NBER WORKING PAPER SERIES

ECONOMIC GROWTH UNDER TRANSFORMATIVE AI

Philip Trammell
Anton Korinek

Working Paper 31815
http://www.nber.org/papers/w31815

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2023, Revised September 2025

We thank Tim Blomfield, Nicholas Emery-Xu, and Don Suh for research assistance. An earlier version of this review, written in part for audiences with technical backgrounds outside of economics, explains the relevant economic concepts as they are introduced (e.g. elasticity of substitution) and explores some of the scenarios in more detail. It may be found at https://philiptrammell.com/static/egtai_old.pdf. Trammell performed part of this work as a Visiting Graduate Fellow at the Digital Economy Lab, Stanford University. Korinek performed part of this work while he was a Visiting Fellow at the Brookings Institution. The views expressed in this work do not necessarily reflect the views of the Brookings Institution. Korinek also acknowledges support from the Darden School of Business at the University of Virginia. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Economic Growth under Transformative AI
Philip Trammell and Anton Korinek
NBER Working Paper No. 31815
October 2023, Revised September 2025
JEL No. E2, O3, O4

ABSTRACT

Recent advances in AI may herald the near arrival of systems that can automate essentially all work. We review the macroeconomic implications of this scenario, in a framework synthesizing several strands of the relevant literature. Robustly, fully automating production alone (so that machines can self-replicate) would dramatically raise the growth rate and lower the labor share, breaking the Kaldor Facts that have long characterized frontier growth. Automating R&D (so that machines can self-improve) would accelerate the transformation, but may not produce it in isolation. Wages—multiplying exploding output and a plummeting labor share—may rise or fall, depending on the returns to scale, the importance of natural resources, and the direction of technical change.

Philip Trammell
Stanford University
trammell@stanford.edu

Anton Korinek
University of Virginia
Department of Economics and EconTAI
and NBER
akorinek@virginia.edu

1 Introduction

Recent advances in artificial intelligence have raised the possibility of an economic transformation as profound as the Industrial Revolution. AI models have demonstrated rapidly expanding capabilities in cognitive tasks, from natural language processing to scientific discovery, with slower but ongoing progress in robotics extending these capabilities to physical labor. The 2024 Nobel Prize in Chemistry awarded to AI pioneer Demis Hassabis for an AI system that can predict protein folding exemplifies how AI is already accelerating scientific research. Future models may reach the level of Artificial General Intelligence (AGI)—the ability to perform all cognitive tasks at least as well as skilled humans—and in the process, unlike any technology yet developed, could fully automate the process of research and development itself, including the creation of better AI systems.

If AI capabilities approach this threshold, the resulting economic transformation could break the stylized facts that have characterized industrial growth for more than two centuries. Since Kaldor (1957) first identified them, two empirical regularities have served as cornerstones of growth theory: an approximately constant long-run growth rate in output per capita, and an approximately constant labor share of output. These "Kaldor Facts" have proven remarkably robust across industrialized countries and subsequent decades, shaping much of modern growth theory: see e.g. Acemoglu (2009) for an overview. Yet when viewed over longer time horizons, the phenomenon of constant exponential growth appears as a recent historical anomaly. The global economic growth rate unambiguously accelerated with the onset of the Industrial Revolution; we have no reason to insist that it cannot accelerate again.

This paper evaluates the channels through which sufficient advances in AI may affect economic growth by synthesizing several strands of the literature within a common framework. We examine two fundamental questions: Under what conditions would AI break the regime of constant exponential growth that has characterized the industrial era? And what would be the implications for factor shares and wages? We distinguish between the automation of production and the automation of research and development and explore their individual and combined effects.

Our analysis reveals several robust findings. First, if AI enables capital to become a gross substitute for labor in production alone, allowing machines to self-replicate, the growth rate would dramatically accelerate and the labor share would plummet toward zero, decisively breaking both Kaldor Facts. The magnitude of this acceleration depends critically on the degree of substitutability and the role of fixed factors like natural resources. Second, automating R&D in isolation has more ambiguous effects: it may accelerate growth if automated research proves highly parallelizable, but it could also leave growth rates bounded if fundamental constraints on the research process persist. Third, the impact on wages is deeply uncertain, depending on the returns to scale, the importance of natural resources, and most critically, the direction of technical change.

In recent years, economists have begun to engage earnestly in theoretical explorations of these transformative possibilities.[1] We aim to synthesize the findings of the most influential papers on the topic, covering both the insights of different categories of models and the underlying mathematical intuition. We have simplified some models and adopted common notation to crystallize the mechanisms at work within a unified framework. While we do not offer assessments of the probability of particular scenarios, our analysis clarifies the economic conditions under which AI could trigger a transformation as great as the Industrial Revolution or greater.

The remainder of this paper is organized as follows. Section 2 examines the implications of automating production, showing how increased substitutability between capital and labor could drive explosive growth and a vanishing labor share. Section 3 analyzes the automation of R&D, exploring conditions under which it could sustain or accelerate growth. Section 4 considers the combined effects of automating both sectors. Section 5 summarizes. Section 6 concludes with discussions of the transition from current AI capabilities to transformative automation, policy implications for managing radical inequality, and the existential risks that may accompany economically transformative AI.

2 Automating production

Consider the standard constant returns to scale (CRS), two-factor production function in capital K and labor L, with factor-augmenting technology terms A and B. Fixing

[1] Sandberg (2013) presented an "overview of models of technological singularity" over a decade ago, before mainstream economists started to analyze the transformative potential of AI. Most of the models he summarizes therefore do not attempt to spell out how AI or other transformative technologies would fit into standard economic models to produce the results in question. Our review fills this gap.

labor supply at $\bar{L}$,

$$Y_t = Y(A_t K_t, B_t \bar{L}).$$

Today, labor and capital are at least *weakly gross complements*: their elasticity of substitution ε is not greater than one, or equivalently, their substitution parameter $\rho \equiv \frac{\varepsilon-1}{\varepsilon} \leq 0$ (Gechert et al., 2022). While this remains, sustained growth requires not only accumulating (effective) capital but developing labor-augmenting technology, as the developed world has done since the Industrial Revolution at a rate of ~2% per year.[2] Because capital accumulates with output whereas labor does not, and output is bottlenecked by the scarcer factor, labor-augmenting technology tends to be the more profitable kind to develop (Acemoglu, 2003). If all technology growth is labor-augmenting and the saving rate is constant, the effective capital stock grows with output. CRS ensures that Y, AK, and $B\bar{L}$ all grow at the same rate, and that the [competitive] labor share does not change (Uzawa, 1961).

Regardless of the microfoundation, therefore, the central question is whether and when AI will allow capital to become a close enough substitute for labor to drive growth on its own. We call such AI *transformative AI*, or TAI. Under almost any standard assumptions, what follows is a radical acceleration to growth and decline in the labor share.

2.1 Growth and labor share

The implications of gross substitution for output and factor shares are most straightforwardly illustrated by the extreme case of perfect substitution:

$$Y_t = A_t K_t + B_t \bar{L}.$$

A now has a simple interpretation: how many machines a machine can build per year. Even fixing A, capital accumulation drives exponential output growth. Unless the saving rate is negligible or A is very low, so that capital replicates slowly, output grows

[2] In broad strokes. Jones (2022) attributes 1.3% of the US's long-run 2% growth in output per person to labor-augmenting technology growth, and most of the rest to growth in human capital.

rapidly with AK and the labor share

$$\frac{B\bar{L}}{AK + B\bar{L}}$$

falls to zero.[3]

Furthermore, there are now strong incentives to develop capital-augmenting technology. Indeed, because capital is ultimately more plentiful and is no longer bottlenecked by labor, it is typically efficient *only* to develop capital-augmenting technology, if both kinds are equally costly (again, see Acemoglu (2003)). If A grows, output grows superexponentially: the robots double their numbers in year one, triple their numbers in year two, and so on. The labor share falls to zero even more rapidly.

Saving. By no-arbitrage, the interest rate has to equal the return on capital, A. As the interest rate rises, the saving rate may rise or fall. The rise of automation may also move the saving rate by shifting income from low to high savers or vice-versa: Berg et al. (2018) emphasize that those who make most of their incomes from wages currently empirically exhibit lower saving rates than those who make most of their incomes from capital. Sachs and Kotlikoff (2012), Sachs et al. (2015), and Benzell and Ye (2024) focus on the case in which retirees with high discount rates hold most capital.

Given any constant (or bounded) inverse elasticity of intertemporal substitution η, however, the saving rate cannot fall quickly enough to defuse superexponential growth. This follows from an investor's Euler equation. Letting δ denote the discount rate, c_t denote consumption per person at t, and R_t denote the cumulative interest factor earned by investing from 0 to t,

$$c_0^{-\eta} = e^{-\delta t} R_t c_t^{-\eta} \quad \implies \quad c_t / c_0 = e^{-\frac{\delta}{\eta} t} R_t^{\frac{1}{\eta}}.$$

Observe that, with constant $g_A > 0$,

$$R_t = e^{\int_0^t A_\tau d\tau} = e^{\frac{A_0}{g_A}(e^{g_A t} - 1)}.$$

Since R_t grows double-exponentially, so does c_t / c_0. This is defused only if the elastic-

[3]This pair of points has been made in many contexts. See Prettner (2019) for an especially straightforward treatment.

ity of intertemporal substitution ultimately falls as consumption rises, with $\eta(c) \to \infty$ sufficiently quickly.

Gross substitution in capital goods, Cobb-Douglas in consumption. We have implicitly assumed that the production functions for capital and consumption are the same. In this case, capital-driven growth requires gross substitution in production across the board, and any acceleration to growth is coupled with a decline in the labor share, as we have seen. This coupling is quite robust, for reasons Berg et al. (2018) elaborate, but it is not strictly necessary. This is because rapid growth can be driven by self-replicating *capital*, as Mookherjee and Ray (2022) emphasize, even if the capital is not quite a gross substitute for labor in the consumption goods sector.

Suppose that capital comes to be a gross substitute for labor in the production of capital goods: e.g., for simplicity, that

$$\dot{K} = sAK - dK,$$

where s and d are the saving and depreciation rates and A is high enough to sustain rapid growth ($sA - d \gg 0$). If capital and labor are also gross substitutes in the production of consumption goods, the labor share falls as capital accumulates, as noted above; if they are gross complements, consumption growth remains bottlenecked by labor-augmenting technology. However, in the edge-case in which the production of consumption goods takes the Cobb-Douglas form

$$C = ((1 - s)K)^{\alpha} \bar{L}^{1-\alpha},$$

the labor share is constant at $1 - \alpha$, even as consumption grows at rate $\alpha g_K = \alpha(sA - d)$: so, exponentially if A is fixed and superexponentially if A rises. Because capital here does not face diminishing returns to its own replication, the proportional capital depreciation process does not force the capital stock to a finite steady state in the absence of technological development.

Rapid, capital-driven growth alongside a constant labor share is thus a technical possibility. This scenario might even be considered a limiting case of the tendency over recent decades for manufacturing to grow more capital-intensive while demand for labor in consumption services remains high (Autor and Dorn, 2013). Nevertheless, because it relies heavily on the Cobb-Douglas edge case, we will put it aside going forward, focusing on the case that machines capable of driving rapid growth are gross substitutes for labor in general, and thus drive the labor share to zero.

2.2 Wages

The impact of TAI on the wage bill is more deeply ambiguous.

Rising wages given capital-augmenting technology. Under a CES production function

$$\begin{aligned} Y_t &= \left(a(AK)^\rho + (1-a)(BL)^\rho\right)^{1/\rho} \\ \implies dY_t/dL_t &= Y_t^{1-\rho} B_t^\rho \, (1-a)\bar{L}^{\rho-1}, \end{aligned}$$

wages stagnate, or grow at their original slow rate g_B, only under perfect substitution ($\rho = 1$). Under gross but imperfect substitution ($\rho \in (0,1)$), as Nordhaus (2021) notes, wages rise rapidly with output even fixing B. Capital grows immensely plentiful and complements labor, however slightly.

More generally, in any $\mathcal{D}^2$ CRS production function $Y(AK, BL)$ with (weakly) diminishing returns to capital, increasing the quantity or augmentation of capital increases (does not affect) the marginal product of labor. Letting $\hat{K} \equiv AK$, $\hat{L} \equiv BL$ and differentiating $\hat{K}F_{\hat{K}} + \hat{L}F_{\hat{L}} = Y$ with respect to $\hat{K}$,

$$\begin{aligned} F_{\hat{K}} + \hat{K}F_{\hat{K}\hat{K}} + LF_{\hat{L}\hat{K}} &= F_{\hat{K}} \\ \implies F_{\hat{L}\hat{K}} &= -\frac{\hat{K}}{\hat{L}} F_{\hat{K}\hat{K}} \geq 0. \end{aligned}$$

The introduction of humanoid robots is analogous to an increase in population, due e.g. to immigration. If immigrants have the same distribution of skills as the native population ($\rho = 1$), under constant returns to scale wages do not change, once the capital needed to complement the larger population has accumulated. If immigrants have a different distribution of skills ($\rho < 1$), comparative advantage allows gains from trade. Wages may fall for native workers in the occupations at which the immigrants are most relatively skilled, but the aggregate wage bill for natives must rise.

Labor-depleting technology. The observation above might suggest that falling wages due to technological progress are only ever a partial equilibrium phenomenon. However, automation can also be labor-depleting.

For illustration, suppose that goods can be produced in two sectors: (i) an AK robot-only sector and (ii) a Cobb-Douglas sector employing labor and equipment. Letting z denote the equipment share of capital,

$$Y = A(1-z)K + (zK)^{\alpha}(B\bar{L})^{1-\alpha}. \tag{1}$$

The marginal product of capital in sector (i) equals A, and in sector (ii) it equals

$$\alpha(zK)^{\alpha-1}(B\bar{L})^{1-\alpha}. \tag{2}$$

Let z^* denote the efficient equipment share of capital. If $A \leq$ (2) even at $z = 1$, then $z^* = 1$. On a balanced growth path, with B growing and A fixed, the marginal product of capital in each sector does not change, and $z^* = 1$ is maintained. Following a sufficient increase in robot productivity A, however, z^* equalizes the marginal product of capital across sectors:

$$A = \alpha(z^*K)^{\alpha-1}(B\bar{L})^{1-\alpha} \quad \Longrightarrow \quad z^* = \left(\frac{A}{\alpha}\right)^{\frac{1}{1-\alpha}} \frac{B\bar{L}}{K}.$$

Substituting z^* for z into (1) and simplifying,

$$Y = AK + \tilde{B}L,$$

where $\tilde{B} \propto A^{-\frac{\alpha}{1-\alpha}}B$.

Further increasing robot productivity thus expands the production possibilities frontier in a way that amounts not only to capital augmentation but to labor *depletion*, since capital is reallocated in equilibrium away from labor-augmenting uses. Wages rise ($g_{\tilde{B}} > 0$) only if productivity improvements in sector (ii) proceed quickly enough to offset any such reallocation: that is, only if $g_B > \frac{\alpha}{1-\alpha} g_A$.

Analogy to an expanding varieties model. For another angle on why the impact on the wage bill is ambiguous, consider an expanding-varieties growth model such as Romer (1990). The share of any particular variety falls to zero as variety expands, but expenditures on it are constant, even though the new varieties are modeled as gross substitutes for the old.

The precise constancy is of course not a logical necessity, but it reveals that the development of new gross substitutes does not necessarily raise or lower real expenditures on pre-existent goods and services. Horses were used for a larger share of tasks in the past, before being eclipsed by motor vehicles, and sailboats were used for a larger share of tasks before being eclipsed by motorboats; but total real expenditures on horses and sailboats have since fluctuated with technology and taste without approaching zero (see e.g. American Horse Council, 2023). Korinek (2026) conjectures that humans may remain in demand after full automation in services where consumers value authentic human connection (e.g., childcare); competitive or performative activities in which the human identity of the performer matters (sports and the arts); religious services; and human oversight as the ultimate arbiter of values for AI alignment. In a world of TAI we will doubtless be old-fashioned, but economic theory cannot tell us the future relative prices of the old-fashioned services provided by humans.

Finally, it must be emphasized that even if the aggregate wage bill stagnates or rises, wage inequality may grow such that many people cannot earn enough labor income to live on. Sustained demand for the world's best athletes, for example, could keep the aggregate wage bill positive or rising while the majority, competing with robots as in the model of labor-depleting technology above, find their wages driven to zero.

2.3 Endogenous automation

So far we have treated TAI as an exogenous technological advance. The example above illustrates that this advance may consist of a decline in the cost of robots, but it may also follow endogenously from rising wages.

Suppose

$$Y = A(zK)^{\alpha}\big(L + \theta(1-z)K\big)^{1-\alpha}, \quad \theta > 0, \tag{3}$$

where z again denotes the equipment share of capital. On a balanced growth path, fixing $z = 1$, the marginal product of equipment is constant, whereas the marginal product of "labor"—human or robotic—grows exponentially. Since the cost of a robot is fixed at $1/\theta$ units of output, they are built ($z^* < 1$) only once the marginal value of converting equipment to robots is positive. In particular, once

$$\frac{K}{L} > \frac{1}{\theta}\frac{\alpha}{1-\alpha}, \tag{4}$$

capital and labor become perfect substitutes at the margin. Maintaining an efficient allocation of capital, the production function reduces to

$$Y = A\big(\theta^{1-\alpha}K + \theta^{-\alpha}L\big)\frac{(1-\alpha)^{1-\alpha}}{\alpha^{\alpha}}.$$

As usual, fixing robot productivity θ, output grows superexponentially—due to a combination of rising A and exponential capital accumulation at fixed A—and the labor share falls to zero. Capital accumulation no longer raises wages, as capital and labor are now fungible on the margin. Technological progress does raise wages, fixing θ, but more generally wages depend on the direction of development; they fall if robot productivity increases so much faster than TFP A that $g_A - \alpha g_\theta < 0$.

The point that rising wages incentivize automation dates at least to Habakkuk (1962). Hémous and Olsen (2022) offer a recent task-based treatment of the effect along with a review of the relevant evidence. Hanson (2001) and Sasaki (2023) observe that this logic applies no less to hypothetical machines that can automate all tasks, inspiring the model above.

In short, though anticipations of full automation in the near term are driven by the possibility of rapid increases in machine capabilities per unit cost (roughly θ), these are not necessary for full automation to occur eventually. As long as it is possible at *some* cost to build machines that can do all work, rising wages will eventually render them profitable, and then ever more so. This suggests that a regime of rapid growth and a plummeting labor share will likely prevail at some point in the future, whether or not economic data suggest that it is imminent. It will be difficult to permanently fulfill the hope (Korinek and Stiglitz, 2020; Acemoglu and Restrepo, 2020; Brynjolfsson, 2022) of steering technological progress so that both the growth rate and the labor share stay bounded above zero, let alone with the labor share remaining close to its current level.

2.4 A fixed factor

A common objection to the feasibility of rapid, automation-driven growth is that natural resource constraints will halt growth in general. As Korinek and Stiglitz (2019) discuss, a third factor in fixed supply which is a gross complement in production eventually bottlenecks output growth even under TAI; and as the effective quantity of machinery grows, the share of the fixed factor rises to one.

Such a hard resource constraint is unlikely to bind in the foreseeable future. The natural resource share is below 5% and not appreciably rising (Eden and Kuruc, 2024). We continue to consume energy less than one five-thousandth as quickly as sunlight reaches the earth's surface (World Energy Council, 2013),[4] and the "rare-earth metals" used pervasively in electronic and other industrial equipment are prevalent in low concentrations throughout the Earth's crust, so should not permanently bottleneck growth as long as there is the energy to mine and refine them (Gosen et al., 2017). Nevertheless, the natural resource share is positive, and it is optimistic to suppose that resource constraints will place no drag on growth at all, as in the case of gross substitution. Following Hanson (2001), we therefore focus on the Cobb-Douglas case.

Growth and labor share. Let $\bar{W}$ denote the fixed factor, and call it *land* for brevity. Then output equals

$$Y = A\,F(K, \bar{L})^{\bar{\alpha}}\ \bar{W}^{1-\bar{\alpha}}, \tag{5}$$

where $F(\cdot)$ is CRS, and its substitution parameter between K and L is non-positive before the introduction of TAI and positive (perhaps one) after.

This case is equivalent to the Solow model given fixed population, with *extended capital* $F(K, \bar{L})$ taking the place of capital and land taking the place of labor. Given constant $g_A > 0$, output grows exponentially (fixing saving and depreciation rates): technology growth no longer replaces an exponential capital accumulation process with a superexponential one, but a bounded process with an exponential one. The land and extended capital shares are fixed at $1 - \bar{\alpha}$ and $\bar{\alpha}$ respectively, and because labor constitutes an ever smaller share of extended capital, its share falls to zero.

Though the land constraint prevents transformative AI from growing output superexponentially, if $\bar{\alpha}$ is near one this is a technicality. Because capital accumulates in line with past output, the steady-state growth rate satisfies

$$g_Y = g_A + \bar{\alpha} g_Y \quad \Longrightarrow \quad g_Y = \frac{1}{1-\bar{\alpha}} g_A,$$

[4] See also Our World in Data (2025) for updated energy consumption figures.

which rises without bound as $\bar{\alpha} \to 1$. Under a land share of 5%, automation raises the capital share from 33% to 95% and the steady-state growth rate by more than an order of magnitude.

Wages. The implications for wages are less sanguine. Without a land constraint, Section 2.2 noted that:

- If any technological progress is labor-augmenting, wages grow even under perfect substitution.
- Even if all technological progress is capital-augmenting, the exploding capital stock drives rapid wage growth if capital and labor are gross but imperfect substitutes.

With a land constraint, any complementarity between labor and capital is largely offset by the fact that a rising capital stock lowers the marginal product of extended capital:

- Even if technological progress augments both labor and capital equally, as in (5), the marginal product of (extended) capital—and thus the wage—is stagnant in steady state, governed by savers' discount rates and elasticities of intertemporal substitution.
- If $F(AK, B\bar{L}) = \left(a(AK)^\rho + (1-a)(B\bar{L})^\rho\right)^{1/\rho}$ and all technological progress is capital-augmenting, wages

$$Y^{1-\rho/\bar{\alpha}}\,(1-a)\bar{B}^\rho \bar{L}^{\rho-1}$$

  rise only if

$$\bar{\alpha} > \rho.$$

  Perfect substitution ($\rho = 1$) guarantees that, as the marginal product of *effective* capital is asymptotically constant, the marginal products of capital and labor fall.

Skilled labor as a fixed factor. The logic above applies equally if we interpret certain worker skills as the fixed factor $\bar{W}$ (with unskilled labor as $\bar{L}$). Using a similar model but with this reinterpretation, Hémous and Olsen (2022) attribute rising wage inequality to automation. As discussed in Section 2.2, an increase in effective capital must raise the aggregate wage bill when labor and capital are the only factors, but can lower wages for those doing jobs for which capital is more substitutable than average. If workers can shift away from these occupations, their wages rise in the long run, but this is not possible if the occupations with rising wages demand skills they cannot acquire.

Many have attributed growing wage inequality throughout the developed world in recent decades to automation. It is worth emphasizing that differences in capital's substitutability for different kinds of work can just as easily make wages more equal. Indeed, AI to date appears have to widened the productivity distribution in some domains (including debate (Roldán-Monés, 2024) and developing-world entrepreneurship Otis et al. (2024)) but narrowed it in others (including coding (Peng et al., 2023), customer service (Brynjolfsson et al., 2025), writing (Noy and Zhang, 2023), product design (Dell'Acqua et al., 2023), advertising (Chen and Chan, 2025), and legal analysis (Choi and Schwarcz, 2025)).

2.5 A task-based framework

Task-based models provide a more granular framework for understanding how automation reshapes production by decomposing production into distinct activities that can be performed by either labor or capital. Building on pioneering work by Zeira (1998) and Acemoglu and Restrepo (2018), this literature examines how the progressive automation of tasks affects factor demand, income distribution, and economic growth. The long-run implications of automation for growth, factor shares, and wages in a task-based model are most clearly illustrated by Aghion et al. (2019).

Suppose output requires performing a (unit) continuum of symmetric tasks:

$$Y = \left(\int_0^1 Y_i^\rho \, di \right)^{1/\rho}, \quad \rho < 0. \tag{6}$$

Tasks $i \leq b \in [0, 1]$ are automatable, in that they can be performed by capital or labor, and tasks $(b, 1]$ are not, in that they can only be performed by labor:

$$Y_i = \begin{cases} K_i + L_i, & i \leq b; \\ L_i, & i > b. \end{cases}$$

Given factor stocks K and $\bar{L}$, so long as there is at least as much capital per automatable task as labor per non-automatable task—i.e. as long as

$$\frac{K}{b} \geq \frac{\bar{L}}{1-b} \tag{7}$$

—it is inefficient to use labor on an automatable task. Then, by symmetry, capital is spread equally across automatable tasks so K/b units of capital are used on each, and by the same token, $\bar{L}/(1-b)$ units of labor on each non-automated task. Substituting for Y_i in (6) yields

$$Y = \left[b\left(\frac{K}{b}\right)^{\rho} + (1-b)\left(\frac{L}{1-b}\right)^{\rho} \right]^{1/\rho} = [(AK)^{\rho} + (BL)^{\rho}]^{1/\rho}, \tag{8}$$

where $A = b^{\frac{1-\rho}{\rho}}$ and $B = (1-b)^{\frac{1-\rho}{\rho}}$.

Since automatability allows capital to perform more tasks, one might imagine that it is equivalent to capital-augmenting technology. Here, however, it is actually equivalent to *labor*-augmenting—and capital-depleting!—technology.[5] As b rises from 0 to 1, A falls from ∞ to 1 and B rises from 1 to ∞. In particular, $g_B = \frac{1-\rho}{\rho} g_{1-b}$, so that B and output grow exponentially if a constant fraction of the remaining non-automatable tasks are automated each year.

For intuition, observe that each time $1-b$ halves, capital can be spread more thinly across the widened range of automatable tasks, but labor is concentrated twice as heavily in each non-automatable task. Automation therefore allows capital to serve as a *better complement* to labor. A marginal unit of labor is spread across fewer non-automatable tasks, producing a larger increase to the supply of each. Given abundant capital, this produces a larger increase to output.

Transformative automation. Automating all tasks is plainly transformative, rendering capital and labor perfect substitutes. More subtly, Korinek and Suh (2024) show that if automation outpaces capital accumulation, it renders the two factors perfect substitutes on the margin so that (7) fails (or holds with equality). If it does so permanently, growth proceeds as in Section 2.1. Once automation is complete, or advancing rapidly enough that accelerating it would no longer allow resources to be allocated more efficiently, it will be valuable to turn instead to increasing the replication rate of

[5] Jones and Liu (2024) offer a similar model in which explicit capital-augmenting technology renders the reduced-form capital-augmenting technology term constant at all times, not just asymptotically.

capital.

Wages after transformative automation. Korinek and Suh (2024) calibrate a production function similar to (8) to macroeconomic data and study the implications of b rising to 1 in finite time. They find that wages rise and then fall, stagnating at a low level once the factors are substitutes.

The conclusion that wages rise initially is robust, since on the current margin, an increase to b is an increase to labor-augmenting technology. The point that wages may ultimately fall—simply by accelerating the process that has driven centuries of wage growth—is also important to recognize, but here depends on precisely how the transition to full automation unfolds. Because we do not observe factor-augmenting technology terms directly, any calibration of B^ρ and ρ in the reduced-form production function (8)—suggesting that the wage will equal B once $\rho = 1$—can also be interpreted as a calibration of $a_1(a_2B)^\rho$ and ρ in a production function of the form

$$Y = [(AK)^\rho + a_1(a_2BL)^\rho]^{1/\rho}$$

for arbitrary $a_1 > 0$ and $a_2 = a_1^{-1/\rho}$. The second calibration, however, would suggest that the full-automation wage will equal a_1a_2B.

New tasks and the limits of task creation. While the framework above treats the task space as fixed, Acemoglu and Restrepo (2018) and subsequent work emphasize that technological progress has historically created new tasks alongside automating existing ones. Autor et al. (2024) documents how the emergence of new work has offset automation's displacement effects. We perform myriad tasks today that would have been inconceivable a century ago—from web development to data science—though on some level this may represent specialization within broader categories: most contemporary white-collar professions could be viewed as specialized descendants of what might once have been termed "scribe" or "philosopher". Naive extrapolation from current automation rates of O*NET-catalogued tasks may substantially underestimate the distance to transformative AI, by failing to account for the continuous emergence of new, often more complex tasks.

Yet there exists a fundamental constraint: humans possess some finite set of cognitive and physical capabilities, and once machines can replicate these underlying capabilities, they will necessarily be able to perform all feasible human tasks, including any newly created ones. The task-based framework above clarifies that once this capability threshold is crossed and condition (7) binds or fails, the economy transitions to a regime of rapid growth and declining labor share, regardless of whether new tasks continue to emerge.

3 Automating R&D

So far we have taken technological development to be constant and exogenous (beyond noting that once labor is no longer a bottleneck, it will be efficient to invest in developing capital-augmenting technology). Even so, we have seen that AI that allows capital to sufficiently automate production would probably greatly accelerate growth. The growth impact of automating technological development but not production is more ambiguous. Through this channel in isolation, AI is more likely to have only a moderate impact on long-run growth, but also more likely to have an impact more explosive than any considered in Section 2.

Many model-based analyses of the growth impacts of AI have tended to focus on impacts on production rather than R&D (see e.g. Acemoglu, 2025; Aghion and Bunel, 2024). This is in part because production is better understood, but also in part because R&D tends to be high-skilled, and historically, low-skilled work has been more exposed to automation. In the case of AI, this historical regularity may soon no longer apply. High-skilled occupations are disproportionately LLM-exposed, and research occupations in particular: Eloundou et al. (2024) find that out of 11 job groupings, “scientists and researchers” are most exposed and “technologists” are next. At the firm level, AI contributes to growth primarily by speeding product innovation (Babina et al., 2024). Many of AI’s most important achievements to date have been scientific, such as the protein folding breakthrough that won AI pioneer Demis Hassabis a Nobel Prize in chemistry. Nor should this be surprising: the output of an AI model query is a non-rival packet of information, not a widget.

3.1 Automating R&D in a semi-endogenous model

Assume that capital and labor remain gross complements in production, so that approximately

$$Y \propto AL.$$

In a standard semi-endogenous growth model (Jones, 1995),

$$g_A \propto A^{-\beta} L^{\lambda} \ : \tag{9}$$

$\beta > 0$ captures the extent to which ideas are harder to find when the frontier is further advanced, and $\lambda < 1$ captures the returns to scale of research work, i.e. the extent to which research is imperfectly parallelizable. Because $\beta > 0$, long-run technology growth requires population growth. In steady state, $g_A = g_L \, \lambda/\beta$.

If AI allows capital to serve as a gross substitute for labor in R&D, we have roughly

$$\begin{aligned} g_A &\propto A^{-\beta} K^{\lambda} \\ &\propto A^{-\beta} (AL)^{\lambda} = A^{\lambda-\beta} L^{\lambda}, \end{aligned} \tag{10}$$

since the capital stock grows with output. Capital now drives sustained growth as long as $\beta \leq \lambda$—and hyperbolic growth if the inequality is strict.

Bloom et al. (2020) estimate $\lambda/\beta \approx 1/3$. This parameter estimate suggests that automating R&D alone would not sustain exponential growth, let alone explosive growth, given a constant population.

Acceleration despite the need for population growth. Given population growth, however, automating R&D raises the steady-state growth rate from $g_L \, \lambda/\beta$ to $g_L \, \lambda/(\beta - \lambda)$: i.e., if $\lambda/\beta = 1/3$, from $g_L/3$ to $g_L/2$.

Furthermore, unless the capital requirement for an automated researcher is very high, automating R&D permits a transitory period of rapid growth as capital pours into the research sector. Carlsmith (2020) estimates that the human brain performs approximately 10^{15} computations per second. AI software as capable and efficient as the brain would let the stock of Nvidia GPUs in early 2026 support about 10 million virtual researchers, based on estimates by Emberson and Owen (2025). This capacity is about equal to the current global researcher population (UNESCO Institute for Statistics, 2024) and has doubled around every ten months for the last six years (Emberson and Owen, 2025). Going forward, therefore, even if only a small fraction of computing capacity is used for research purposes, total research effort could rapidly expand. Given $\beta > \lambda$, a growing supply of hands to build these chips will be necessary to sustain growth in the very long run; but as spending on information processing equipment constitutes less than 2% of US GDP,[6] automating R&D could permit a large increase in research effort before this constraint binds.

Since machines capable of fully substituting for human labor in all production tasks are feasible in principle, R&D acceleration shortens the development horizon for such systems. It may be, therefore, that the large acceleration to technological development induced by automating R&D would be temporary if R&D remained automated in isolation, but will be permanent in practice by hastening the arrival of full automation more generally.

3.2 Explosive growth in a semi-endogenous model?

It is also possible that, if research is automated, $\beta < \lambda$ will obtain, so that technology-driven explosive growth will follow even if production always requires labor. This is for at least three reasons.

Larger vs. more numerous systems. While labor remains the primary input to R&D, each additional unit of resources must take the form of a discrete brain, able only very slowly to communicate with the world outside its skull. The returns to increasing R&D investment may diminish more slowly, or even increase, when larger investments may take the form of larger (not just more numerous) well-integrated systems (Sotala, 2012). That is, "λ" may prove higher in (10) than in (9).[7] By analogy, we would underestimate the scale effect of increasing population size if we imagined that people required their own farmsteads and could not agglomerate in labs or cities.

Similarly, Agrawal et al. (2019) observe that, in a Weitzman (1998)-style model in which technological progress is made by discovering fruitful combinations of existing

[6] In 2024, \$0.5T out of \$29T according to FRED's national accounts.

[7] Indeed, AI progress to date has benefited greatly not only from increases in the raw computation used to train models, but from the Raina et al. (2009) discovery that GPUs allow more of the relevant computation to be done in parallel, and from the ongoing development of ever larger chips that permit even more parallel computations with low latency (see e.g. Kundu et al. (2025)).

ideas, ideas may get harder to find due partly to the human brain's inability to consider large combinations of ideas at once. If so, as the authors show, automated R&D carried out by large and growing AI models may structurally face a lower value of "β".

Copies of Einstein. Ekerdt and Wu (2024) argue that, as the share of the population in research has grown, average researcher ability has declined. In the AI context, this selection effect can be avoided by running multiple copies of the most capable models. Absent selection, Ekerdt and Wu estimate $\lambda/\beta \approx 3/5$.

Returns to R&D diminish slowly in hardware and software. Bloom et al. find wide variation in the returns to R&D across industries. Moore's Law—the doubling of semiconductor efficiency roughly every 18 months—has been supported by a growing workforce in semiconductor R&D, but this workforce has grown only one fifth as quickly. If semiconductor R&D has not benefited from spillovers from other industries, this implies that the industry has exhibited $\lambda/\beta \approx 5$. Thus if R&D were automated and semiconductors were its sole input, growth could easily be explosive.

Furthermore, the number of effective artificial researchers increases not only in the quantity and efficiency of our hardware but also in the efficiency of our software. Eth and Davidson (2025), reviewing the literature across language models, computer vision models, and other domains, estimate that the raw computational requirements of many AI capabilities have tended to roughly halve each year. As the number of AI researchers has not grown quite as quickly, this implies an estimate of $\lambda/\beta > 1$ in the domain of software efficiency as well. At least before accounting for spillovers and other inputs, this suggests that the number of effective virtual researchers could explode, and drive explosive growth more generally, even absent hardware improvements.

Productivity improvements, however, may require not only explicit R&D effort but also a process of "learning by doing". That is, the information generated by real-world economic activity may be a gross complement to R&D labor and capital, so that the useful research outputs of even a large number of virtual Einsteins would eventually be bottlenecked without growth in the information generated by the economy at large.[8]

[8] Observe that the "scaling laws" describing performance of AI models across a very wide range of domains (see Villalobos (2023) for a review) imply that compute and data are gross complements in AI training.

Arguably, the capabilities of generative AI models have advanced rapidly in recent years due to the discovery of ways to harness the vast pool of data that constitutes the Internet, and further significant advances will require gathering, codifying, and continuously updating training data capturing much of what people learn in the course their work. In this case, the bottlenecking input to technological progress, once R&D is automated, may ultimately be proportional to output (as in (10)) rather than effective computation.

3.3 Automating R&D in fully endogenous models

The rate of technological development has remained roughly constant in frontier economies for more than a century though the number of researchers has greatly grown. On a semi-endogenous account, this is because the number of researchers required to sustain a given growth rate rises as the technological frontier advances. That is, the primary bottleneck to faster technological development is a lack of research labor. As detailed above, it follows straightforwardly that if this bottleneck is relieved, with capital effectively increasing the supply of researchers, growth accelerates.

On an endogenous growth account, by contrast, a fixed population can sustain a constant growth rate. As a result, the endogenous growth literature since Young (1998) and Aghion and Howitt (1998, ch. 12.2) has proposed forces that constrain the growth rate despite a growing population. If these forces are strong enough, they prevent automating research from delivering explosive growth.

Extreme non-parallelizability. If technology growth takes roughly the form (9), constant g_A is incompatible with positive g_L if $\beta = 0$ and $\lambda > 0$. On the semi-endogenous interpretation, $\beta > 0$. The alternative endorsed by some endogenous accounts is effectively that $\beta = 0$ and $\lambda = 0$. It is difficult to disentangle the variables empirically: Bloom et al. (2020) estimate λ/β, for example, but cannot identify whether both variables are very low or very high.

Concretely, suppose that

$$g_A \propto f(L), \tag{11}$$

where $f(\cdot)$ is increasing but $\lim_{L\to\infty} f(L) = \bar{f} < \infty$. Then the growth rate is constant with a constant population but can rise only to $\bar{f}$ if the population grows,[9] perhaps due to the need for serial experiments. Automating R&D allows the growth rate to approach $\bar{f}$ more quickly, or allows the growth rate to rise somewhat when a fixed population would otherwise have rendered it constant, as $K \propto AL$ replaces L in (11). Nevertheless, the long-run growth rate is constant.

As discussed in the previous section, automated research systems may face a lower penalty to parallelizing R&D than populations of human researchers, and the penalty may fall over time with advances in hardware. That is, in the present context, we might imagine that automating research will deliver roughly

$$g_A \propto h(A)f(K),$$

so that the maximum growth rate achievable by scaling research investment at a given time equals $h(A)\bar{f}$ and rises with A. If so, whether automating R&D will deliver explosive growth depends on the shape of $h(\cdot)$, a question about which today we can only speculate.

Creative destruction. In other endogenous growth models, explosive growth is technologically feasible with a growing population,[10] but R&D efforts of the kind that could deliver it are limited by "Schumpeterian" dynamics. An innovator earns profits from developing a more efficient way to produce a given product for only as long as this process remains on the technological frontier. An accelerating stream of investments in productivity improvements cannot be sustained in equilibrium, as it would shrink the expected duration of each innovator's monopoly to zero.

Automating R&D may therefore fail to raise the growth rate. Aghion et al. (2019, Appendix) discuss this point via a model in which partially automating R&D actually causes the growth rate to fall,[11] but the point can also be seen directly from the fact that many Schumpeterian endogeneous growth models are in fact "lab equipment" models already[12]: capital is modeled as the only input to R&D. This unrealistic as-

[9] This is, very roughly, the model proposed by Young (1998) and Lashkari (2023): see Trammell (2025).

[10] Or even a fixed population: see Trammell (2025).

[11] In their model, technological development proceeds in two stages. AI accelerates only one but thereby shortens the monopoly earned by successfully investing in the other. The resulting disincentive yields slower growth on balance.

[12] Including seminally Aghion and Howitt (1998, ch. 12.2); see Trammell (2025) for an up-to-date

sumption can be made for simplicity, without predicting explosive growth, because the central bottleneck to growth is not a lack of resources available to devote to research but an arbitrage equation ensuring that innovations arrive slowly enough to justify developing them.

4 Full automation

Naturally, AI will have the most radical effects on growth if it renders capital a gross substitute for labor in both production and R&D.

If the long-run rate of technological development is constrained to be constant, e.g. for the reasons discussed in Section 3.3, we are in the world of Section 2 with constant $g_A > 0$. With capital serving as a gross substitute for labor in production, so that $Y \approx AK$, growth is double-exponential: the growth rate itself grows at rate g_A. We will assume the semi-endogenous model for the remainder of this section.

The Cobb-Douglas edge case. In a semi-endogenous model, as we have seen, the growth rate is constrained primarily by R&D inputs, so under full automation it can rise arbitrarily rapidly alongside the capital stock. In fact, the feedback loop from technology to output to capital to technology is strong enough that even a Cobb-Douglas production function, in combination with a Cobb-Douglas R&D function, yields hyperbolic growth if the exponents on capital are sufficiently high (Aghion et al., 2019).

Fixing population, we have

$$\begin{aligned} Y &\propto AK^{\alpha}, \\ g_A &\propto A^{-\beta} K^{\xi}. \end{aligned} \tag{12}$$

To intuit the parameter condition necessary to trigger hyperbolic growth,[13] observe that if the capital stock is proportional to output,

$$K \propto AK^{\alpha} \quad \Longrightarrow \quad K \propto A^{\frac{1}{1-\alpha}}. \tag{13}$$

review.

[13]For a proof, see Trammell (2023).

Substituting (13) into (12) yields

$$g_A \propto A^{\frac{\xi}{1-\alpha}-\beta}. \tag{14}$$

Technology, capital, and output thus grow hyperbolically if

$$\frac{\xi}{1-\alpha} > \beta \tag{15}$$

and exponentially only if the terms are equal. If the inequality is reversed, the feedback loop at hand delivers only power-functional growth, so that exponential growth requires growth in population.[14]

Gross substitution. Gross substitution in R&D is equivalent to the case of $\xi = \lambda$ (or more, as discussed in Section 3.1). Whether it delivers hyperbolic growth remains ambiguous, but condition (15) with $\xi = \lambda$—i.e. $\frac{\lambda}{1-\alpha} > \beta$—is of course a weakening of the $\lambda > \beta$ required given gross complementarity in production.

Gross substitution in production is equivalent here to the case of $\alpha = 1$. It delivers hyperbolic growth for any value of ξ or β.

A fixed factor. Under a Cobb-Douglas natural resource constraint, as discussed in Section 2.4, the capital share can rise only to $\bar{\alpha} < 1$, where $1-\bar{\alpha}$ is the natural resource share. Because, again, the natural resource share is not even 5%, this constraint will not prevent hyperbolic growth in the event of fully automated production given any non-negligible capital share in R&D.

With "L" in place of "K", this model also describes a world in which labor remains necessary for production and R&D but population grows proportionally to output. Kremer (1993) uses just such a model to explain the superexponential—indeed, arguably roughly hyperbolic (Roodman, 2020)—population growth that obtained across one million years of the Malthusian past.

Growth with fixed resources and fixed technology presumably cannot continue forever, even under full automation. Natural resources may thus eventually come to bottleneck output more strongly. That is, fixing resources, sustained growth may re-

[14] Davidson et al. (2025) offer a generalization of condition (15) in a model in which both hardware and software improvements can increase the number of effective "virtual researchers", as motivated by the discussion in Section 3.2.

quire the development of resource-augmenting technology "Z". Suppose that, absent automation,

$$Y = F(K, AL, Z\bar{W}),$$

where $F(\cdot)$ is CRS and K, AL, and $Z\bar{W}$ are gross complements.[15] In this case, automating production yields approximately

$$Y = F(AK, AK, Z\bar{W}) \approx Z\bar{W}$$

in the long run, so that output grows with technology as in Section 3. Here, if $\lambda < \beta$, growth eventually slows—even, in this case, with both production and research fully automated.[16]

Because growth does not begin to decelerate until the resource share has risen to the point of reversing inequality (15), however, automating both sectors in the near future would still yield a substantial period of accelerating growth.

5 Summary

Once production is sufficiently automated, output growth is likely to rise, and this is likely to continue for the foreseeable future. Automating R&D may not deliver explosive growth on its own (except by speeding the development of ways to automate production); whether it does so depends largely on whether automated research proves highly parallelizable and on whether it requires inputs that scale with economic activity rather than with effective computation. Even if fully automating R&D in isolation would not dramatically speed up growth, however, more R&D automation could greatly enhance the acceleration to growth induced by automating production.

Automating R&D in isolation would probably preserve a positive labor share, as would developing machines that could self-replicate but could not well substitute for

[15] We must also suppose that historically $g_Z = g_A + g_L$, maintaining constant factor shares. The elasticity of substitution between energy (or natural resources in general) and other resources is often estimated to be much less than one over short time horizons but to rise significantly with the horizon (Koetse et al., 2008; van der Werf, 2008), suggesting that the resource share has indeed remained low only due to the ongoing development of resource-augmenting technology.

[16] If there is a finite pool of technologies as well, β must eventually rise without bound, so that $g_A = 0$ at high A.

labor in producing consumption goods. Otherwise, the burst in output made feasible by a qualitative advance in automation will almost all accrue to capital, driving the labor share to zero. It will probably be very inefficient (and so very difficult) to steer technology in a labor-augmenting direction permanently.

Whether the wage bill will rise or fall in the limit of automation depends on subtleties about the long-run direction of technology and tastes, and the degree of returns to scale, about which little can be said today.

The long-run scenarios may be summarized as follows. In each case we assume that capital grows more quickly than population and that, given gross complementarity between labor and capital in production, technological development shifts from being labor- to capital-augmenting.

Scenario[17]	Growth	Labor share	Wages
Production automated			
2 factors; GS	D	$\to 0$	D
2 factors; PS	D	$\to 0$	S or $\to 0$
CD for consump., GS/PS for capital	D	$\sim$	D
3 factors (CD resource constr.); GS	+	$\to 0$	$\sim$
3 factors (CD resource constr.); PS	+	$\to 0$	$\to 0$
R&D automated, or both automated given a GC resource constraint[18]			
Semi-endog., weak feedback	+		
Semi-endog., intermed. feedback	D		
Semi-endog., strong feedback	H		
Hard parallelizability constraint	+		
Schumpeterian	$\sim$		
Both automated			
Semi-endog., very weak feedback and CD resource constr.	+ or D		
Semi-endog., otherwise	H		

6 Conclusion

From AI to automation. Interest in the question of what will follow from near-complete automation has grown recently because the rapid pace of recent AI progress

[17] "CD", "GS", and "PS" stand for Cobb-Douglas, gross but not perfect substitution, and perfect substitution. "+", "D", and "H" refer to cases in which AI produces increases to a variable's long-run growth rate, double-exponential growth, and hyperbolic growth. "S" refers to cases in which AI causes the variable to stagnate above zero. "$\sim$" refers to cases in which AI has an ambiguous effect on the variable.

[18] "Weak", "intermediate", and "strong" research feedback refer to the $\lambda/(\lambda-\beta) < 1$, $= 1$, and > 1 cases. We here focus only on growth, on the view that wages and the labor share will still mainly be governed by production.

suggests that the automation of production and/or R&D may be on the horizon.

AI to date has not had the consequences that might be expected to precede such a drastic economic transformation. As Nordhaus (2021) documents, capital and labor remain (at least weakly) gross complements, the growth rate is not rising, the IT share of capital is not skyrocketing, and any increases in the capital share are modest and easily attributable to other forces.[19] Humlum and Vestergaard (2025) find that LLMs have been adopted rapidly—over a billion users within 2.5 years of their introduction—but so far seem to have delivered only modest productivity improvements.

Nevertheless, the capabilities of AI systems have advanced at extraordinary rates. Recent analysis by METR documents that the maximum task length AI systems can handle has been doubling approximately every seven months (Kwa et al., 2025). Amodei (2025) projects that coding could be fully automated by 2026. Should capabilities continue advancing at comparable rates, the comprehensive automation of computer-based work may follow shortly thereafter—potentially providing sufficient technological foundation to devise methods for automating remaining economic activities.

An alternative analytical lens for understanding the trajectory toward full automation focuses on improvements in *sample efficiency*. Deep learning has enabled AI systems to match or exceed human performance in domains such as coding, language processing, and strategic gaming, where either substantial public data exists or synthetic data generation is feasible. Nevertheless, the human brain remains orders of magnitude more sample-efficient than current machine learning algorithms, requiring far less data to acquire comparable non-trivial skills. Eth and Davidson (2025) review algorithmic progress across language models, computer vision, and related domains and finds that data requirements for training AI capabilities have approximately halved annually in recent years. If the current human-machine sample efficiency gap spans roughly four orders of magnitude, these rates suggest AI systems could achieve human-level learning efficiency within one to two decades.

AI is starting to help contribute to R&D, and the high capital share in AI R&D specifically (Besiroglu et al., 2024) suggests that capital is, in some sense, contributing significantly to its own improvement. As AI R&D is automated ever more fully,

[19] Work in progress by one of the authors confirms that updating Nordhaus's analysis to 2025 does not change these conclusions, and that, on examining the supply chains for robotics and computer equipment, there is similarly little evidence for a capital-goods-driven growth explosion of the kind described at the end of Section 2.1.

AI capabilities could advance ever more rapidly. Erdil et al. (2025) and Kokotajlo et al. (2025) calibrate models of this feedback mechanism—albeit with some speculative elements—and predict near-complete R&D automation within several years, precipitating rapid economic growth thereafter. Informed partly by such models, the median forecast on the prediction platform Metaculus places AGI in 2033. A notable characteristic of the R&D-initiated "singularity" hypothesis is its prediction of rapid, nearly discontinuous transformation in economic conditions. Consequently, the absence of observable effects in current economic data cannot definitively refute the possibility of imminent dramatic change: the theory itself predicts minimal advance warning in conventional economic indicators.

Preparing for economic transformation. While Piketty and Zucman (2014)'s controversial claim that capital and labor are already gross substitutes conflicts with the consensus view—supported by direct evidence including Gechert et al.'s (2022) comprehensive review and theoretical arguments that gross substitutability should manifest in accelerating growth and declining labor shares—AI may render the claim true in the future, and vindicate Piketty (2014)'s concern about twenty-first century capital concentration. This prospect necessitates considering policy responses that could mitigate extreme distributional consequences (Korinek and Lockwood, 2025).

Preparing for dangerous AI. Concerns about existential risk from advanced AI systems, articulated in early theoretical work (Bostrom, 2014) and recently amplified through public statements like Center for AI Safety (2023), represent perhaps the most extreme form of economic transformation: complete cessation of economic activity through catastrophic failure. Both proponents and skeptics of such risks employ AI's capacity for generating radical economic growth as a proxy measure for its potential to inflict systemic damage (see Davidson (2021) and Narayanan and Kapoor (2025) for contrasting perspectives). The scenarios analyzed throughout this paper carry differential risk profiles: a "software-only singularity" driven purely by R&D automation presents unique challenges through its potential for rapid, unobservable acceleration beneath conventional economic indicators, while scenarios involving automated physical production imply AI systems capable of tangibly affecting the material world without requiring human intermediation, raising distinct safety concerns about autonomous physical systems.

The economic valuation of existential risk mitigation—quantifying optimal consumption sacrifices to reduce catastrophic probabilities—has generated substantial recent analysis (Jones, 2024, 2025; Shulman and Thornley, 2025), though excessive precaution may paradoxically increase long-term risks through delayed beneficial development (Trammell and Aschenbrenner, 2025). Policy responses center on established approaches: advancing technical AI safety research to ensure system alignment and robustness, and developing governance frameworks capable of managing rapid capability transitions while maintaining incentives for innovation.

The imperative of preparation. The economic transformation analyzed throughout this paper may represent humanity's most consequential transition since the Industrial Revolution, or perhaps since the Neolithic. While considerable uncertainty remains regarding both the timing and precise mechanisms, the potential magnitude of change demands serious preparatory efforts commensurate with the stakes involved. The scenarios we have examined, from explosive growth driven by self-replicating or self-improving capital to radical redistributions of economic returns, are not merely theoretical curiosities but plausible trajectories that could materialize within decades or even years. The computational and algorithmic progress documented here suggests we may have limited time to design institutions capable of managing unprecedented growth rates, distributional upheaval, and the accompanying risks.

Yet this moment also presents extraordinary opportunity: if navigated successfully, the automation of cognitive and physical labor could usher in an era of prosperity that liberates humanity from millennia of drudgery (Korinek and Juelfs, 2024). The frameworks we have reviewed provide tools for reasoning about these transitions. Whether the coming transformation proves to be humanity's greatest triumph or its final chapter may depend critically on the decisions made in the narrow window before these technologies reach their transformative thresholds. Given both the promise and peril inherent in transformative AI, even modest investments in understanding and shaping this transition represent perhaps the highest-return activities our profession can engage in.

References

Acemoglu, Daron, "Labor- and Capital-Augmenting Technical Change," *Journal of the European Economic Association*, 2003, *1* (1), 1–37.

— , *Introduction to Modern Economic Growth*, Princeton University Press, 2009.

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

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

— **and** — , "The Wrong Kind of AI? Artificial Intelligence and the Future of Labour Demand," *Cambridge Journal of Regions, Economy and Society*, 2020, *13* (1), 25–35.

Aghion, Philippe and Peter Howitt, *Endogenous Growth Theory*, Cambridge, MA: MIT Press, 1998.

— **and Simon Bunel**, "AI and Growth: Where Do We Stand?," 2024. Working paper.

— **, 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*, National Bureau of Economic Research, 2019, pp. 237–282.

Agrawal, Ajay, John McHale, and Alex Oettl, "Finding Needles in Haystacks: Artificial Intelligence and Recombinant Growth," in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds., *The Economics of Artificial Intelligence: An Agenda*, National Bureau of Economic Research, 2019, pp. 149–174.

American Horse Council, "2023 Economic Impact of the Horse Industry in the USA," 2023. Accessed Aug. 15, 2025.

Amodei, Dario, "CEO Speaker Series with Dario Amodei of Anthropic," Council on Foreign Relations, CEO Speaker Series March 2025. Uncorrected transcript. Quote: "I think we'll be there in three to six months, where AI is writing 90% of the code. And then in twelve months, we may be in a world where AI is writing essentially all of the code.".

Autor, David, Caroline Chin, Anna Maria Salomons, and Bryan Seegmiller, "New Frontiers: The Origins and Content of New Work, 1940–2018," *Quarterly Journal of Economics*, 2024, *139* (3), 1399–1455. NBER Working Paper No. 30389, August 2022.

Autor, David H. and David Dorn, "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review*, 2013, *103* (5), 1553–97.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, "Artificial Intelligence, Firm Growth, and Product Innovation," *Journal of Financial Economics*, 2024, *151* (103745).

Benzell, Seth G. and Victor Yifan Ye, "Simulating the Global Effect of Transformative AI: Growth, Welfare, Economic Power, and Policy Responses," 2024. Working paper.

Berg, Andrew, Edward F Buffie, and Luis-Felipe Zanna, "Should We Fear the Robot Revolution? (The Correct Answer Is Yes)," *Journal of Monetary Economics*, 2018, *97*, 117–148.

Besiroglu, Tamay, Nicholas Emery-Xu, and Neil Thompson, "Economic impacts of AI-augmented R&D," *Research Policy*, 2024, *53* (7).

Bloom, Nicholas, Charles I Jones, John Van Reenen, and Michael Webb, "Are Ideas Getting Harder to Find?," *American Economic Review*, 2020, *110* (4), 1104–1144.

Bostrom, Nick, *Superintelligence*, Oxford University Press, 2014.

Brynjolfsson, Erik, "The Turing Trap: The Promise and Peril of Human-Like Artificial Intelligence," *Dædalus*, 2022, *151* (2), 267–84.

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

Carlsmith, Joseph, "How Much Computational Power Does It Take to Match the Human Brain?," *Open Philanthropy Report*, 2020.

Center for AI Safety, "Statement on AI Risk," 2023.

Chen, Zenan and Jason Chan, "Large Language Model in Creative Work: The Role of Collaboration Modality and User Expertise," 2025. Working paper.

Choi, Johnathan H. and Daniel Schwarcz, "AI Assistance in Legal Analysis: an Empirical Study," *Journal of Legal Education*, 2025, *73* (2).

Davidson, Tom, "Could Advanced AI Drive Explosive Economic Growth?," Technical Report, Open Philanthropy 2021.

_ **, Basil Halperin, Tom Houlden, and Anton Korinek**, "Is Automating AI Research Enough for a Growth Explosion?," 2025.

Dell'Acqua, Fabrizio, Edward McFowland III, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R. Lakhani, "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality," 2023. Working paper.

Eden, Maya and Kevin Kuruc, "The Long-run Relationship Between Per Capita Incomes and Population Size," 2024. GPI Working Paper Series No. 29-2024.

Ekerdt, Lorenz K.F. and Kai-Jie Wu, "Self-Selection and the Diminishing Returns of Research," 2024. Working paper.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, "GPTs are GPTs: Labor market impact potential of LLMs," *Science*, 2024, *384* (6702), 1306–8.

Emberson, Luke and David Owen, "The stock of computing power from NVIDIA chips is doubling every 10 months," Technical Report, Epoch 2025.

Erdil, Ege, Tamay Besiroglu Andrei Potlogea, Edu Roldan, Anson Ho, Jaime Sevilla, Matthew Barnett, Matej Vrzla, and Robert Sandler, "GATE: An Integrated Assessment Model for AI Automation," Technical Report, Epoch AI 2025.

Eth, Daniel and Tom Davidson, "Will AI R&D Automation Cause a Software Intelligence Explosion?," Technical Report, Forethought 2025.

Gechert, Sebastian, Tomas Havranek, Zuzana Irsova, and Dominika Kolcunová, "Measuring Capital-Labor Substitution: The Importance of Method Choices and Publication Bias," *Review of Economic Dynamics*, 2022, *45*, 55–82.

Gosen, Bradley S. Van, Philip L. Verplanck, Robert R. Seal II, Keith R. Long, and Joseph Gambogi, "Rare-earth elements," in Klaus J. Schulz, John DeYoung Jr., Robert R. Seal II, and Dwight Bradley, eds., *Critical mineral resources of the United States—Economic and environmental geology and prospects for future supply*, United States Geological Survey, 2017.

Habakkuk, H. John, *American and British Technology in the Nineteenth Century*, Cambridge, UK: Cambridge University Press, 1962.

Hanson, Robin, "Economic Growth Given Machine Intelligence," 2001. Working paper.

Hémous, David and Morten Olsen, "The Rise of the Machines: Automation, Horizontal Innovation, and Income Inequality," *American Economic Journal: Macroeconomics*, 2022, *14* (1), 179–223.

Humlum, Anders and Emilie Vestergaard, "Large Language Models, Small Labor Market Effects," 2025. NBER working paper 33777.

Jones, Benjamin F. and Xiaojie Liu, "A Framework for Economic Growth with Capital-Embodied Technical Change," *American Economic Review*, 2024, *114* (5), 1448–87.

Jones, Charles I, "R&D-based Models of Economic Growth," *Journal of Political Economy*, 1995, *103* (4), 759–784.

Jones, Charles I., "The Past and Future of Economic Growth: ASemi-Endogenous Perspective," *Annual Review of Economics*, 2022, *14*, 125–52.

__ , "The A.I. Dilemma: Growth versus Existential Risk," *American Economic Review: Insights*, December 2024, *6* (4), 575–590.

__ , "How Much Should We Spend to Reduce A.I.'s Existential Risk?," NBER Working Paper 33602, National Bureau of Economic Research March 2025.

Kaldor, Nicholas, "A Model of Economic Growth," *The Economic Journal*, 1957, *67* (268), 591–624.

Koetse, Mark J., Henri L.F. de Groot, and Raymond J.G.M. Florax, "Capital-energy Substitution and Shifts in Factor Demand: A Meta-analysis," *Energy Economics*, 2008, *30*, 2236–51.

Kokotajlo, Daniel, Scott Alexander, Thomas Larsen, Eli Lifland, and Romeo Dean, "AI 2027," Technical Report, AI Futures Project 2025.

Korinek, Anton, "Economic Policy Challenges for the Age of AI," in Pierre-Olivier Gourinchas, Maurice Obstfeld, and Petia Topalova, eds., *Rethinking Economic Policy: Steering Structural Change*, MIT Press, 2026, chapter 18. Forthcoming.

__ **and Donghyun Suh**, "Scenarios for the Transition to AGI," 2024. NBER working paper 32255.

__ **and Joseph E Stiglitz**, "Artificial Intelligence and Its Implications for Income Distribution and Unemployment," in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds., *The Economics of Artificial Intelligence: An Agenda*, National Bureau of Economic Research, 2019, pp. 349–390.

__ **and Joseph E. Stiglitz**, "Steering Technological Progress," *Conference paper, NBER Economics of AI Conference*, 2020.

— **and Lee Lockwood**, “"Public Finance in the Age of AI", journal = "Brookings CRM Working Paper",” 2025.

— **and Megan Juelfs**, “Preparing for the (Non-Existent?) Future of Work,” in Justin Bullock et al., ed., *Oxford Handbook of AI Governance*, Oxford University Press, 2024.

Kremer, Michael, “Population Growth and Technological Change: One Million BC to 1990,” *The Quarterly Journal of Economics*, 1993, *108* (3), 681–716.

Kundu, Yudhishthira, Manroop Kaur, Tripty Wig, Kriti Kumar, Pushpanjali Kumari, and Vivek Puri, “A Comparison of the Cerebras Wafer-Scale Integration Technology with Nvidia GPU-based Systems for Artificial Intelligence,” Technical Report 2025.

Kwa, Thomas, Ben West, Joel Becker, Amy Deng, Katharyn Garcia, Max Hasin, Sami Jawhar, Megan Kinniment, Nate Rush, Sydney von Arx, Ryan Bloom, Thomas Broadley, Haoxing Du, Brian Goodrich, Nikola Jurkovic, Luke Harold Miles, Seraphina Nix, Tao Lin, Neev Parikh, David Rein, Lucas Jun Koba Sato, Hjalmar Wijk, Daniel M. Ziegler, Elizabeth Barnes, and Lawrence Chan, “Measuring AI Ability to Complete Long Tasks,” *arXiv:2503.14499*, March 2025.

Lashkari, Danial, “Innovation Policy in Theories of Knowledge Diffusion and Selection,” 2023. Working paper.

Mookherjee, Dilip and Debraj Ray, “Growth, Automation, and the Long-run Share of Labor,” *Review of Economic Dynamics*, 2022, *46*, 1–22.

Narayanan, Arvind and Sayash Kapoor, “AI as Normal Technology,” Technical Report, Knight First Amendment Institute 2025.

Nordhaus, William D, “Are We Approaching an Economic Singularity? Information Technology and the Future of Economic Growth,” *American Economic Journal: Macroeconomics*, 2021, *13* (1), 299–332.

Noy, Shakked and Whitney Zhang, “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence,” *Science*, 2023, *381* (6654).

Otis, Nicholas G., Rowan Clarke, Solene Delecourt, David Holtz, and Rem Koning, “The Uneven Impact of Generative AI on Entrepreneurial Performance,” 2024. Working paper.

Our World in Data, “Global primary energy consumption by source,” 2025.

Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer, "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot," 2023. Technical report.

Piketty, Thomas, *Capital in the Twenty-First Century*, Cambridge, MA: Harvard University Press, 2014.

— **and Gabriel Zucman**, "Capital is Back: Wealth-Income Ratios in Rich Countries 1700–2010," *Quarterly Journal of Economics*, 2014, *129* (3), 1255–310.

Prettner, Klaus, "A Note on the Implications of Automation for Economic Growth and the Labor Share," *Macroeconomic Dynamics*, 2019, *23* (3), 1294–1301.

Raina, Rajat, Anand Madhavan, and Andrew Y. Ng, "Large-scale Deep Unsupervised Learning using Graphics Processors," in "Proceedings of the 26th Annual International Conference on Machine Learning" ICML '09 Association for Computing Machinery New York, NY, USA 2009, pp. 873–80.

Roldán-Monés, Antonio, "When GenAI increases inequality: evidence from a university debating competition," 2024. Working paper.

Romer, Paul, "Endogenous Technological Change," *Journal of Political Economy*, 1990, *98* (5), S71–S102.

Roodman, David, "On the Probability Distribution of Long-term Changes in the Growth Rate of the Global Economy: An Outside View," 2020. Open Philanthropy.

Sachs, Jeffrey D and Laurence J Kotlikoff, "Smart Machines and Long-term Misery," Technical Report 2012. National Bureau of Economic Research, Working Paper No. 18629.

Sachs, Jeffrey D., Seth G. Benzell, and Guillermo LaGarda, "Robots: Curse or Blessing? A Basic Framework," 2015. NBER working paper 21091.

Sandberg, Anders, "An Overview of Models of Technological Singularity," in Max More and Natasha Vita-More, eds., *The Transhumanist Reader: Classical and Contemporary Essays on the Science, Technology, and Philosophy of the Human Future*, Wiley Online Library, 2013, pp. 376–394.

Sasaki, Hiroaki, "Growth with Automation Capital and Declining Population," *Economics Letters*, 2023, *222* (110958).

Shulman, Carl and Elliott Thornley, "How Much Should Governments Pay to Prevent Catastrophes? Longtermism's Limited Role," in Hilary Greaves, Jacob Barrett, and David

Thorstad, eds., *Essays on Longtermism: Present Action for the Distant Future*, Oxford: Oxford University Press, 2025.

Sotala, Kaj, "Advantages of Artificial Intelligences, Uploads, and Digital Minds," *International Journal of Machine Consciousness*, 2012, *4* (1), 275–91.

Trammell, Philip, "Growth given Cobb-Douglas Automation," 2023. Note.

—, "The Extreme Inefficiency of Expanding Variety in Endogenous Growth Theory," 2025. Working paper.

— **and Leopold Aschenbrenner**, "Existential Risk and Growth," 2025. Working paper.

UNESCO Institute for Statistics, "The Sustainable Development Goals Extended Report 2024, 9: Industry, Innovation and Infrastructure," Technical Report 2024.

Uzawa, Hirofumi, "Neutral Inventions and the Stability of Growth Equilibrium," *The Review of Economic Studies*, 1961, *28* (2), 117–124.

van der Werf, Edwin, "Production Functions for Climate Policy Modeling: An Empirical Analysis," *Energy Economics*, 2008, *30*, 2964–79.

Villalobos, Pablo, "Scaling Laws Literature Review," 2023. Epoch AI blog post.

Weitzman, Martin L, "Recombinant Growth," *The Quarterly Journal of Economics*, 1998, *113* (2), 331–360.

World Energy Council, "World Energy Resources: 2013 Survey," 2013. Chapter 8: Solar Energy.

Young, Alwyn, "Growth without Scale Effects," *Journal of Political Economy*, 1998, *106* (1), 41–63.

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