Clause: Distributing the Benefits of AI for the Common Good," in "Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society" ACM New York NY USA February 2020, pp. 327–331.

Otis, Nicholas G., Rowan Clarke, Solène Delecourt, David Holtz, and Rembrand Koning, "The Uneven Impact of Generative Artificial Intelligence on Entrepreneurial Performance: Evidence from a Field Experiment in Kenya," *Management Science*, July 2026, p. mnsc.2024.06909.

Patel, Dwarkesh, Alex Imas, and Philip Trammell, "What remains scarce after AGI?," June 2026.

Porzio, Tommaso, Federico Rossi, and Gabriella Santangelo, "The Human Side of Structural Transformation," *American Economic Review*, August 2022, *112* (8), 2774–2814.

Pritchett, Lant, "Divergence, Big Time," *The Journal of Economic Perspectives*, 1997, *11* (3), 3–17.

Ransohoff, Nan, "The third wave of American philanthropy," May 2026.

Ravanilla, Nico, "Generative AI and Labor Displacement in the Global South:Evidence from Philippine Business Process Outsourcing," August 2026.

Restrepo, Pascual, "We Won't Be Missed: Work and Growth in the AGI World," in Ajay Agrawal, Erik Brynjolfsson, and Anton Korinek, eds., *The Economics of Transformative AI*, Chicago: University of Chicago Press, 2026. Forthcoming.

Rockall, Emma, Marina M Tavares, and Carlo Pizzinelli, "AI Adoption and Inequality," 2025.

Rossi, Federico, "The Relative Efficiency of Skilled Labor Across Countries: Measurement and Interpretation," *American Economic Review*, 2022, *112* (1), 235–266.

Sakai, Ando, Divya Kirti, Willa Lin, and Assyifa Szami Ilman, "How is Labor Demand Changing as AI Improves? New Global Evidence," Technical Report, IMF 2026.

Schaal, Jacob, "Distributing AGI's wealth worldwide is a very tricky problem," September 2026.

Shahidi, Peyman, Gili Rusak, Benjamin S. Manning, Andrey Fradkin, and John J. Horton, "The Coasean Singularity? Demand, Supply, and Market Design with AI Agents," in "The Economics of Transformative AI," University of Chicago Press, December 2025, pp. 145–165.

Shanahan, Matthew, "The State of Mobile Internet Connectivity 2026," Technical Report, GSMA 2026.

Trammell, Philip and Anton Korinek, "Economic Growth Under Transformative AI," *Annual Review of Economics*, 2026, *18*, 589–611.

UN, "Innovative voluntary financing options for artificial intelligence capacity-building," Technical Report A/79/966 2025.

Verasight, "Most Americans Use AI, but Oppose Building More Data Centers," 2026. https://data.verasight.io/ai/most-americans-use-ai-but-oppose-building-data-centers/.

Verhoogen, Eric, "Firm-Level Upgrading in Developing Countries," *Journal of Economic Literature*, December 2023, *61* (4), 1410–64.

World Bank, "The Promise of Artificial Intelligence," WDR 2026.

Yelizarova, Anna, "The Missing Institution: A Global Dividend System for the Age of Transformative AI," December 2025. Book Title: The Digitalist Papers, Volume 2: The Economics of Transformative AI.

Zeira, Joseph, "Workers, Machines, and Economic Growth," *The Quarterly Journal of Economics*, 1998, *113* (4), 1091–1117.

A. Details on simulations

Here we expand more precisely on our process for drawing the parameter combinations that generate the simulations presented and discussed throughout the paper.

We consider the case of a "typical" advanced country and a "typical" developing country. The TFP difference $\frac{Z_R}{Z_P}$ is set to generate a factor seven GDP gap between the advanced and developing country and the skill endowment $\frac{H_C+H_S}{H_L}$ is set to differ by a factor of seven, which roughly corresponds to differences between the US and India.

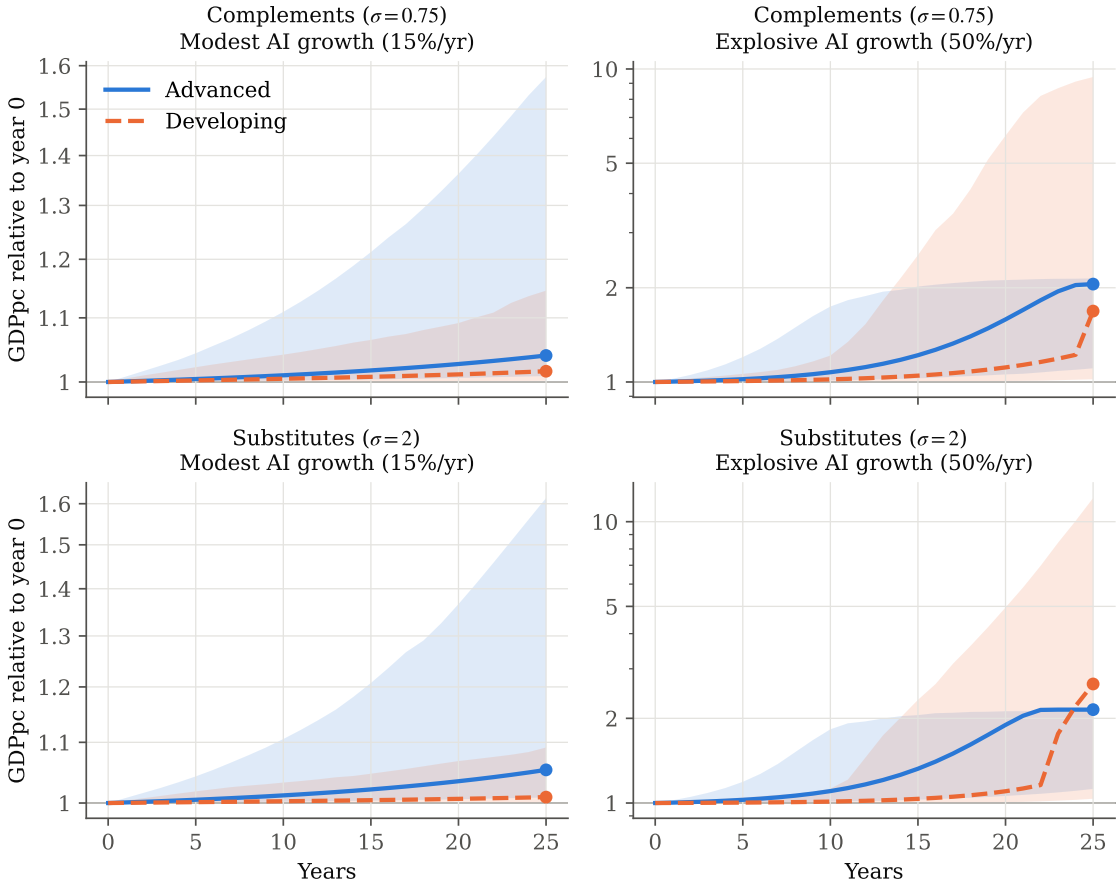
We assume that η_L and η_S are given by 1 plus a standard log-normal distribution so that they are strictly greater than 1 (as needed) and approximately 75 percent of draws fall between 1 and 3. $1/\eta_C$ is drawn from the same distribution so that η_C is strictly less than 1 and 75 percent of values fall between $1/3$ and 1. Note that while this procedure ensures that η_C is less than unity and η_L, η_S are greater than unity, both assumptions we make in deriving our analytic results, it does not guarantee that $\eta_S > \eta_L$, another assumption, holds. We assume that χ follows a uniform distribution, reflecting maximum uncertainty over the extent to which AI will substitute for or complement skills.

The final parameters to select are the various parameters α_ℓ. These parameters have no meaning on their own and largely map to (hypothetical) data via their impact on the AI cost shares θ_ℓ (conditional on prices). These could be measured today, but technology diffusion is still likely in progress — it is possible (maybe even likely) that AI cost shares would continue to rise relative to today *even if capabilities or pricing stopped advancing* simply because more firms and industries would figure out how to integrate existing capabilities into their workflows. Thus there is uncertainty even here. Our approach is to make draws for θ_ℓ and then select α_ℓ to match these values, and we select a log-normal distribution with log-mean $\ln(0.01)$ and log-standard deviation 1 so that the median cost share is 1 percent with a 95th percentile of 5 percent.[19]

B. Appendix Tables and Figures

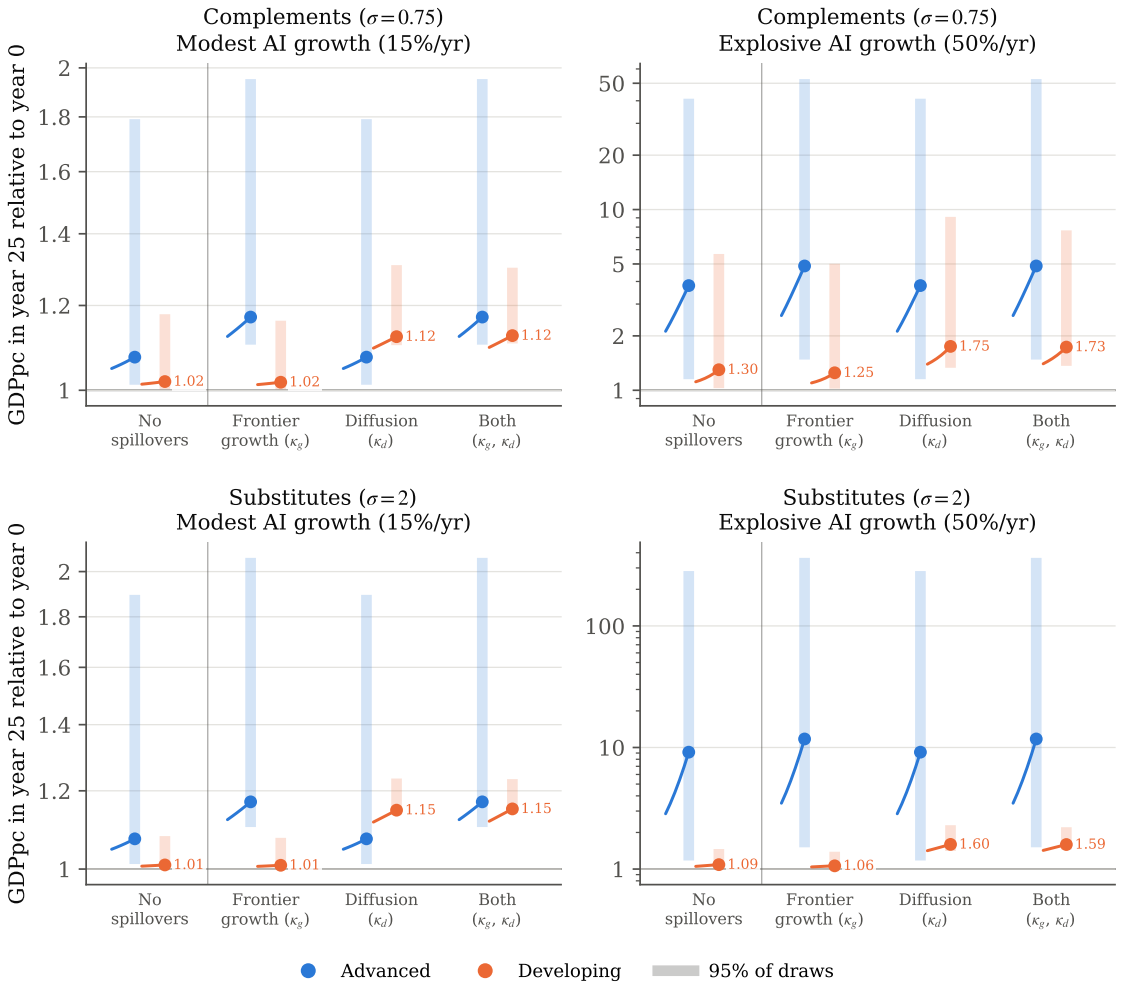
[19]Recent reports suggest that Anthropic and OpenAI are reaching annualized revenues of roughly 120 billion USD collectively. US GDP of about 32 trillion USD and a BLS labor share of about 53 percent suggest that total labor compensation is about 17 trillion USD, so a median of 1 percent is at least the correct order of magnitude.

Figure B1: GDP Growth: Advanced vs Developing (with Satiation) under $\varepsilon = 1.5$



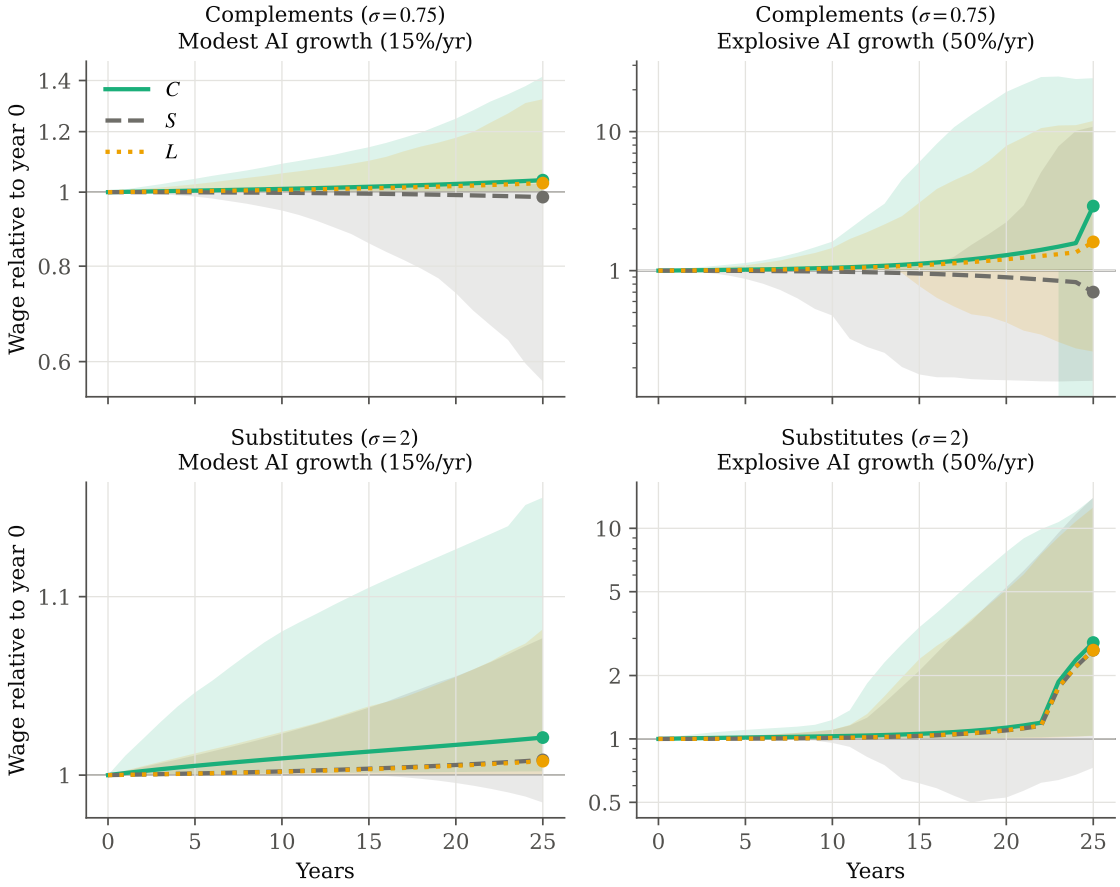
This figure plots GDP relative to its $t = 0$ value over 25 years in both the advanced and developing countries under the assumptions of modest vs explosive growth (left vs right columns) and gross complementarity or substitutability of AI (top vs bottom rows). Each path corresponds to the "representative" draw as described in the text, while the shaded areas contain 95 percent of all simulated scenarios. The parameters here feature satiation in demand for AI-produced goods and services at 1.5 times the advanced country's $t = 0$ consumption level. For this figure, we choose $\varepsilon = 1.5$.

Figure B2: Impact of Productivity Scenarios on Year 25 GDP



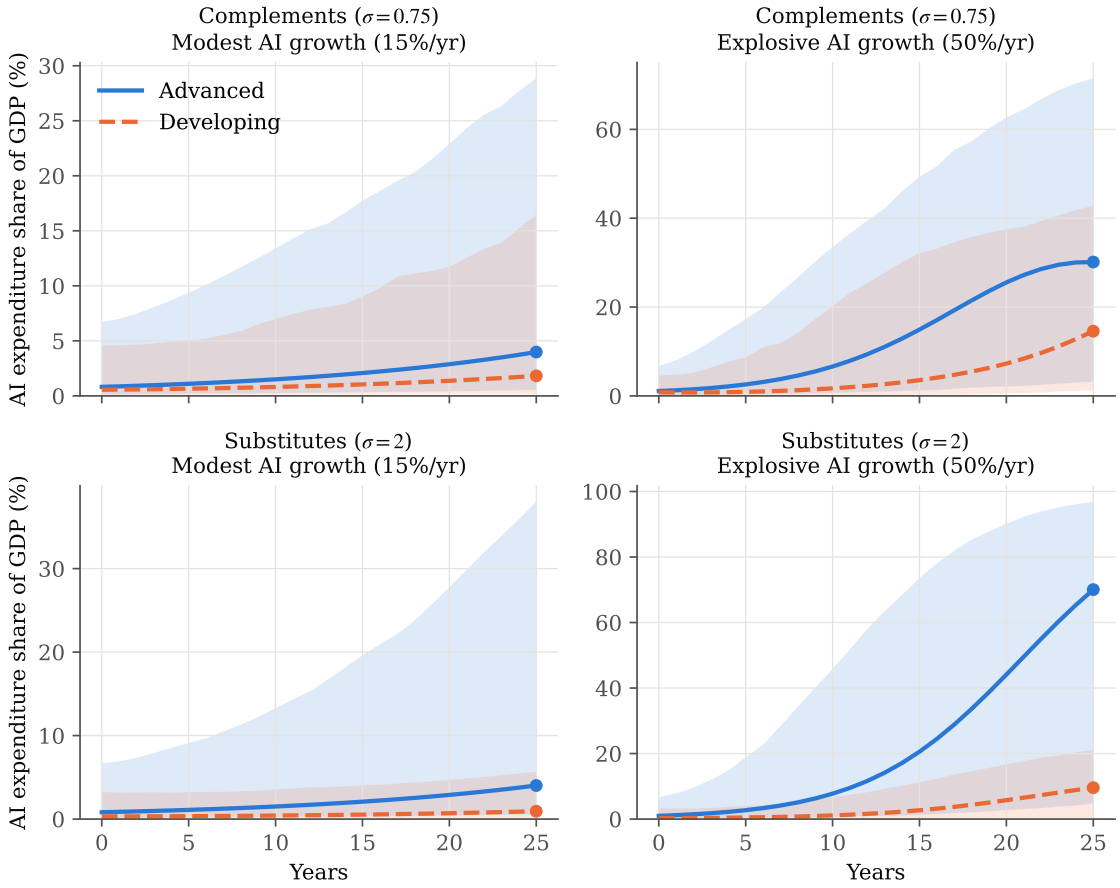
This figure displays the impact of AI growth on the $t = 25$ GDP (relative to the $t = 0$ level) for both the advanced and developing economies in the representative simulation under our extension adding frontier and diffusion growth in TFP. Each panel displays the results under a different set of assumptions regarding modest vs explosive growth (left vs right column) and gross complementarity or substitutability of AI (top vs bottom rows). The shaded bars contain 95 percent of simulations. The parameters here feature no satiation in demand for AI-produced goods and services.

Figure B3: Wage Growth in the Developing Economy (with Satiation) under $\varepsilon = 1.5$



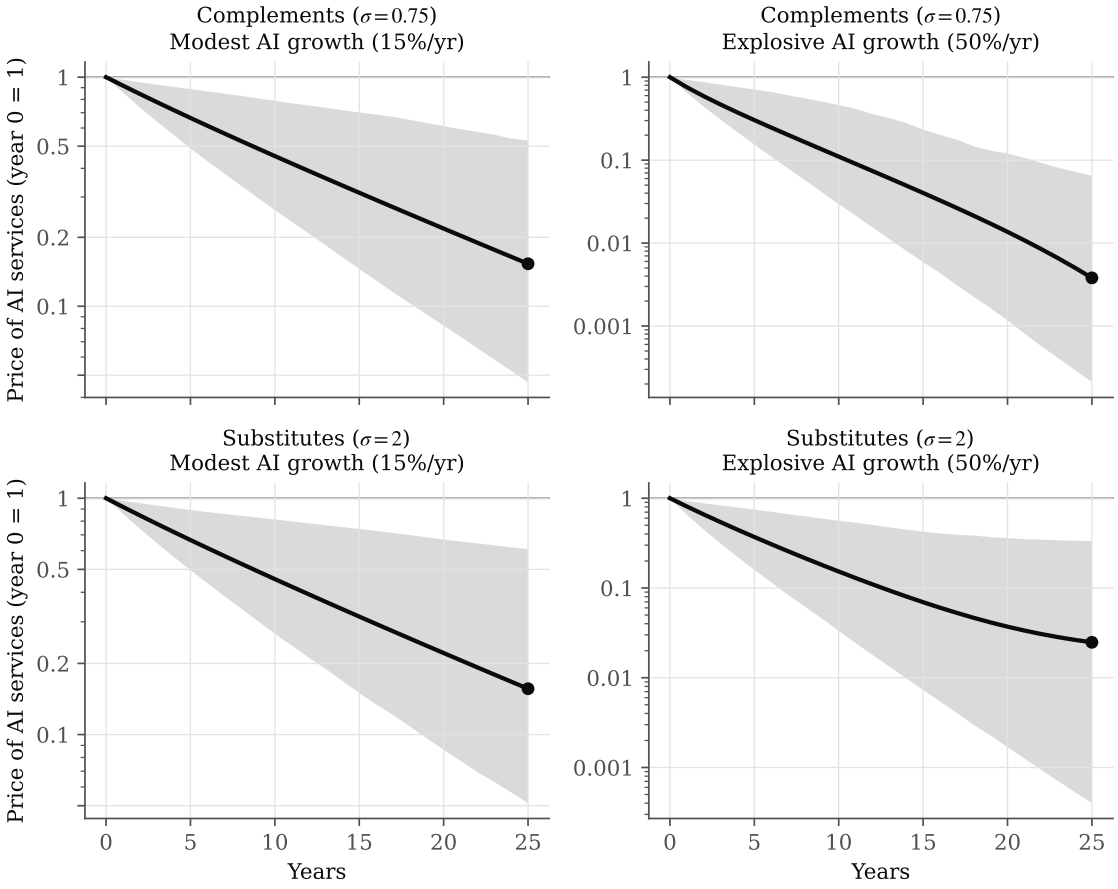
This figure plots wages relative to their $t = 0$ value over 25 years for all three labor types in the developing country under our assumptions of modest vs explosive growth (left vs right column) and gross complementarity or substitutability of AI (top vs bottom rows). Each path corresponds to the "median" draw as described in the text, while the shaded areas contain 95 percent of all simulated scenarios. The parameters here feature satiation in demand for AI-produced goods and services at 1.5 times the advanced country's t = 0 consumption level.

Figure B4: AI Expenditure Share of GDP: Advanced vs Developing (No Satiation)



The figure displays spending on AI services as a share of total spending on production (i.e. input costs) in both the advanced and developing economy. Each panel displays the results under a different set of assumptions regarding modest vs explosive growth (left vs right column) and gross complementarity or substitutability of AI (top vs bottom rows). The shaded bars contain 95 percent of simulations. The parameters here feature no satiation in demand for AI-produced goods and services.

Figure B5: Global Price of AI Services (No Satiation)



The figure displays the global price of AI services in our simulations. Each panel displays the results under a different set of assumptions regarding modest vs explosive growth (left vs right column) and gross complementarity or substitutability of AI (top vs bottom rows). The shaded bars contain 95 percent of simulations. The parameters here feature no satiation in demand for AI-produced goods and services.