Past Automation and Future A.I.: How Weak Links Tame the Growth Explosion

Charles I. Jones
Stanford GSB and NBER

Christopher Tonetti*
Stanford GSB and NBER

March 26, 2026 — Version 0.4

Preliminary — the numbers may change

Abstract

How much of past economic growth is due to automation, and what does this imply about the effects of A.I. in the coming decades? We perform growth accounting using a task-based model for key sectors in the U.S. economy. Historically, TFP growth is driven primarily by improvements in capital productivity. At the task level, capital productivity has grown at least 3 percentage points per year faster than labor productivity. The main benefit of automation is therefore that we use rapidly-improving machines instead of slowly-improving humans on an increasing share of tasks. Looking to the future, we develop an endogenous growth model in which the production of both goods and ideas is endogenously automated and calibrate the model based on our historical accounting. Automation leads economic growth to accelerate, but the acceleration is remarkably slow because of the prominence of "weak links," i.e. an elasticity of substitution among tasks substantially less than one. Even when most tasks are automated by rapidly-improving capital, output is constrained by the tasks performed by slowly-improving labor.

*We are grateful to Sebastian Di Tella, Pete Klenow, Pascual Restrepo, Michael Peters, Phil Trammell, and seminar participants at the CEPR Macro/Growth meeting, LBS, LSE, the Markus Academy, METR, National University of Singapore, Pacific Macro, SED, the Stanford-Berkeley Macro Conference, and Yale for helpful comments.

1. Introduction

Artificial intelligence is the latest form of automation, a process that has been ongoing for centuries. Farmers once threshed grain by hand; now a single combine harvester replaces dozens of workers. Elevator operators, typists, and travel agents were once ubiquitous; today, software and simple robots handle most of these tasks. In automobile plants, spot-welding and spray-painting have moved from human workers to industrial robots. Frontier LLMs are increasingly able to write computer code to replace some of the tasks that software engineers perform. Specialized A.I. models can even automate parts of the research process — think of AlphaFold solving the protein-folding problem or an A.I. model that assists researchers in making new discoveries (Bubeck et al., 2025).

How much of past economic growth is due to automation? Theoretical models have advanced our conceptual understanding of the automation process (Zeira, 1998; Acemoglu and Restrepo, 2018). There has, however, been much less progress on measuring its empirical contribution to economic growth. The first half of this paper fills that gap for the aggregate U.S. economy over the past 70 years and for select industries over the past 40 years. In the second half, we build a model in which both goods and idea production are endogenously automated over time. We calibrate the model based on our historical evidence and simulate the future to shed light on the possible consequences of continued automation through artificial intelligence.

Our framework is based on a canonical task model, but several ingredients are worth highlighting. First, we set the elasticity of substitution across tasks less than one: all tasks are essential, and even infinite output of some tasks leaves total production finite. Intuitively, our model features "weak links" and total output is constrained by the scarcest tasks. For example, consider $Y = \left(\sum_i Y_i^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}$ with $\sigma < 1$ so that the exponent $\frac{\sigma-1}{\sigma}$ is negative. If $Y_i = \infty$ for all tasks other than task j, then $Y = Y_j$. If the other tasks are merely finite rather than infinite, aggregate output is lower still. Overall output is therefore less than the quantity of the scarcest task. Production is constrained by the weakest links.

Second, our model features three types of productivity. The production of task i is $Y_{it} = \psi_{kit} K_{it} + \psi_{\ell it} L_{it}$. Each task can be produced with capital or labor as perfect sub-

stitutes. Each factor has its own productivity term that can change in a heterogeneous way over time, and this is true for each task. Automation is the process of switching task production from using labor to using capital, which occurs when the productivity of capital rises by enough relative to the productivity of labor. In addition, the production function is multiplied by an overall productivity index Z_t that captures other sources of TFP, for example due to quality improvements, new varieties, or changes in misallocation. Thus, the model features a rich structure of heterogeneity and multiple sources of productivity growth that could be driving TFP.

Growth accounting in this framework begins with readily available production account data from the BEA and the BLS for various U.S. industries. To identify the different sources of growth, we need to make one key assumption and one heroic measurement. The key assumption is this: while we place no restriction on the level or growth of any of the productivity variables, we assume that, on the margin, the tasks that get automated at a given point in time are those that have higher than average labor costs (or, equivalently, lower than average labor productivity). On average, the tasks at which labor is expensive are the ones that get automated. Our heroic measurement is that we measure the set of tasks that are automated in a given year via queries to ChatGPT's Deep Research. Both of these aspects of identification will be discussed in detail.

With this accounting framework, we derive several results:

1. When the automation process is continuous, firms switch from using labor to using capital to produce a task at exactly the point where the costs are equal. This means that the switching process itself generates no productivity growth.

2. Instead, the key gain from automation is that it allows production of a task to shift from slowly-improving human labor to rapidly-improving machines.

3. Our historical analysis suggests that the sum of "other" TFP growth and the average rate at which people are getting more productive, $\hat{Z}_t + \hat{\psi}_{\ell t}$, is small — for example equal to -0.1% per year for the private business sector since 1950. In contrast, the average excess rate at which machines are getting better is large. Across sectors, this gap, $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$, is at least 3 percentage points per year.

4. Finally, we calculate how much TFP growth would have been lost if the set of tasks that are automated had been "frozen" at some point in the distant past, but

capital, labor, and other productivity growth occurred at their historical rates. For the private business sector, if we freeze the set of automated tasks in 1950, 87% of TFP growth between 1950 and 2023 would have been eliminated. Using rapidly-improving machines on a fixed set of tasks is limited by weak links; most growth comes from switching to rapidly-improving machines on more tasks.

The final part of the paper augments our accounting framework to endogenize the automation process by incorporating the production of new ideas that raise ψ_{kit} and $\psi_{\ell it}$. This idea production function itself benefits from the automation of tasks. We calibrate the model to our historical evidence and simulate the model into the future to consider the possible consequences of continued automation — including via A.I. — for economic growth. While the results of this simulation are inherently speculative, they are grounded in our measurement of historical automation.

We consider two calibrations: "A.I. as a continuation of historical patterns" (based on the private business sector) and "A.I. as a break with the past" (based on the computer sector, where automation has been historically fast). Simulating the endogenous automation model forward in time we find:

1. Despite the stability of past economic growth, future growth accelerates for at least the next 50 years as the automation process endogenously speeds up. There is a flywheel effect: automation leads to more ideas which in turn leads to more automation.

2. When A.I. is a continuation of broad historical patterns, growth rates reach only 2.5% by 2075. But in our other calibration in which A.I. effectively makes the entire economy grow like the computer industry, the acceleration is substantially faster; growth explodes with incomes becoming infinite in finite time. Nevertheless, this explosion is notably slow: infinite income does not occur until around 2060 even in this extreme calibration. Both cases feature the flywheel effect that accelerates growth. However, it is only when the last weak links are automated away that the explosion occurs.

How to read this paper. This paper is long. Sections 2 – 4 present the theory and historical accounting to make the key point that automation contributes the majority

of past economic growth by letting us switch from slowly-improving labor to rapidly-improving capital on an increasing share of tasks. Spend half your time on these sections, and then spend the remainder of your time on Section 5. That section endogenizes growth and automation — including of the idea production function — and simulates the future consequences of A.I. for economic growth.

Related literature

This paper contributes to the literature on task-based models of economic growth begun by Zeira (1998), Acemoglu and Restrepo (2018, 2020, 2022), and Hemous and Olsen (2022) and to the research agenda on the economic impacts of A.I. outlined by Agrawal, Gans, and Goldfarb (2019) and Brynjolfsson, Korinek, and Agrawal (2025).

Our paper builds on Aghion, Jones, and Jones (2019) and B. Jones and X. Liu (2024). Aghion, Jones, and Jones (2019) use the task approach to study A.I. as automation. That paper presents a model in which the productivity of capital and labor in performing tasks is constant, and automation follows an exogenous law of motion. The paper emphasizes that bottlenecks may constrain the effects of automation on growth but notes that explosive growth is possible if A.I. fully automates both goods and idea production.

B. Jones and X. Liu (2024) incorporate heterogeneous productivity improvements in capital into the Aghion, Jones, and Jones (2019) framework. They showed that a balanced growth path could emerge even when automation is far from complete because automation raises the capital share while "better machines" lower the capital share. They go on to embed this setup in a fully endogenous growth model in which automation is an endogenous outcome of innovation. Farboodi, Koh, and Xia (2025) build on this work to study an endogenous automation process driven by data.

Trammell and Korinek (2020), Davidson (2021), Erdil and Besiroglu (2023), Aschenbrenner (2024), Korinek and Suh (2024), Davidson, Halperin, Houlden, and Korinek (2025), and Epoch AI (2025) all highlight the possibility of explosive economic growth that results from A.I. automating goods and idea production. B. Jones (2025) suggests that bottlenecks may constrain the growth impact of A.I. even when automating research and development.

All of these papers discussed so far highlight theoretical possibilities. Our paper is most clearly distinguished in using theory combined with industry-level data to mea-

sure automation and to quantify its consequences, both historically and in the future.

Young (2025) estimates a nested CES production function with capital, labor, and intermediates. He finds an elasticity of substitution between capital and labor of around 0.4–0.5. Young (2025) then finds intriguing evidence that capital-augmenting technical change is negative and suggests that a task-based model of technical change could drive this empirical finding. We confirm this result, but in value added terms. Our main contribution starts from this interesting fact and performs structural growth accounting to understand the nature of automation.

Although not the focus of their paper, B. Jones and X. Liu (2024) provide a time series for the fraction of tasks that have been automated and for average task-specific capital productivity; for manufacturing, they find that this latter series is roughly stationary and shows little growth since 1960. They back these out from industry-level data under the assumption that these are the only sources of productivity growth. Building on B. Jones and X. Liu (2024), Caunedo and Keller (2024) quantify the role of capital-embodied technical change for structural transformation. B. Jones and X. Liu (2024) allow capital productivity to vary across tasks but treat labor productivity as homogeneous and constant. Caunedo and Keller (2024) allow labor productivity to vary across tasks but treat capital productivity as homogeneous. This allows Cuanedo and Keller to measure improvements to capital using the relative price of investment. Their main finding is that CETC is the main driver of the reallocation of labor out of agriculture and accounts for one third of the reallocation of labor into services.

Instead, we seek to answer the question how much of past economic growth was due to automation while allowing for a rich set of sources of productivity growth. We allow both capital productivity and labor productivity to vary across tasks and over time arbitrarily. In addition, we allow for factor neutral productivity improvements. While our model is richer, identification requires more data and some alternative assumptions about patterns of automation. The payoff is that we provide a detailed accounting of the sources of TFP growth.

In terms of other papers that attempt to quantify the growth impacts of automation and A.I., Acemoglu (2024) suggests that the macroeconomic impacts of A.I. may be very modest in the next decade, raising TFP growth by less than 0.1pp per year. Aghion and Bunel (2024) respond by questioning some of the empirical choices made by Acemoglu

and calculate a larger impact over the next decade, raising TFP growth by 0.7pp per year.

2. Framework

Consider the following economic environment, which we typically think of as describing a sector like agriculture or motor vehicles:

$$Y_t = Z_t \left(\int_0^1 \alpha_i Y_{it}^{\frac{\sigma-1}{\sigma}} \, di \right)^{\frac{\sigma}{\sigma-1}} \qquad \text{where } \sigma < 1 \tag{1}$$

$$Y_{it} = \tilde{\psi}_{kit} K_{it} + \tilde{\psi}_{\ell it} L_{it} \tag{2}$$

$$K_t = \int_0^1 K_{it} \, di \tag{3}$$

$$L_t = \int_0^1 L_{it} \, di \tag{4}$$

where all parameters are positive.

A unit measure of complementary tasks are used to produce output. The heterogeneous share parameters α_i capture the fact that some tasks are more important than others. One unit of capital can produce $\tilde{\psi}_{kit}$ units of task i, while one unit of labor can produce $\tilde{\psi}_{\ell it}$ units of the task. We define $\psi_{kit} = \alpha_i^{\frac{\sigma}{\sigma-1}} \tilde{\psi}_{kit}$ and $\psi_{\ell it} = \alpha_i^{\frac{\sigma}{\sigma-1}} \tilde{\psi}_{\ell it}$.

Our setup therefore permits three different types of productivity improvements: higher ψ_{kit}, higher $\psi_{\ell it}$, and higher Z_t. We refer to Z_t as "other productivity." It can capture quality improvements in the sector, but it is also possible that new varieties or increased misallocation could impact Z_t.

2.1 Discussion of the Economic Environment

Complementarity and substitution. We set $\sigma < 1$ so that tasks are complements in production. This restriction is important and we discuss supporting evidence at the end of this section. Notice that the model features both complementarity and substitution. The complementarity arises from $\sigma < 1$ while the substitution arises from task production $Y_{it} = \tilde{\psi}_{kit} K_{it} + \tilde{\psi}_{\ell it} L_{it}$. The interplay between complementarity and

substitution is what allows a simple framework to give rise to a rich set of outcomes. It contrasts with the either/or aspect of the more traditional CES models such as $Y = F(BK, AL)$ which permits only complementarity *or* substitution rather than allowing both.[1]

Weak links. With $\sigma < 1$, tasks are "weak links" in the sense of Kremer (1993) and Jones (2011). Every task is essential to production and having infinite output of any task or even any measure of tasks below 100% still only leads to finite production. Total output can be no larger than the output of the weakest link — the task with the lowest output.[2] Aghion, Jones, and Jones (2019) referred to this feature as "bottlenecks," but we find the "weak links" interpretation to be more appropriate. Many of the most important conceptual insights of the paper are a direct result of the weak links production structure.

New tasks? Our model features a fixed measure of tasks, but we now have tasks such as "repair the computer" or "enter data into a spreadsheet" that did not always exist. In a world of substitutes, it is easy to see how adding new tasks could increase output. Indeed, that is essentially the mechanism underlying the Romer (1990) growth model. However, in our world of complements, adding new tasks could easily reduce output—production involves weak links rather than love-of-variety. Our approach in this paper to incorporating new methods of production is to add "new procedures" to our current setup. With a fixed unit measure of tasks, each task must be something that has always been done. In agriculture, this might be "till the soil" or "plant the seed." Over time, we invent new procedures for performing these tasks. For example, in the distant past we tilled the soil with manual labor, then with an ox and a plow, and now with a fancy GPS-enabled tractor.

Allow each task to be produced by a bunch of different procedures: $Y_{it} = \tilde{\psi}^1_{kit} K^1_{it} + \cdots + \tilde{\psi}^{N_{kt}}_{kit} K^{N_{kt}}_{it} + \tilde{\psi}^1_{\ell it} L^1_{it} + \cdots + \tilde{\psi}^{N_{\ell t}}_{\ell it} L^{N_{\ell t}}_{it}$. Adding new procedures is then isomorphic to increasing $\tilde{\psi}_{kit}$ or $\tilde{\psi}_{\ell it}$ in the baseline model — you only use the procedure that produces a task with the lowest cost.

We also explored adding another CES layer with love-of-variety *above* our current

[1] Nested CES specifications can also feature this richness as, e.g., Krusell, Ohanian, Ríos-Rull, and Violante (2000) utilized.

[2] Strictly speaking, this requires a discrete number of tasks rather than the continuum.

task CES. Then new varieties of goods could be invented and production of those new varieties could require tasks that did not previously exist. This approach complicates the model substantially while delivering many of the same predictions as our current setup. This would be a useful direction to explore in future research. The new procedures approach yields a substantially simpler model, so we use that in this paper.

2.2 Allocating Inputs to Tasks

Resources are allocated via a competitive equilibrium. A representative firm chooses how to allocate a given amount of capital and labor across tasks in order to maximize profits, taking output and factor prices as given:

$$\max_{\{K_{it},L_{it}\}} P_t Y_t - w_t \int_0^1 L_{it}\, di - r_t \int_0^1 K_{it}\, di \tag{5}$$

subject to (1), (2), (3), and (4).

It is optimal to use capital to produce task i whenever

$$\frac{\psi_{kit}}{r_t} \geq \frac{\psi_{\ell it}}{w_t} \tag{6}$$

and to use labor when the inequality goes the other way. We therefore define the set of tasks using capital and labor as

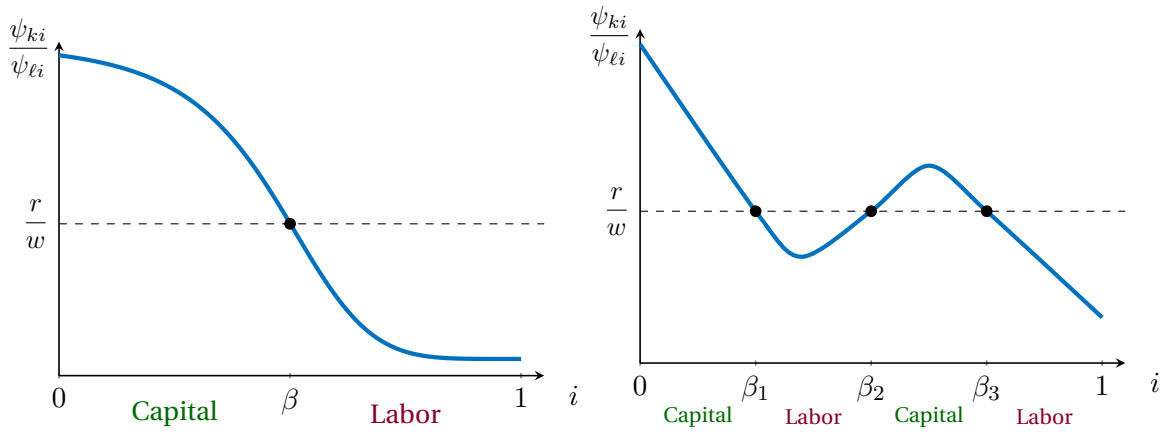
$$\Omega_{kt} = \{i \in [0,1] \mid \psi_{kit}/\psi_{\ell it} \geq r_t/w_t\}$$
$$\Omega_{lt} = \{i \in [0,1] \mid \psi_{kit}/\psi_{\ell it} < r_t/w_t\}$$

Notice that any task β that is just at the margin of being automated — that is, a task at the boundary of the two sets — satisfies the automation condition

$$\frac{\psi_{k\beta t}}{\psi_{\ell\beta t}} = \frac{r_t}{w_t}. \tag{7}$$

Single versus multiple points of automation. We assume that at any point in time, there are a finite number of such indifference points. In that case, Ω_{kt} consists of the union of a finite number of subintervals of $[0,1]$. Let β_t denote the measure of tasks that are produced with capital — the measure of Ω_{kt} — and $1-\beta_t$ denote the measure

Figure 1: Automation and Comparative Advantage: Examples



(a) Single automation point

(b) Multiple automation points

of tasks that are produced with labor. In the special case in which $\frac{\psi_{kit}}{r_t} = \frac{\psi_{\ell it}}{w_t}$ occurs for only a single task, β_t is the task where the crossing occurs and capital is used on the interval $[0, \beta_t]$ while labor is used on $[\beta_t, 1]$. See the left panel of Figure 1. This is a canonical example that is helpful to keep in mind in understanding the model.

We also allow for the more general case in which there are multiple points of automation as in the right panel of Figure 1. Let M_t denote the number of points of automation and call those marginal tasks β_{mt} for $m = 1, \ldots, M_t$.

For convenience, we make an assumption that will hold throughout the paper:

Assumption 1: Technological change is such that there is only automation and no "de-automation."

This assumption states that once tasks transition from being produced with labor to being produced with capital, they never switch back. Particularly since our empirical work is based on time periods of a decade or longer, this strikes us as a plausible assumption worth the simplification in notation and improved expositional clarity.

2.3 Production Function CES Representation

The task-basked production function can be represented as a standard CES-like production function, which provides a link between our growth accounting framework and

traditional measures.

Proposition 1 (*Reduced-form production function*)**.** In equilibrium, output Y_t can be represented as a familiar CES-like production function:

$$\begin{aligned} Y_t &= F(B_t K_t, A_t L_t) \\ &= \left((B_t K_t)^{\frac{\sigma-1}{\sigma}} + (A_t L_t)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \end{aligned}$$

where

$$B_t = Z_t \left(\int_{\Omega_{kt}} \psi_{kit}^{\sigma-1} \, di \right)^{\frac{1}{\sigma-1}} \quad \text{and} \quad \psi_{kit} = \alpha_i^{\frac{\sigma}{\sigma-1}} \tilde{\psi}_{kit}$$

$$A_t = Z_t \left(\int_{\Omega_{\ell t}} \psi_{\ell it}^{\sigma-1} \, di \right)^{\frac{1}{\sigma-1}} \quad \text{and} \quad \psi_{\ell it} = \alpha_i^{\frac{\sigma}{\sigma-1}} \tilde{\psi}_{\ell it}.$$

Factor shares are

$$s_{Kt} \equiv \frac{r_t K_t}{P_t Y_t} = \left(\frac{B_t K_t}{Y_t} \right)^{\frac{\sigma-1}{\sigma}} \quad \text{and} \quad s_{Kit} \equiv \frac{r_t K_{it}}{P_t Y_t} = \left(\frac{\psi_{kit} Z_t K_{it}}{Y_t} \right)^{\frac{\sigma-1}{\sigma}}$$

$$s_{Lt} \equiv \frac{w_t L_t}{P_t Y_t} = \left(\frac{A_t L_t}{Y_t} \right)^{\frac{\sigma-1}{\sigma}} \quad \text{and} \quad s_{Lit} \equiv \frac{w_t L_{it}}{P_t Y_t} = \left(\frac{\psi_{\ell it} Z_t L_{it}}{Y_t} \right)^{\frac{\sigma-1}{\sigma}}$$

According to Proposition 1, our task-based approach has a reduced-form CES-like representation with capital-augmenting productivity B_t and labor-augmenting productivity A_t, where the heterogeneous share parameters, α_i, are folded into the productivity parameters. The endogenous automation share β_t enters both B_t and A_t. Since β_t changes whenever w_t/r_t changes, A_t and B_t change as well; in other words, the elasticity of substitution between capital and labor is no longer given by σ when automation is allowed to adjust. Still, the representation remains useful since A_t, B_t, and β_t do not depend on K_t and L_t, so the standard CES factor share formulas are valid.

Are computers an A_t or a B_t? A question that has long been puzzling in the growth literature is how to interpret A_t and B_t. For example, there is a long tradition of specifying a CES production function like that in Proposition 1 as the primitive and thinking about capital-augmenting versus labor-augmenting technical change. But this leads to obvious questions that do not have obvious answers. For example, is a better computer an increase in A_t or B_t? Is it like having twice as many old computers ($\uparrow B_t$) or does

it effectively increase the user's time endowment ($\uparrow A_t$)? Much of the literature has answered this question by saying that better computers and information technology show up as investment-specific technological change — the same as an increase in B_t for our purposes.[3] This literature uses hedonics and sharply-declining information technology prices to measure changes in B_t — examples include Greenwood, Hercowitz, and Krusell (1997), Herrendorf, Rogerson, and Valentinyi (2020), and Caunedo, Jaume, and Keller (2023). But it is not obvious that this is the right thing to do.

An advantage of the task model is that it provides a framework in which the answer to this question is clear: better computers are an increase in ψ_{kit} for the tasks that have been automated using computers. For tasks that are performed purely with labor, a computer does not make labor better at that task. Of course because those tasks are complementary with other tasks that use a computer, a worker's productivity and wage can rise with automation. But here, a better computer is clearly an increase in ψ_{kit}. This insight will be useful in interpreting the results from our applications.

2.3.1 Growth rates of B_t and A_t

Our main growth accounting exercise is in terms of growth in primitives ψ_{kit}, $\psi_{\ell it}$, and Z. First, we find it instructive to perform growth accounting in the more standard $F(BK, AL)$ representation. Recall that

$$B_t = Z_t \left(\int_{\Omega_{kt}} \psi_{kit}^{\sigma-1} \, di \right)^{\frac{1}{\sigma-1}}$$

Taking the time derivative of B_t and expressing everything in terms of growth rates ($\hat{X}_t \equiv \dot{X}_t / X_t$) leads to our next result.

Proposition 2 (*Growth of B_t and A_t*)**.** Under the "no de-automation" Assumption 1,

$$\hat{B}_t = \hat{Z}_t + \hat{\psi}_{kt} - \frac{1}{1-\sigma} \bar{\omega}_{k\beta t} \dot{\beta}_t \quad \text{where} \quad \hat{\psi}_{kt} \equiv \int_{\Omega_{kt}} \hat{\psi}_{kit} \omega_{kit} \, di \tag{8}$$

$$\hat{A}_t = \hat{Z}_t + \hat{\psi}_{\ell t} + \frac{1}{1-\sigma} \bar{\omega}_{\ell \beta t} \dot{\beta}_t \quad \text{where} \quad \hat{\psi}_{\ell t} \equiv \int_{\Omega_{\ell t}} \hat{\psi}_{\ell it} \omega_{\ell it} \, di \tag{9}$$

[3] There is, of course, a difference between capital-augmenting technical change and investment-specific technical change. The latter only affects new capital while the former affects all capital. However, this distinction is not important for the points made here.

where $\omega_{\ell it}$ is task i's labor cost share, ω_{kit} is task i's capital cost share, $\bar{\omega}_{k\beta t}$ is the share of capital costs for the tasks that are being automated, $\bar{\omega}_{\ell\beta t}$ is the share of labor costs for the tasks that are being automated, and $\dot{\beta}_t$ is the total flow of automation that occurs across the different automation points.

Proof. See Appendix B.

The growth rate of B_t is the sum of three terms. First is the general TFP growth via $\hat{Z}_t$. Second is $\hat{\psi}_{kt}$, which is a weighted average of the growth rates of ψ_{kit} on the already automated tasks, capturing the gains from better computers, machine tools, software, etc. The importance of each task is given by the weight in the average, which is its cost share. The third and final term is the automation effect associated with an increase in the fraction of tasks that have been automated, β_t. This third term is negative because of a "capital depletion" effect: spreading a given amount of capital over a larger number of tasks reduces capital per task and shows up as a decrease in productivity when $\sigma < 1$ (Aghion et al., 2019). For intuition from a simple model see Appendix B.3 for the special case with homogeneous productivities.

The same logic applies to A_t. There are again three terms, capturing general TFP growth ($\hat{Z}_t$), average productivity improvements within the tasks that use labor ($\hat{\psi}_{\ell t}$), and an automation effect. In this case, the automation effect is positive. Mathematically, the share of tasks using labor is $1-\beta_t$, which leads to the additional negative sign. Economically, a given amount of labor is being concentrated onto fewer tasks, so labor per task rises. Because the tasks are complements instead of substitutes, more labor per task raises productivity.

2.3.2 Total Factor Productivity Growth

To see the overall growth consequences, we now turn to the growth rate of total factor productivity. The standard Solow approach implies

$$\hat{Y}_t = s_{Kt}\left(\hat{B}_t + \hat{K}_t\right) + s_{Lt}\left(\hat{A}_t + \hat{L}_t\right).$$

Table 1: Data and Sources

Short name	NAICS code	Sector full name	Years
Private business	—	Private business sector	1950–2023
Agriculture	—	Agriculture	1950–2021
Computers	334	Computer and electronic products	1987–2017
Motor vehicles	3361–63	Motor vehicles, bodies and trailers, and parts	1987–2017
Retail trade	44–45	Retail trade	1987–2017
Software	511, 516	Publishing industries (includes software)	1987–2017

Rearranging this expression leads to TFP growth:

$$\underbrace{\hat{Y}_t - s_{Kt}\hat{K}_t - s_{Lt}\hat{L}_t}_{\widehat{TFP}_t} = s_{Kt}\hat{B}_t + s_{Lt}\hat{A}_t. \tag{10}$$

TFP growth is the weighted average of growth in B_t and A_t where the weights are the production elasticities (which equal the factor shares).

2.4 Data and Empirics

Our baseline data source is the BEA/BLS Integrated Industry-level Production Account (KLEMS) that covers around 60 sectors of the U.S. economy from 1987 to 2021. For the aggregate economy, we use the private business sector multifactor productivity data from 1950 to 2023 from U.S. Bureau of Labor Statistics (2025). For agriculture, our data are from the U.S. Department of Agriculture for 1950 to 2021 (Wang et al., 2024) . Table 1 lists the sources of the data for the various sectors.

The elasticity of substitution, σ. The elasticity of substitution, σ, is a key parameter of the model. For our baseline calculation, we assume a value of $\sigma = 0.3$ and explore robustness to this choice. Two comments are warranted regarding this choice.

First, as discussed above in our near-CES representation result in Proposition 1, σ would be the elasticity of substitution between capital and labor in our reduced-form

representation if the automation set were held fixed. Allowing the automation set to change means that the elasticity of substitution between capital and labor is greater than σ. A large literature estimating the elasticity of substitution between capital and labor almost invariably finds values less than one. The meta analysis of Gechert et al. (2022) argues for a mean estimate of 0.3. Recent papers support this view. Oberfield and Raval (2021) estimates values between 0.5 and 0.7, while Young (2025) finds estimates of around 0.4-0.5. These are estimates of the elasticity of substitution between capital and labor, which is an upper bound on σ, justifying our baseline of $\sigma = 0.3$.

Second, it is useful to consider the following question: We know that the share of factor income paid to capital has risen in recent years. What has happened to the share of factor income paid to computers? On the one hand, computers are everywhere. The number of transistors on a computer chip today is 50 million times more than it was in the 1970s. On the other hand, the price of compute has plummeted, suggesting that the marginal product of computing power has as well. Which effect dominates?

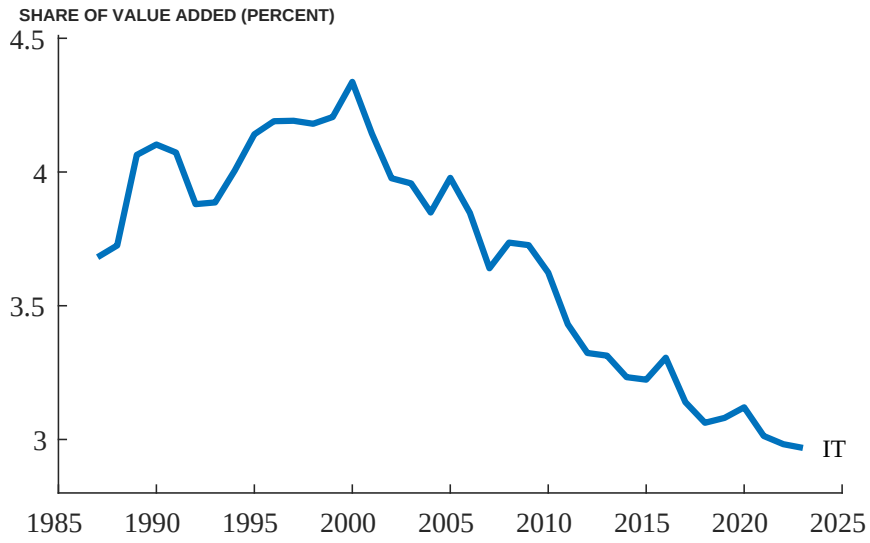
Figure 2 shows the answer. During the dot-com era of the late 1990s, the factor share of income for computers rose from around 3.7 percent to 4.3 percent. But since 2000, the share has fallen substantially to 3.0 percent. In other words, even though the amount of computing power has exploded, we pay less of our GDP as a return to computers today than in the past. This is exactly what a production function with an elasticity of substitution less than one would predict. And this fact may itself be very informative about the effects of future A.I.-driven automation on the economy.

Measurement. We identify and measure the key variables as follows. First, our data sources provide us with a measure of total factor productivity growth, labor productivity, and the share of factor payments to capital and labor.

Second, we assume there is perfect competition in markets so that the production function elasticities are equal to the shares of factor payments. With our reduced-form CES production function, this means that

$$s_{Lt} \equiv \frac{w_t L_t}{P_t Y_t} = \frac{\partial \log Y_t}{\partial \log L_t} = \left(\frac{A_t L_t}{Y_t}\right)^{\frac{\sigma-1}{\sigma}} \tag{11}$$

Figure 2: The Share of Factor Income Paid to Computers



Note: The factor income share of information technology in the private business sector has declined over the past 25 years. *Source:* Bureau of Labor Statistics (2024).

We therefore recover A_t from data as

$$A_t = s_{Lt}^{\frac{\sigma}{\sigma-1}} \cdot \frac{Y_t}{L_t} \tag{12}$$

In other words, A_t is just labor productivity adjusted by the labor share. Then we recover the growth rate of B_t so that the TFP growth accounting equation (10) holds exactly. We emphasize this point so that it is transparent that the calculation of the factor augmenting growth rates of A_t and B_t is straightforward.

Growth in TFP, A_t, and B_t. Table 2 shows the growth in total factor productivity, A_t, and B_t for the various sectors as well as the average factor shares.

The first row shows the data for the "aggregate" sector, corresponding to the private business sector in the BLS multifactor productivity data. TFP growth between 1950 and 2023 averaged 1.2% per year. The capital share averaged 0.35 and the labor share averaged 0.65.

More interesting is the breakdown into growth in A_t versus B_t. For the private business sector, the growth rate of labor augmenting productivity A_t was 2.3% per year,

Table 2: TFP Growth: Basic Data

Sector	TFP growth	Growth in B_t	Growth in A_t	Factor share of capital	Factor share of labor
Private business	1.2	-1.0	2.3	0.35	0.65
Agriculture	3.3	2.6	4.4	0.57	0.43
Computers	12.8	9.5	15.0	0.41	0.59
Motor vehicles	1.7	0.1	2.9	0.43	0.57
Retail trade	1.7	-1.7	2.5	0.20	0.80
Software	1.8	-1.2	4.6	0.47	0.53

Note: Growth rates are average annual log changes. Agriculture and the private business sector start in 1950. For the other sectors, the data cover 1987 to 2017. Factor shares are averages over the entire period.

a conventional number. However, the growth rate of B_t, the capital augmenting component, was actually negative, at -1.0% per year. TFP growth is a weighted average of $\hat{B}$ and $\hat{A}$, and the $\hat{A}$ term is large enough that a negative $\hat{B}$ is needed to explain the 1.2% TFP growth rate: $1.2 = .35 \cdot (-1.0) + .65 \cdot 2.3$. Notice that to the extent that factor shares are stable, this calculation is largely invariant to the elasticity of substitution, σ.

Looking at the sectoral data in Table 2, several results are worth noting. First, TFP growth ranges from a low of 1.7% per year in motor vehicles and retail trade to a high of 12.8% per year in the computer sector. Second, the growth rate of A_t is always substantially higher than TFP growth, which means that the growth rate of B_t will be lower. For agriculture, computers, and motor vehicles, the growth rate of B_t is positive, while for the other sectors the growth rate of B_t is negative, as it was for the aggregate.

Young (2025) first provided detailed empirical evidence for negative growth in B_t and positive growth in A_t using gross output KLEMS data for the United States. Our evidence confirms his result for value-added based TFP measures.

The model developed so far — especially Proposition 2 — helps us make sense of the negative growth in B_t that we often observe. As in Aghion, Jones, and Jones (2019), automation is simultaneously labor augmenting ($\uparrow A_t$) and capital depleting ($\downarrow B_t$). Declining B_t can be a sign of automation: average capital per automated task declines, which reduces "effective capital" and therefore reduces B_t. The positive growth in B_t

for agriculture and computers can be explained by the neutral productivity term Z_t — positive growth in Z_t will increase both $\hat{B}_t$ and $\hat{A}_t$ — or by rapid ψ_{kt} growth and a modest increase in the share of tasks that are automated.

3. Automation Growth Accounting: Theory and Evidence

We now use the structure of the model to uncover the consequences of automation for total factor productivity growth, both in the theory and in the data.

Combining equation (10) with equations (8) and (9) yields our main decomposition:

$$\widehat{TFP}_t = \hat{Z}_t + \underbrace{s_{Kt}\hat{\psi}_{kt}}_{\text{Better capital}} + \underbrace{s_{Lt}\hat{\psi}_{\ell t}}_{\text{Better labor}} + \underbrace{\frac{\dot{\beta}_t}{1-\sigma}\left(s_{Lt}\bar{\omega}_{\ell\beta t} - s_{Kt}\bar{\omega}_{k\beta t}\right)}_{\text{Automation effect}} \tag{13}$$

Total factor productivity growth can be decomposed into four terms: "other" productivity growth $\hat{Z}_t$, improvements in the productivity of capital $\hat{\psi}_{kt}$, improvements in the productivity of labor $\hat{\psi}_{\ell t}$, and the overall effect of automation (the sum of the two automation terms). Each term is weighted by its cost share.

The automation term. We now show that when automation is smooth, the automation effect in the TFP decomposition is zero. Automation is smooth when the time derivatives $\dot{\psi}_{kit}$, $\dot{\psi}_{\ell it}$, $\dot{r}_t$, and $\dot{w}_t$ exist, which ensures the automation indifference condition (7) holds at all points of automation. In particular, use the definition of the weights in Proposition 2 to notice that at any point of automation, β, the automation effect term depends on

$$s_{Lt}\omega_{\ell\beta t} - s_{Kt}\omega_{k\beta t} = \frac{w_t L_t}{P_t Y_t}\cdot\frac{w_t L_{\beta t}}{w_t L_t} - \frac{r_t K_t}{P_t Y_t}\cdot\frac{r_t K_{\beta t}}{r_t K_t} \tag{14}$$

$$= \frac{w_t L_{\beta t}}{P_t Y_t} - \frac{r_t K_{\beta t}}{P_t Y_t} \tag{15}$$

That is, the automation term depends on the cost of doing the marginally-automated task with labor versus the cost of doing it with capital. But at the margin, these costs must be the same, so this difference is zero. This is true at any point of automation and so it is also true for the sum over all the points. Hence, the overall automation effect

term is zero as well.

We therefore have the following useful result:

Proposition 3 (*Zero TFP growth from smooth automation*)**.** When the marginal task that is automated satisfies the indifference condition (7), the automation effect in the TFP decomposition in equation (13) is zero. Therefore, TFP growth equals

$$\begin{aligned} \widehat{TFP}_t &= s_{Kt}\hat{B}_t + s_{Lt}\hat{A}_t \\ &= \underbrace{s_{Kt}\hat{\psi}_{kt}}_{\text{Capital productivity}} + \underbrace{s_{Lt}\hat{\psi}_{\ell t}}_{\text{Labor productivity}} + \underbrace{\hat{Z}_t}_{\text{Other TFP growth}} \end{aligned} \tag{16}$$

Discussion. This proposition states that the contribution of automation to TFP growth is zero. What is going on? To understand, notice that the ψ_{kit} and $\psi_{\ell it}$ are the fundamental primitives of the problem and β_t is an endogenous variable. As the ψ's change, tasks get automated. When automation is continuous, there will be marginal tasks for which it is equally costly to use capital or labor. This is the indifference condition in (7) that implicitly pins down β_t. But that means that whenever automation occurs, this indifference condition is satisfied, and therefore each instant of automation cannot lead to TFP growth.

Instead, the proposition states that, apart from other sources in $\hat{Z}_t$, TFP growth is simply the weighted average of productivity improvements on the capital tasks and the labor tasks. It is entirely the "within" terms and there is no composition effect. The weights, in turn, are the standard factor cost shares, s_{Kt} and s_{Lt}.

So one of our key answers to the question "How much of growth is due to automation?" is zero! We will provide alternative definitions of automation below that lead to different answers, but this result makes clear that the answer depends on what precisely we mean by automation.

There is a parallel here between the firm dynamics literature and the contribution of new entrants to growth. How much of growth is due to the entry of firms that turn out to be superstars, such as Apple or Google? Well, the answer depends on how much of these firms subsequent growth is attributed to entry. For example, if all new firms since the year 1900 are counted as entrants, then entry accounts for nearly 100% of growth.

On the other hand, if the growth of new entrants is counted only during the first year (and after that Apple and Google are treated as "incumbents"), then very little growth is due to new entry. The issue is similar here: how much of the subsequent growth from better computers gets included in the automation term? If none, then the contribution of automation is zero. But in this alternative decomposition, the length of a period determines how much of the "better computers" gets attributed to automation. In empirical case studies that compare a treatment firm that automated to a control firm that did not automate, the measured effect of automation reflects the improvements in the productivity of capital since the point of automation (plus the initial jump in productivity if automation is not smooth). We provide an alternative measurement of the effect of automation on growth drawing on this perspective in Section 4.

3.1 The Effect of Automation on TFP Growth

Given that the automation "composition effect" is zero, it is helpful to consider what else in our environment might be related to automation and contribute to TFP growth. To see our next important result, it is helpful to substitute $s_{Lt} = 1 - s_{Kt}$ into the TFP decomposition in Proposition 3 to get

$$\widehat{TFP}_t = \underbrace{\hat{Z}_t + \hat{\psi}_{\ell t}}_{\text{Baseline TFP growth}} + \underbrace{s_{Kt}\left(\hat{\psi}_{kt} - \hat{\psi}_{\ell t}\right)}_{\text{Automation effect: boost from machines getting better}} \tag{17}$$

Equation (17) is one of the key equations of the paper. TFP growth is the sum of two terms. The first reflects baseline TFP growth from "other" sources ($\hat{Z}_t$) and the minimum improvement on all tasks from people getting better ($\hat{\psi}_{\ell t}$). The second term is the additional boost that comes from automation: machines may get better at a faster rate than humans, boosting TFP growth by $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$. But this boost only applies to a fraction of the tasks. Intuitively, the proper way to measure this fraction is by the cost share of tasks using capital — the factor income share s_{Kt}.

What we show in the remainder of this section is that the "Baseline TFP growth" term in (17) is small in most sectors. Therefore the majority of TFP growth is due to the

"Automation effect," the boost that comes from machines getting better at a faster rate than humans.

3.2 Identification

To implement equation (17) empirically, we need to identify the terms on the right-hand side. From the previous section, we have measures of TFP growth, $\hat{A}_t$, $\hat{B}_t$, and factor income shares. We now explain how we identify the remaining variables.

Step 1. First, recall equation (9):

$$\hat{A}_t = \hat{Z}_t + \hat{\psi}_{\ell t} + \frac{1}{1-\sigma}\bar{\omega}_{\ell\beta t}\dot{\beta}_t \tag{18}$$

Consider the $\bar{\omega}_{\ell\beta t}$ part of the equation. From equation (32), $\omega_{\ell\beta t} = \frac{L_{\beta t}}{L_t}$. That is, the weight for a particular task is just the fraction of labor used in the newly-automated tasks. Then, $\bar{\omega}_{\ell\beta t}$ is a weighted average of the ω's at the points of automation (from equation 34). That is,

$$\bar{\omega}_{\ell\beta t} = \frac{L_{\bar{\beta}t}}{L_t}$$

where $L_{\bar{\beta}t}$ is a weighted average of the labor used in the tasks that are being automated.[4]

Step 2. Next, labor is used on the share $1-\beta_t$ of the tasks, so the average amount of labor used per task is $L_t/(1-\beta_t)$. This leads to our second assumption:

Assumption 2: $L_{\bar{\beta}t} \geq \frac{L_t}{1-\beta_t}$

This assumption says that the marginal tasks that are automated on average use at least as much labor as the average of the tasks that are not yet automated. This assumption is important and it merits discussion. For the moment, let's accept it and derive its consequences. After stating our identification result, we will discuss the assumption in detail.

[4] In particular, $L_{\bar{\beta}t} = \sum_{m=1}^{M_t} L_{\beta_m t}\frac{|\dot{\beta}_{mt}|}{\dot{\beta}_t}$.

Assumption 2 implies that $\bar{\omega}_{\ell\beta t} \geq \frac{1}{1-\beta_t}$. Substituting this into equation (18) gives

$$\begin{aligned}\hat{A}_t &= \hat{Z}_t + \hat{\psi}_{\ell t} + \frac{1}{1-\sigma}\bar{\omega}_{\ell\beta t}\dot{\beta}_t \\ &\geq \hat{Z}_t + \hat{\psi}_{\ell t} + \frac{1}{1-\sigma}\frac{\dot{\beta}_t}{1-\beta_t}\end{aligned} \tag{19}$$

Step 3. Now consider the last term in the preceding equation and make the following notational definition:

Definition 1: *The automation rate is defined as*

$$x_t \equiv \frac{\dot{\beta}_t}{1-\beta_t} = -\frac{d\log(1-\beta_t)}{dt} \tag{20}$$

The variable x_t has an elegant economic interpretation as the rate of automation. Notice that $\dot{\beta}_t = x_t(1-\beta_t)$. That is, x_t is the fraction of the labor tasks that get automated in period t.

Substituting this definition into equation (19) and rearranging, we get

$$\hat{A}_t - \frac{1}{1-\sigma}x_t \;\geq\; \hat{Z}_t + \hat{\psi}_{\ell t} \tag{21}$$

This equation is important in that — provided we can construct an empirical measure of the automation rate x_t — it gives an upperbound on the baseline term in TFP growth from equation (17).

Step 4. Finally, substituting this expression back into that TFP growth equation (17) allows us to obtain a lower bound on $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$.

Our next proposition collects two important bounds on the terms that make up TFP growth in equation (17):

Proposition 4 (*Identifying key bounds on $\hat{\psi}_{kt}$ and $\hat{\psi}_{\ell t}$*)**.** Provided we can measure the automation rate x_t,

$$\hat{A}_t - \frac{1}{1-\sigma}\, x_t \;\geq\; \hat{Z}_t + \hat{\psi}_{\ell t}$$

is an upper bound on the baseline component of TFP growth in equation (17) and

$$\hat{\psi}_{kt} - \hat{\psi}_{\ell t} \geq \frac{1}{s_{Kt}} \left(\widehat{TFP}_t - \left[\hat{A}_t - \frac{1}{1-\sigma}\, x_t \right] \right)$$

is a lower bound on the automation boost associated with the average rate at which machines are getting better.

The intuition for how the proposition works can be seen most easily by considering the homogeneous case in which $\psi_{\ell it} = \psi_{\ell t}$:

$$A_t^{homog} = \left(\frac{1}{1-\beta_t} \right)^{\frac{1}{1-\sigma}} Z_t \psi_{\ell t}$$

The $\left(\frac{1}{1-\beta_t}\right)^{\frac{1}{1-\sigma}}$ term appears as a given amount of labor gets spread over $1-\beta_t$ tasks: concentrating labor on a smaller range of tasks raises labor per task, increasing labor productivity – the opposite of the "love of variety" effect since $\sigma < 1$. Taking logs and derivatives of this equation we see that

$$\hat{A}_t^{homog} - \frac{1}{1-\sigma}\, x_t = \hat{Z}_t + \hat{\psi}_{\ell t}$$

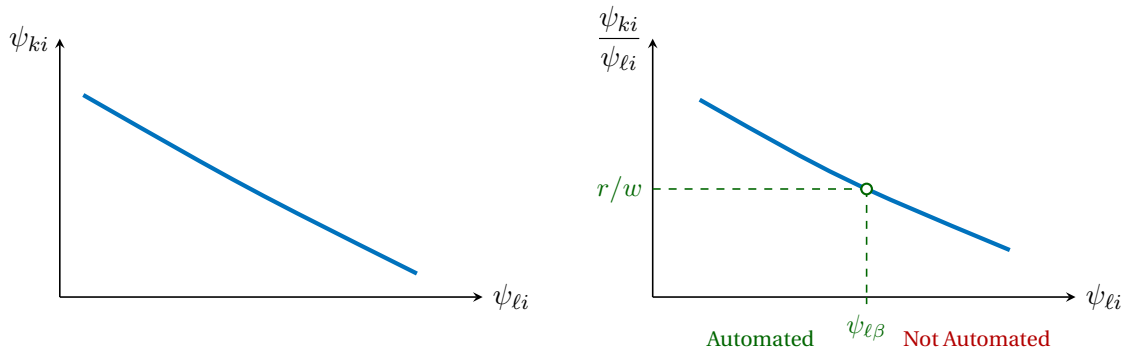
In Proposition 4, we get a very similar relationship in the heterogeneous $\psi_{\ell it}$ case but as an inequality. Substituting this inequality into the TFP growth equation allows us to get a lower bound on $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$.

This discussion justifies the following additional result:

> *Corollary* — In the special case in which $\psi_{\ell it} = \psi_{\ell t}$ — i.e., in the case in which the labor productivity terms are homogeneous while heterogeneity occurs in the ψ_{kit} only — the inequalities in Proposition 4 become equalities.

Most of the remainder of the paper is devoted to applying these results empirically. But first we need to discuss why Assumption 2 is a reasonable assumption.

Figure 3: Moravec's Paradox



(a) Moravec's Paradox

(b) Comparative advantage

Note: Tasks that are hard for humans tend to be easy for machines and tasks that are easy for humans tend to be hard for machines. Comparative advantage is also therefore negatively related to ψ_ℓ.

3.3 Discussion of Assumption 2

Assumption 2 requires that among the tasks that we have not yet automated, we first automate the tasks that have high labor costs, at least on average.

This assumption turns out to be a natural consequence of what is known as *Moravec's Paradox*: tasks that are hard for humans tend to be easy for machines and tasks that are easy for humans tend to be hard for machines (Moravec, 1988). Evolution optimized walking, dexterity, and vision in humans over millions of years, so it is hard for machines to do better, while evolution did not optimize for playing chess or solving complex math problems. Figure 3a illustrates this stylized fact by showing a negative relationship between ψ_{kit} and $\psi_{\ell it}$.

Because automation is about comparative advantage, it is helpful to make the same plot but with $\psi_{ki}/\psi_{\ell i}$ on the vertical axis. Notice that a negative relationship in Figure 3a implies the negative relationship in Figure 3b: dividing by ψ_ℓ only reinforces the negative relationship.

In the case in which there is only a single automation point, Assumption 2 states

that

$$\omega_{\ell\beta t} \equiv \frac{\psi_{\ell it}^{\sigma-1}}{\int_{\Omega_{\ell t}} \psi_{\ell it}^{\sigma-1}\, di} = \frac{L_{\beta t}}{L_t} \geq \frac{1}{1-\beta_t}$$

With $\sigma < 1$, this is a natural consequence of Moravec's Paradox. The marginal task that is automated has the lowest $\psi_{\ell i}$ and therefore the highest $\psi_{\ell i}^{\sigma-1}$.

One final point of clarification relates to another difference between comparative and absolute advantage. In our initial setup (e.g., see Proposition 1), $\psi_{\ell it} = \alpha_i^{\frac{\sigma}{\sigma-1}} \tilde{\psi}_{\ell it}$ where the α_i is a common share parameter on task i regardless of whether or not it is performed by humans or machines whereas the $\tilde{}$ variables are the input-specific productivity terms.

The discussion so far implicitly assumed that $\alpha_i = 1$ for all tasks. To generalize to the case in which these share parameters are present, notice two things. First, because the share parameters enter both ψ_k and ψ_ℓ, they do not affect comparative advantage: $\frac{\psi_{ki}}{\psi_{\ell i}} = \frac{\tilde{\psi}_{ki}}{\tilde{\psi}_{\ell i}}$. But the share parameters do affect absolute advantage: $\psi_{\ell it} = \alpha_i^{\frac{\sigma}{\sigma-1}} \tilde{\psi}_{\ell it}$. A natural assumption is that α_i is uncorrelated with $\tilde{\psi}_{\ell it}$. In that case, the presence of the share parameters simply causes the curve in Figure 3b to "flatten," making the inequality in Assumption 2 less tight (closer to equality).

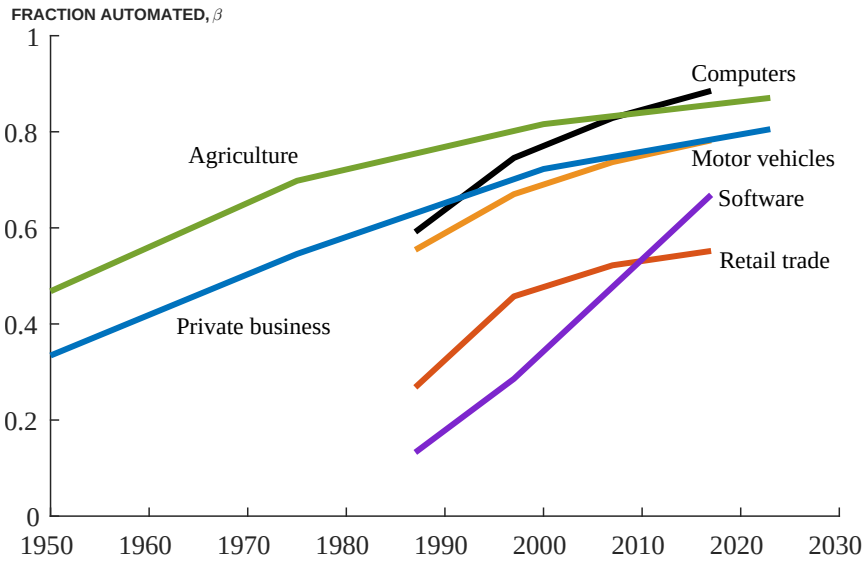
The bottom line is that Moravec's Paradox suggests that Assumption 2 is a reasonable assumption. Importantly, this is also an assumption that future empirical work can seek to measure and test.

3.4 Measuring β_t

Almost all of the terms in Proposition 4 can be measured with our BEA/BLS data. The exception is the share of tasks that have been automated, β_t, and in particular, the automation rate $x_t \equiv \frac{\dot{\beta}_t}{1-\beta_t}$.

This brings us to a crucial part of our analysis: how we measure the share of tasks that have been automated at each point in time. We considered hiring a team of RA's to scour the voluminous literature for each of our sectors to construct this measure. However, this would have been a major undertaking and we would have had to iteratively develop a detailed rubric to keep the RA's methods uniform across people, time, and tasks. It then occurred to us that this is a perfect task for OpenAI's Deep Research state-

Figure 4: Share of Tasks that are Automated, β_t



Note: See Appendix A for the details of how we measure β_t.

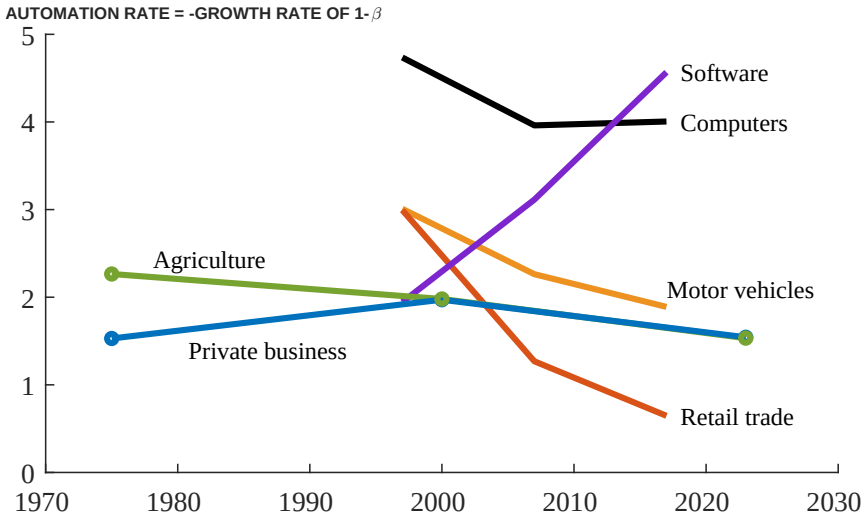
of-the-art LLM. The instructions we used for this task are reported in Appendix A.[5]

Figure 4 show the data on the share of tasks that are automated, β_t, by sector and over time. There is substantial heterogeneity, both across sectors and over time. For example, in recent years, retail trade has the lowest share of tasks that have been automated — around 50% — whereas computers/electronics has the highest share — more than 85%. In contrast, in 1950, less than half of tasks in agriculture were automated and only around 1/3 of tasks in the private business sector were automated.

Finally, Figure 5 shows the automation rate. Recall that $x_t \equiv \frac{\dot{\beta}_t}{1-\beta_t}$, so the automation rate measures the fraction of tasks performed by labor that get automated during a particular period of time. Automation rates for the private business sector and for agriculture are relatively stable over time since 1950 and equal to around 2% per year, having slowed slightly since 2000. For the BEA/BLS sectors, automation rates range from a low of under 1% for retail trade to a high of more than 4% for software and

[5]In the future, we will refine these instructions and repeat the exercise many times to ensure consistency and create bootstrap standard errors. Also, recall that in equation (1), there are share parameters (α_i) that govern the importance of each task. Some tasks can be very important while others can be relatively minor. These share parameters get combined with the levels of the technology parameters (in ψ_{kit} and $\psi_{\ell it}$), so in constructing the fraction of tasks that are automated, our model implies we should focus on an equally-weighted average of tasks; we do not need to adjust here for how important each task is.

Figure 5: Automation Rates, $x_t \equiv \frac{\dot{\beta}_t}{1-\beta_t}$



Note: See Appendix A for the details of how we measure β_t.

computers/electronics. Interestingly, software shows a substantial increase in the automation rate over time.

3.5 Empirics: The Contribution of $\hat{\psi}_{kt}$ and $\hat{\psi}_{\ell t}$ to TFP Growth

We can now turn to the empirical analysis of the contribution of $\hat{\psi}_{kt}$ and $\hat{\psi}_{\ell t}$ to TFP growth using Proposition 4. In particular, that proposition derives an upper bound on $\hat{Z}_t + \hat{\psi}_{\ell t}$ and a lower bound on $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$. Now that we've measured the automation rate x_t, we can calculate these bounds.

Table 3 shows the empirical implementation of these two key bounds. The left side of the table begins by calculating the upper bound on $\hat{Z}_t + \hat{\psi}_{\ell t}$ given by $\hat{A}_t - \frac{1}{1-\sigma} x_t$. The striking finding here is that the upper bound on $\hat{Z}_t + \hat{\psi}_{\ell t}$ is small in most sectors. For example, for the private business sector, the upper bound is -0.1% per year. The implication is that for much of the economy, the rate at which people are getting better is small unless other TFP growth, $\hat{Z}_t$, is substantially negative. It is possible that humans are getting worse at tasks so that $\hat{\psi}_{\ell t}$ is negative and this could be offsetting some growth in Z_t. Alternatively, a rise in misallocation could make $\hat{Z}_t$ negative. Either way, it is noteworthy that these two sources of growth combine to contribute little to TFP growth

Table 3: Bounds on $\hat{Z}_t + \hat{\psi}_{\ell t}$ and $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$

Sector	Growth rate of A_t	Automation rate, x_t	Upperbound on $\hat{Z}_t + \hat{\psi}_{\ell t}$	TFP Growth	Capital share, s_{Kt}	Lowerbound on $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$
Private business	2.3	1.7	-0.1	1.2	0.35	3.6
Agriculture	4.4	2.0	1.5	3.3	0.57	3.2
Computers	15.0	4.2	9.0	12.8	0.41	9.5
Motor vehicles	2.9	2.4	-0.5	1.7	0.43	5.5
Retail trade	2.5	1.6	0.2	1.7	0.20	8.5
Software	4.6	3.2	0.0	1.8	0.47	3.8

Note: See Proposition 4 for the equations describing the bounds.

outside of agriculture and computers.

The right side of the Table 3 calculates the lower bound on $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ as in Proposition 4, basically by subtracting the upper bound on $\hat{Z}_t + \hat{\psi}_{\ell t}$ from the TFP growth rate and scaling by the capital share. The key finding is that the lower bound on $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ is large. For the private business sector, the rate is 3.6% per year, and the lower bound is at least 3.2% per year across all sectors. In other words, the average rate at which machines are getting better across automated tasks is substantially higher than task-specific labor productivity growth. Machines get better much faster than people do.

These two finding are of course related. The fact that the $\hat{Z}_t + \hat{\psi}_{\ell t}$ is a small number, possibly even negative, means that the $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ term must explain the bulk of TFP growth.

It may at first seem surprising that these numbers suggest a low value for $\hat{\psi}_{\ell t}$. But two points are worth noting. First, these numbers are contributions to TFP growth; the BLS/BEA have already taken out the contribution from rising educational attainment and other forms of labor quality. Second, consider a thought experiment: how much better is a college graduate today at inverting matrices or writing essays — with no access to computers — than the equivalent college graduate in 1950? Probably not much better, except perhaps for gains in nutrition.

4. Counterfactual: Freezing automation in the initial year

In this section, we consider a counterfactual in which automation is "frozen" in place in some early year. That is, the set of tasks that are automated is fixed after some point in time: $\Omega_{kt} = \Omega_{k0}$. How much lower would TFP growth have been in this counterfactual world?

4.1 Theory

To answer this question, recall our key equation for TFP growth:

$$\widehat{TFP}_t = s_{Kt}(\hat{\psi}_{kt} - \hat{\psi}_{\ell t}) + \hat{Z}_t + \hat{\psi}_{\ell t} \tag{22}$$

For the counterfactual, we need to make assumptions about what happens to $\hat{\psi}_{kt}$, $\hat{\psi}_{\ell t}$, and $\hat{Z}_t$. We make a natural assumption:

Assumption 3 (Counterfactual growth rates)**:** In the counterfactual world in which automation is frozen in some early year ($\Omega_{kt} = \Omega_{k0}$), the average of the task-specific growth rates $\hat{\psi}_{kt}$ and $\hat{\psi}_{\ell t}$ are unchanged, as is $\hat{Z}_t$.

In a structural model with endogenous technological change, this need not be the case: research effort presumably shifts when tasks are automated. However, given that we are using averages across a large number of tasks and given the difficulty of doing anything else, this seems like a reasonable starting point.

Under this assumption, notice that equation (22) implies that the only way TFP growth is altered in the counterfactual is because the factor share s_{Kt} changes. In particular, comparing actual TFP growth to counterfactual TFP growth, $\widehat{TFP}_t^{cf}$, we have

$$\widehat{TFP}_t - \widehat{TFP}_t^{cf} = \left(s_{Kt} - s_{Kt}^{cf}\right)(\hat{\psi}_{kt} - \hat{\psi}_{\ell t}) \tag{23}$$

So the key thing to understand is how the capital factor income share s_{Kt} changes in the counterfactual. In our CES task setup, the capital share is given by

$$\frac{s_{Kt}}{1 - s_{Kt}} = \left(\frac{B_t}{A_t}\frac{w_t}{r_t}\right)^{\sigma-1}$$

Intuitively, the idea behind the counterfactual is to freeze the set of tasks that are automated at the set that prevailed in 1950 or 1987. This will affect the time path of B_t and A_t through their laws of motion, as in Proposition 2. In particular, $\dot{\beta}_t > 0$ lowers B_t and raises A_t. So it lowers B_t/A_t in the actual data we observe, which will have raised the capital share since $\sigma < 1$. In the counterfactual, we set $\dot{\beta}_t^{cf} = 0$ to shut down this channel. This means the capital share in the counterfactual will decline over time: the "machines getting better" force lowers the capital share and we've turned off the "more tasks get automated" force that historically raised the capital share.

The evolution of the capital share in the counterfactual is[6]

$$\frac{s_{Kt}^{cf}}{1-s_{Kt}^{cf}} \leq \frac{s_{Kt}}{1-s_{Kt}} \exp\left(-\int_0^t \frac{1}{s_{K\tau}} x_\tau d\tau\right) \tag{24}$$

Intuitively, the counterfactual capital share will be lower than the actual capital share to the extent that in the historical data there was automation — i.e., to the extent that x_t was positive. The counterfactual capital share starts with the actual capital share and then "undoes" the contribution from automation.

Putting all this together, we have the following proposition:

Proposition 5 (*Counterfactual contribution of automation*)**.** Under Assumptions 1 – 3, the lost TFP growth from "freezing" the set of automated tasks in some historical year satisfies

$$\begin{aligned}\widehat{TFP}_t - \widehat{TFP}_t^{cf} &= \left(s_{Kt} - s_{Kt}^{cf}\right)(\hat{\psi}_{kt} - \hat{\psi}_{\ell t}) \\ &\geq \left(s_{Kt} - \text{upperbound on } s_{Kt}^{cf}\right) \times \text{lower bound on } (\hat{\psi}_{kt} - \hat{\psi}_{\ell t})\end{aligned}$$

where the upper bound on s_{Kt}^{cf} is given by equation (24) and the lower bound on $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ is given in Proposition 4.

[6]The derivation of the result is shown in Appendix D. For the sectoral calculations, we assume r_t/w_t follows an unchanged path since the sectors are small relative to the aggregate economy; for the private business sector, we allow r_t/w_t to change endogenously, which changes the equation slightly.

Table 4: Counterfactual Contribution of Automation

Sector	Capital share $s_{K,T}$	Capital share $s^{cf}_{K,T}$	$\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$	Lost Growth $\widehat{TFP}_t - \widehat{TFP}^{cf}_t$	TFP Growth	Lost Growth Share of $\widehat{TFP}_t$
Automation set frozen in 1950:						
Private business	0.420	0.000	3.6	1.0	1.2	87%
Agriculture	0.655	0.127	3.2	0.9	3.3	28%
Automation set frozen in 1987:						
Computers	0.459	0.033	9.5	2.5	12.8	20%
Motor vehicles	0.524	0.161	5.5	1.0	1.7	57%
Retail trade	0.259	0.022	8.5	0.8	1.7	50%
Software	0.463	0.102	3.8	0.9	1.8	50%

Note: The counterfactual contribution of automation supposes the set of tasks that are automated is frozen in some inital year and is calculated according to Proposition 5. As in the proposition, $s^{cf}_{K,T}$ is an upper bound, $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ is a lower bound, and $\widehat{TFP}_t - \widehat{TFP}^{cf}_t$ and the final share column are lower bounds. Growth rates are r percents per year, averaged over the relevant time period.

4.2 Results: Freezing Automation Substantially Reduces Growth

To implement Proposition 5 empirically, we assume that automation is frozen in place in the initial year for each sector (1950 for private business and agriculture; and 1987 for the other sectors). Table 4 shows the results.

The first two columns of the table show the actual and counterfactual capital shares in the final year (e.g., 2017 or 2023). The actual capital shares are relatively high, ranging from about 25% to 65%. The second column shows the capital share in the final year under the assumption that automation has been frozen since 1950 or 1987. As expected, the counterfactual capital shares are much lower because the automation set is frozen while machines continue to get better. For the private business sector, the counterfactual capital share falls to essentially zero versus an actual share of 42.0%. The reason for this is shown in the next column, which reports $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ (which we already showed back in Table 3). In particular, capital productivity growth is very fast.

The fourth column in Table 4 implements Proposition 5 to compute a lower bound on the "lost growth" that comes from freezing automation in the initial year. For the

private business sector, this lost growth is 1.0 percent per year. This can be compared to the actual TFP growth rate of 1.2 percent. This means that freezing automation in place in 1950 would have cost the economy most of economic growth — 87% between 1950 and 2023 — in the private business sector.

Across the other sectors of the economy, the missing growth ranges from 0.8 percent per year to 2.5 percent per year. The last column expresses this as a share of TFP growth. Across the BEA sectors, freezing automation in 1987 would have cost the economy 20% of growth in the computer sector and 50% or more in motor vehicles, retail trade, and software.

A natural question to ask is "Why is the private business sector share of lost growth so much higher than for the component sectors?" The main reason for this is that we freeze automation in 1950 for the private business sector but only in 1987 for the component sectors. Like the firm-entry analogy given earlier, the closer to today that we freeze automation, the less time there is for lost growth to accumulate. The lost growth comes from the product of two terms: the $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ captures the productivity growth boost from automation, and the importance of this boost is governed by the capital share s_{Kt}. The higher is the growth boost, the more rapidly the capital share declines in the counterfactual, as machines are getting better but automation is frozen.[7]

If we do not increase the share of tasks that are automated, we do not switch over to making rapid progress on enough of our weak links, and those remaining, slowly improving, weak links tank growth. Even if capital productivity skyrocketed to infinty on the tasks that were already automated in 1950, that would not have delivered infinite growth because tasks are complements. Historically, long-run growth occurred because we found ways to rapidly improve the productivity of machines and because we increased the set of tasks that benefited from this rapid growth—strengthening more of our weak links. The bottom line from this exercise is that historical automation has been tremendously important to TFP growth in the U.S. economy.

[7]One might also ask about agriculture, where automation is also frozen in 1950 but the share of lost growth is much smaller at 28%. There are two reasons for this. First, the capital share is much higher — 65.5% in the data in 2021, so it is harder for this share to decline to zero. Second, the automation boost $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ is actually lowest in agriculture, at just 3.2% versus 3.6% for the private business sector.

4.3 Robustness

The results just given assume $\sigma = 0.3$ and assume that ChatGPT provides a valid estimate of the automation rate for the different sectors. Here, we relax these assumptions.

The three panels of Table 5 show a set of results for the upper bound on $\hat{Z}_t + \hat{\psi}_{\ell t}$, the lower bound on $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$, and the lower bound on the share of TFP growth that would be lost if automation were frozen in 1950 or 1987. The robustness checks cover different values of σ and an automation rate x_t that is only half of what ChatGPT reports.

The most conservative numbers in the robustness table occur with $\sigma = 0.1$ or when the automation rate is only 50% of what ChatGPT estimates. For the $\sigma = 0.1$ case, the upper bound on $\hat{Z}_t + \hat{\psi}_{\ell t}$ is 0.4% for the private business sector. The lower bound on $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ is 2.2% for the private business sector and even larger than that amount in all sectors other than software. Finally, at least 57 percent of TFP growth in the private business sector would be lost if automation were frozen in 1950.

5. The Future of A.I.

This section develops an endogenous growth model with automation and calibrates it based on our evidence from past automation. The spirit of the exercise is that A.I. is just the latest form of an automation process that has been ongoing for at least a century. A key feature of the model is that the production of ideas can also be automated, and this is a place where A.I. can play an important role.

To understand the results of the full dynamic model, however, it is helpful to begin with some warm-up exercises that consider extreme versions of automation.

5.1 Static Effects: What if A.I. fully automates software?

Consider an extreme version of automation: what if some fixed set of tasks are automated with infinite productivity? A first instinct is that this would produce infinite output. But that instinct comes from production functions with an elasticity of substitution of at least unity. With an elasticity below one, we are in the "weak links" setting. Being infinitely good at some tasks does not lead to infinite output because production

Table 5: Robustness of Automation Results

Sector	$\sigma = 0.1$	$\sigma = 0.3$	$\sigma = 0.5$	x_t halved
		Upper bound on $\hat{Z}_t + \hat{\psi}_{\ell t}$		
Private business	0.4	-0.1	-0.9	1.1
Agriculture	2.0	1.5	0.6	2.9
Computers	10.1	9.0	7.1	12.0
Motor vehicles	-0.2	-0.5	-1.2	1.2
Retail trade	0.5	0.2	-0.5	1.4
Software	0.9	0.0	-1.7	2.3
		Lower bound on $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$		
Private business	2.2	3.6	6.2	0.1
Agriculture	2.3	3.2	4.9	0.7
Computers	6.9	9.5	14.1	2.0
Motor vehicles	4.5	5.5	7.2	1.4
Retail trade	6.5	8.5	12.1	2.2
Software	1.9	3.8	7.3	-1.0
		Lost TFP growth – freeze automation (percent share; lower bound)		
Private business	57	87	134	...
Agriculture	22	28	39	...
Computers	14	20	30	...
Motor vehicles	48	57	74	...
Retail trade	37	50	73	...
Software	25	50	95	...

Note: The table shows the robustness of our automation results to alternative parameter choices. The first three columns consider different values of the elasticity of substitution across tasks, σ; our baseline is $\sigma = 0.3$. The final column assumes the automation rate x_t is one half of what we have measured.

is constrained by the the weakest links.[8]

Start with our familiar CES production function but collect tasks into two groups: those we will infinitely automate (labeled ∞) and those we will leave unchanged (labeled $\varnothing$). These could be software and non-software, or manufacturing and non-manufacturing, or cognitive and non-cognitive.

$$\begin{aligned} Y_t^{\frac{\sigma-1}{\sigma}} &= \int_0^1 \alpha_i Y_{it}^{\frac{\sigma-1}{\sigma}} di \\ &= \int_{\Omega_\varnothing} \alpha_i Y_{it}^{\frac{\sigma-1}{\sigma}} di + \int_{\Omega_\infty} \alpha_i Y_{it}^{\frac{\sigma-1}{\sigma}} di \\ &= \alpha_\varnothing Y_{\varnothing t}^{\frac{\sigma-1}{\sigma}} + \alpha_\infty Y_{\infty t}^{\frac{\sigma-1}{\sigma}} \end{aligned}$$

where $\alpha_\varnothing \equiv \int_{\Omega_\varnothing} \alpha_i di$ and $\alpha_\infty \equiv \int_{\Omega_\infty} \alpha_i di$.

Perfect competition and first order conditions imply the usual factor share equation:

$$s_{jt} \equiv \frac{P_{jt} Y_{jt}}{P_t Y_t} = \alpha_j \left(\frac{Y_{jt}}{Y_t} \right)^{\frac{\sigma-1}{\sigma}} \tag{25}$$

Implementing infinite automation. Now consider a counterfactual in which the tasks in the ∞ sector are automated with infinite ψ_K — extreme but useful. Y_∞ goes to infinity, so $Y_\infty^{\frac{\sigma-1}{\sigma}}$ goes to zero. Intuitively, infinite automation eliminates some of the weak links.

Assuming no other changes,

$$Y_{cf} = \alpha_\varnothing^{\frac{\sigma}{\sigma-1}} Y_{\varnothing t}.$$

[8] Related points appear in Aghion, Jones, and Jones (2019) and B. Jones and X. Liu (2024), but the result is stated in terms of the fraction of *tasks* that are infinitely automated, which is hard to observe. Like us, B. Jones (2025) focuses on cost shares, but that paper studies automating the idea production function.

Now divide both sides by initial GDP, Y_t and use equation (25):

$$\begin{aligned}\frac{Y_{cf}}{Y_t} &= \alpha_{\varnothing}^{\frac{\sigma}{\sigma-1}} \frac{Y_{\varnothing t}}{Y_t} \\ &= s_{\varnothing t}^{\frac{\sigma}{\sigma-1}} \\ &= \left(\frac{1}{1-s_{\infty t}}\right)^{\frac{\sigma}{1-\sigma}} \approx 1 + \frac{\sigma}{1-\sigma}\, s_{\infty t} \end{aligned} \tag{26}$$

where the approximation is valid when the factor's cost share is small. When $\sigma = 1/2$, this approximation tells us that the percent gain in output from automating the factor's tasks with infinite productivity is simply equal to the factor's cost share itself, $s_{\infty t}$. When $\sigma < 1/2$, the gain is even smaller. When $\sigma = 3/4$, the gain is $3 \cdot s_{\infty t}$.

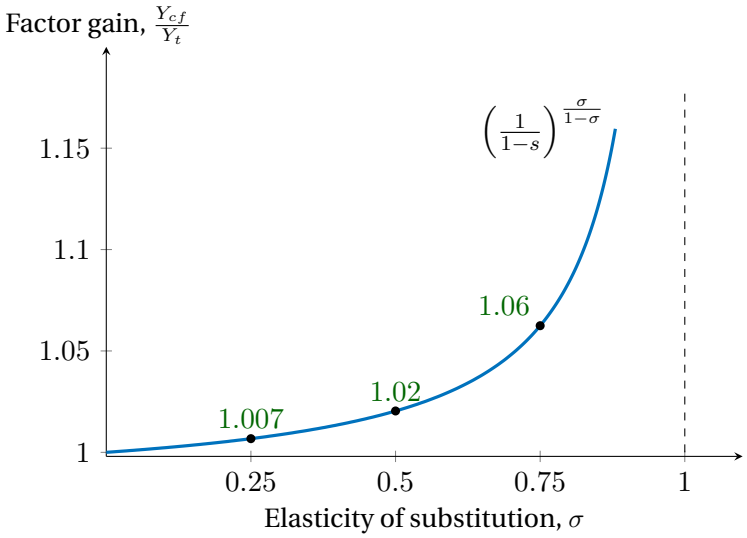
Finally, if $\sigma = 1$ — so that no factor is essential — the gain is infinite. This case underscores the weak-link logic of the model and why $\sigma < 1$ is natural. When $\sigma < 1$, the scarcity of some tasks constrains overall production; when $\sigma \geq 1$, that force disappears, and making even one task infinitely productive is enough to send output to infinity.

Automating software. Given the advances in LLMs at coding, software is generally thought to the one of the first industries that will be largely automated by A.I. The share of software in GDP is around 2%.[9] This means that automating all the tasks that are currently done by software with infinite productivity would only raise GDP by about 2% when $\sigma = 1/2$ and even less for $\sigma = 0.3$. Figure 6 shows how this result changes for different values of σ. The effects are remarkably small.

Automating all cognitive tasks. More speculatively, transformative A.I. is thought to move on to automating all cognitive tasks — anything that could be done by a remote worker with a computer could potentially be done by an A.I. agent. Around two thirds of GDP is paid to labor. We consider what would happen if half of this were fully automated with infinite productivity. With $\sigma = 0.3$, equation (26) gives a gain of $(1/(1-1/3))^{0.3/0.7} \approx 1.18$; that is, infinitely automating 1/3 of GDP would only raise GDP by 18%. At some level, this number seems quite small; after all we have infinite productivity on a third of current GDP. However, the logic is again one of weak links. The economy is constrained by the other two-thirds of tasks that are not automated.

[9] For example, the share of NAICS 511, 516 (Publishing industries, except internet (includes software)) in 2021 was less than 1.7% of nominal GDP.

Figure 6: Automating Software, $s = 2\%$



Note: The figure implements equation (26) to show the proportional gain in GDP from automating with infinite productivity all the tasks currently performed by software, assuming that software accounts for 2% of current GDP.

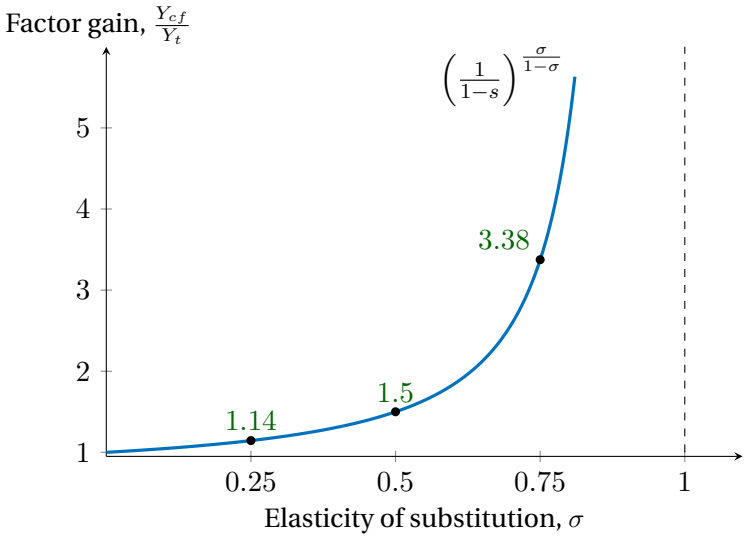
Figure 7 shows how this result changes for different values of σ. In this case, the values start to get large, e.g., when $\sigma = 3/4$.

5.2 Long-Run Growth with Infinite Automation

The static calculations so far freeze the set of tasks that are infinitely automated. We now consider what happens in the long run when this set increases.

Figure 8 shows three possibilities for automation in the long run. In the green line, some fraction of tasks — say 5% — can never be automated. In this case, the infinite automation of the other 95% of tasks removes a large number of weak links, but the economy is still constrained by the 5% of tasks that cannot be automated. This scenario captures the intuition that some tasks seem likely to be performed by people for at least several decades: helping an elderly patient with dementia through a confused night, rewiring the electrical system in a renovated building, running a kindergarten classroom, negotiating a delicate business deal, or playing professional sports. In this case, the production function eventually converges to $Y_t = A_t L_t$ where $\hat{A}_t = \hat{\psi}_{\ell t}$.

Figure 7: Automating All Cognitive Tasks, $s = 1/3$



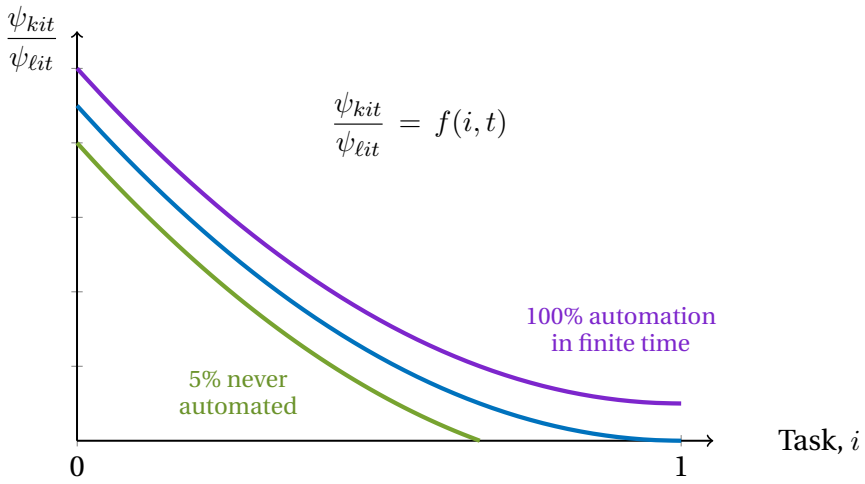
Note: The figure implements equation (26) to show the proportional gain in GDP from automating with infinite productivity all the tasks currently performed by cognitive labor, assuming that currently accounts for 1/3 of GDP.

The infinite automation of 95% of tasks raises output considerably (by $20^{\frac{\sigma}{1-\sigma}}$), but the remaining weak links have two important consequences. First, output remains finite even with the infinite automation, and second, growth eventually slows to the rate at which people get better on the weak links that are never eliminated. The economy ultimately succumbs to the Baumol cost disease.

The purple line provides the other extreme: in this case, 100% of tasks will be automated in finite time (i.e., when the wage rises enough to make capital cheaper on every task). Production converges to $Y_t = B_t K_t$ where $\hat{B}_t = \hat{\psi}_{kt}$. This is an "AK" style model in which the productivity of capital itself grows. So this case results in explosive growth.

Finally, the blue line provides an intermediate case in which $f(1,\cdot) = 0$. That is, the productivity of capital on the last task is always zero. Labor will therefore always be used in at least this task, even as the share of tasks using labor vanishes toward zero over time. It is not obvious what happens in this intermediate case, but we have constructed various examples with a range of outcomes, including both explosive growth and growth that is finite. We will explore this case in more detail in the next section.

Figure 8: Three Types of Automation in the Long Run



Note: The figure shows three possibilities for automation in the long run: (a) In the green line, some fraction of tasks can never be automated, so that $f(\beta,\cdot) = 0$ for $\beta > \bar{\beta}$. (b) In the purple line, $f(1,\cdot) > 0$ so that 100% of tasks are automated in finite time. (c) In the middle blue line, $f(1,\cdot) = 0$ so that there is always some task that is not automated, but the fraction of tasks using labor vanishes to zero asymptotically.

5.3 Dynamics: Automating the Idea Production Function

The previous two subsections show (a) what happens in a static setting in which some tasks are infinitely automated, and (b) what happens to long-run growth depending on whether or not all tasks are eventually automated.

Here we consider the full dynamics of automation and growth when tasks in both the goods production function and the idea production function can be automated. Relative to the model in the first half of the paper, we both enrich the environment in some dimensions and specialize it in others. First, we directly build on the automation model we've already developed, which is useful both directly and for calibrating the parameters. Second, we introduce an idea production function that allows us to endogenize ψ_{kit} and $\psi_{\ell it}$ as well as the automation process itself. Third, the model is a "lab equipment" version of an (either fully- or semi-) endogenous growth model with endogenous automation. Ideas are produced using units of the final good, so that a single automation process incorporates the automation of tasks for producing both goods and ideas. The full model is summarized in Table 6.

Table 6: The Dynamic Model: Automating Goods and Ideas

CES task model	Same as before $\Rightarrow Y_t$ and Ω_{kt}
Idea PF	$\dot{Q}_t = \bar{q} R_t^{\lambda} Q_t^{\phi}$
Resource constraint	$C_t + I_t + R_t = Y_t$
Ideas $\Rightarrow \psi_{kit}$	$\psi_{kit} = Q_t^{\theta_k} f(i)$
Ideas $\Rightarrow \psi_{\ell it}$	$\psi_{\ell it} = Q_t^{\theta_\ell}$ (homogeneous)
Heterogeneity	$f(i) = \frac{(1-i)^{\mu}}{1+\mu_0(1-i)^{\mu}} + \bar{f}$
Capital accumulation	$\dot{K}_t = I_t - \delta K_t$
Population growth	$L_t = L_0 e^{nt}$
Allocations	$R_t = \bar{\iota}_R Y_t$ and $I_t = \bar{\iota}_K Y_t$

Note: The model is an endogenous growth model with endogenous automation. The "lab equipment" structure means that automating the goods production function also automates the production function for ideas. The $f(i)$ function incorporates heterogeneity across tasks in the timing of automation.

In terms of simplifications relative to the model in the first part of the paper, we assume a convenient functional form $f(i)$ for the comparative advantage of capital and labor at different tasks:

$$f(i) = \frac{(1-i)^{\mu}}{1+\mu_0(1-i)^{\mu}} + \bar{f} \tag{27}$$

This functional form permits an "S" shape for $f(i)$. Importantly, the $\bar{f}$ parameter also allows us to control what happens to automation in the long run since $f(1) = \bar{f}$. If $\bar{f} > 0$, then $f(1) > 0$ so that all tasks are automated in finite time. Conversely, if $\bar{f} < 0$, then there is a positive set of tasks that are never automated. Finally, the setup permits the intermediate case of $f(1) = \bar{f} = 0$ so that the fraction of tasks using labor vanishes to zero asymptotically. These correspond to the cases shown earlier in Figure 8.

This functional form is monotonically decreasing, so we get a "single crossing" in the automation condition $\frac{\psi_{kit}}{\psi_{\ell it}} = Q_t^{\theta} f(i) = \frac{r_t}{w_t}$, where $\theta \equiv \theta_k - \theta_\ell$. This means that there is a unique equilibrium β_t such that tasks below β_t use capital and tasks above β_t use labor.

Calibration. We calibrate the model to match many of the facts that we documented in the first part of the paper, as shown in Table 7. Note that because $\psi_{\ell it}$ is homogenous across i, the inequalities in our earlier propositions hold with equality.

The parameters μ_0 and μ capturing the heterogeneity in capital productivity through $f(i)$ are chosen so that the mapping between the automation cutoffs β_t and the capital share, s_{Kt} fits as well as possible for 1950, 1975, 2000, and 2023.[10]

We consider two calibrations that differ in the key parameter θ_k, which governs how fast machines improve as ideas accumulate. The first calibration matches the private business sector, where $\hat{\psi}_{kt} - \hat{\psi}_{\ell t} = 3.6\%$ per year. We view this as "A.I. as a continuation of historical patterns." The second calibration matches the computer sector, where $\hat{\psi}_{kt} - \hat{\psi}_{\ell t} = 9.5\%$ per year. We view this as "A.I. as a break with past" in which automation proceeds at the historically fast pace seen in computers. Both calibrations share the same $\hat{\psi}_{\ell t} = 0.5\%$ per year, chosen via introspection to be consistent with the low values of $\hat{Z}_t + \hat{\psi}_{\ell t}$ in the data.[11]

As anticipated earlier in Figure 8, we consider three values of the fraction of tasks that are never automated: 3%, 0%, and -3%. Our thinking is that we can choose the parameters of the comparative advantage function $f(i)$ to match historical data on β_t and s_{Kt}. However, this does not necessarily tell us what happens in the long run. To be agnostic, we consider these three cases, which permits a "$Y_t = B_t K_t$" explosive growth case, a "$Y_t = A_t L_t$" case in which weak links are a permanent feature of the economy, and a case in between.

A.I. as a continuation of historical patterns (Private Business). We begin with the calibration based on the private business sector. Figure 9a shows the evolution of the capital share s_{Kt} over time for the three cases. As expected, there is a full automation case in which the capital share rises to 100% because all tasks are automated in finite time. Conversely, there is also the "permanent weak links" case in which the capital

[10] In this particular formulation of the dynamic model, there is a one-to-one mapping between the automation cutoffs β_t and the capital share, s_{Kt}:

$$\left(\frac{s_{Kt}}{1-s_{Kt}}\right)^{\frac{1}{1-\sigma}} = \frac{f(\beta_t)}{\xi(\beta_t)}\left(\frac{1}{1-\beta_t}\right)^{\frac{1}{1-\sigma}}$$

where $\xi(\beta_t) = \left(\int_0^{\beta_t} f(i)^{\sigma-1} di\right)^{\frac{1}{\sigma-1}}$.

[11] In simulations, we have $\hat{Z}_t$ starting at the implied historical value but trending to zero slowly over time.

Table 7: Calibration of the Dynamic Model

Moment or Parameter	Value	Source
Moments from first half of the paper		
Capital shares, s_{Kt}	.35, .33, .33, .42	1950, 1975, 2000, 2023
Automation cutoffs, β_t	.33, .55, .72, .81	1950, 1975, 2000, 2023
Task TFP growth, $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$	3.6% or 9.5%	Pvt Business or Computers
Labor task TFP growth, $\hat{\psi}_{\ell t}$	0.5%	Assume humans improve slowly
Chosen to match moments in the data / first half of paper		
ψ_{kit} idea elasticity, θ_k	4.13 or 9.96	Pvt Business or Computers
$\psi_{\ell it}$ idea elasticity, θ_ℓ	0.5	$\hat{\psi}_{\ell t} = 0.5\%$ per year
"Other" TFP growth, $\hat{Z}_t$	-0.6%	From $\hat{A}_t$ and automation rate. Decays to zero
$f(i)$ parameters: μ, μ_0	3.46, 3.93	To match capital share, s_{Kt}, given β_t, 1950 – 2023
Initial ideas, Q_0	1.12 or 1.04	Pvt Business or Computers
Initial "other TFP", Z_0	11.0 or 11.4	Pvt Business or Computers
Initial idea productivity, $\bar{q}\bar{\iota}_R^{\lambda}$	0.0001	To match $\hat{Q} = 1\%$ in 2020
Fraction never automated, $\bar{f}$	3%, 0, -3%	Values for $1 - \bar{\beta}$ chosen to illustrate different possibilities
Chosen from the literature		
Elasticity of substitution, σ	0.3	First half of paper
Idea PF parameters, λ, ϕ	1, -2	BJVW (2020)
Population growth, n	0.01	1% per year
Investment rate, $\bar{\iota}_K$	0.20	20% of GDP
Depreciation rate, δ	0.05	5% per year
Initial capital-output ratio	2.5	$K/Y = \bar{\iota}_K/(g_Y + \delta)$
Initial labor force, L_0	1	Normalization

share falls to zero. Our functional form for $f(i)$ with $\bar{f} = 0$ turns out to deliver a stable capital share around 36%— it is not literally constant along the transition path, but nearly so.

Figure 9b shows the evolution of economic growth over time for the three cases. As anticipated, the full automation case results in explosive growth, and the incomplete automation case results in growth that ultimately falls to a low rate based on $\hat{\psi}_\ell = 0.5\%$. This is the case in which some weak links can never be automated away. We eventually have infinite effective capital on the automated tasks, so production settles down to $Y_t = A_t L_t$ and growth is limited to the rate at which people get better.

The baseline case is the one in which $\bar{f} = 0$ so that the share of tasks using labor vanishes to zero, but only as t goes to infinity. Recall that this case also leads to a stable capital share. The automation rate x_t rises over time and this leads growth to accelerate even though the capital share remains stable.[12]

Another important finding in Figure 9b is that even though the growth paths are eventually very different, the paths are indistinguishable for the next 75 years. A rising automation rate means that growth rates rise over time in all three cases for the next 75 years. The incomplete automation effects only become visible as we approach that constraint in the distant future.

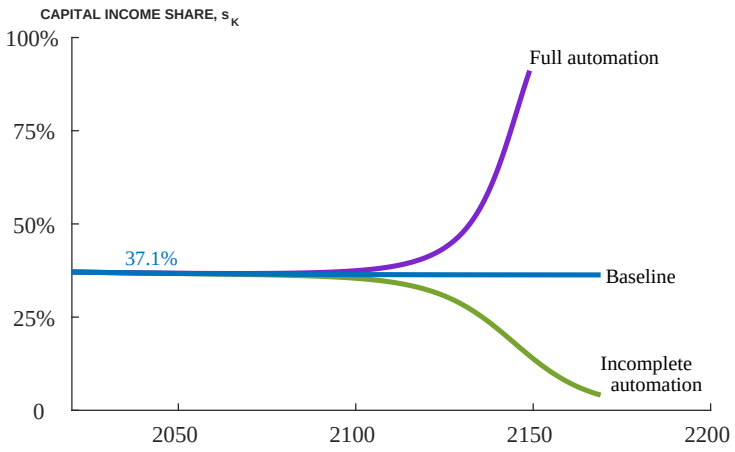
To see how this accelerating economic growth plays out, it is helpful to see the graph of GDP per person, shown in Figure 9c. The red dashed line shows the initial trend line corresponding to constant economic growth. The acceleration in growth is apparent in the rising slope on the logarithmic scale.

Despite the accelerating growth, the effects of A.I. on GDP per person are remarkably small for the next 20 to 40 years. The labels on the dots report the factor gain over the initial trend line. The acceleration is slow because the economy remains constrained by weak links — the relatively few and shrinking set of tasks for which human labor is still essential.

[12]The capital share equation in the preceding footnote can be used to see that the power functions in $f(i)$ lead to a constant capital share. To see how growth can still accelerate, note that the labor share equation is $1 - s_{Kt} = \left(\frac{A_t L_t}{Y_t}\right)^{\frac{\sigma-1}{\sigma}}$ and therefore $y_t = (1 - s_{Kt})^{\frac{\sigma}{1-\sigma}} A_t$ and $\hat{A}_t = \hat{Z}_t + \hat{\psi}_{\ell t} + \frac{1}{1-\sigma} x_t$. A rising automation rate x_t will then raise the growth rate if the capital share is stable.

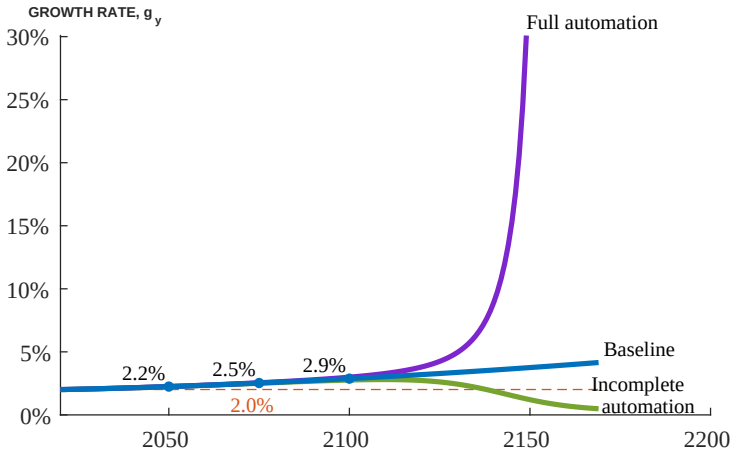
Figure 9: The Future if AI = Private Business

(a) The Capital Share



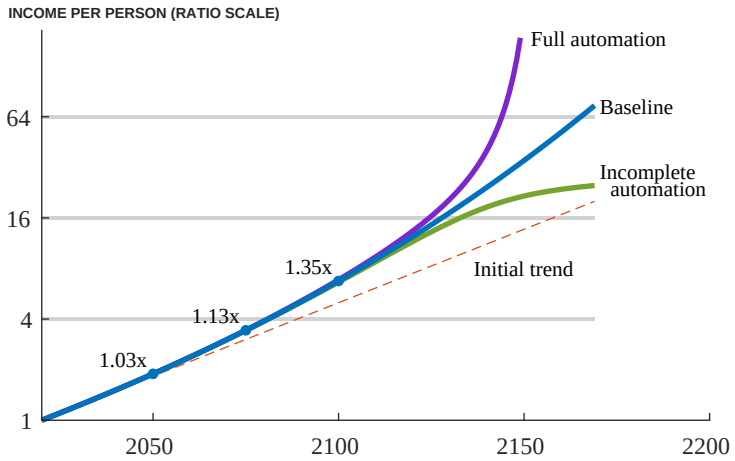
Depending on the ultimate nature of automation, the capital share can rise to 100%, fall to zero, or remain stable at its current value.

(b) Economic Growth



If the capital share reaches 100% in finite time, then growth explodes. If it falls to zero, growth falls to a low rate based on $\hat{\psi}_\ell = 0.5\%$.

(c) GDP per Person



Even though the futures are eventually very different, the paths are indistinguishable for the next 75 years. The labels on the dots report the factor gain over the initial trend line.

A.I. as a break with the past (Computers). We next consider the calibration based on the computer sector, in which the initial automation rate is substantially faster. Figure 10 shows the results.

The intuition for the difference between the two calibrations is straightforward. In the computer sector, machines improve much faster ($\hat{\psi}_{kt} - \hat{\psi}_{\ell t} = 9.5\%$ per year versus 3.6% for the private business sector), which drives faster automation and more rapidly accelerating growth. As we discuss in the next section, the computer calibration features explosive growth ($\Phi > 1$) while the private business calibration does not.

Intuition for slowly accelerating growth. One way to understand the simulation results is to ask: under what parameter conditions does growth explode? The answer depends on the degree of dynamic increasing returns, Φ.

It is helpful to focus on the baseline case in which $\bar{f} = 0$. In this case, the capital share s_{Kt} stabilizes at some value s_K^* rather than going off to 100% or 0%, so this is the case in which one might expect a standard BGP. Appendix C shows this formally: as $\beta_t \to 1$ the capital share settles to $s_K^* = \frac{1}{\mu(1-\sigma)}$. For $\mu = 3.46$ and $\sigma = 0.3$, this gives $s_K^* \approx 0.36$.

The second main result in Appendix C is that the key measure of dynamic increasing returns in this model is Φ, defined as

$$\Phi \equiv \frac{\lambda}{1-\phi}\left(\theta_\ell + \frac{s_K^*}{1-s_K^*}\theta_k\right) \tag{28}$$

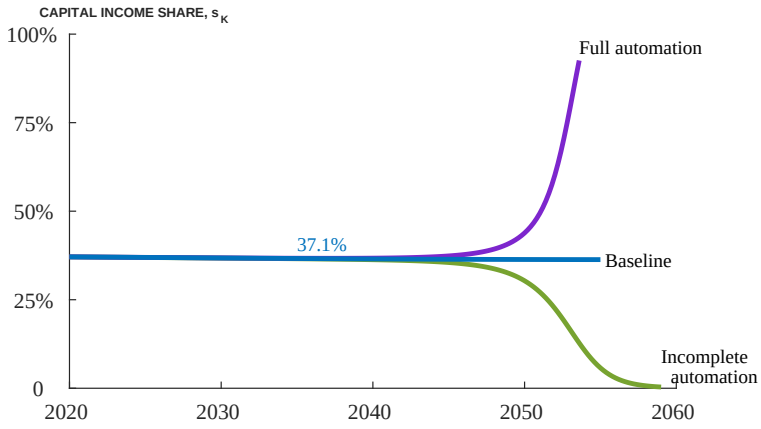
In particular, when $\Phi < 1$, this automation model features a BGP with semi-endogenous growth. The long-run growth rate is given by

$$g_y = \frac{\Phi n}{1-\Phi} \tag{29}$$

This equation has a standard form for a semi-endogenous growth model in which the idea production function uses goods rather than labor as the main input (the so-called "lab equipment" version). The overall degree of dynamic increasing returns, Φ, is itself the product of two terms. The first is $\frac{\lambda}{1-\phi}$, which is the degree of increasing returns in the idea production function, familiar from many SEG models. The second is $\theta_\ell + \frac{s_K^*}{1-s_K^*}\theta_k$, which captures the effect of ideas on $\psi_{\ell it}$ and ψ_{kit}. In addition, the θ_k parameter

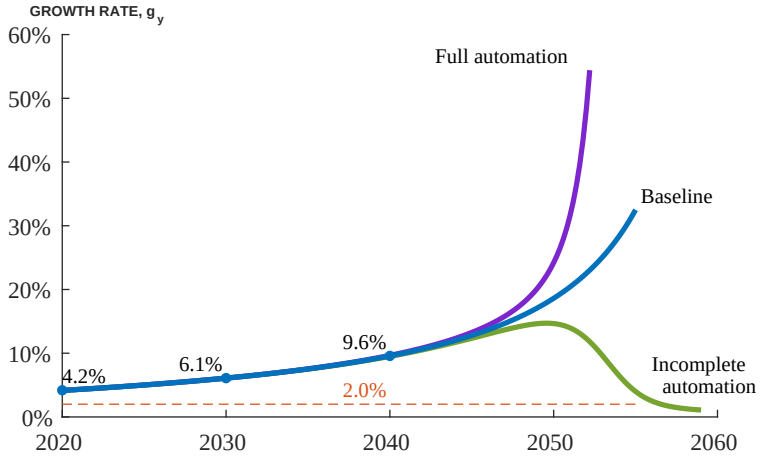
Figure 10: The Future if AI = Computers

(a) The Capital Share



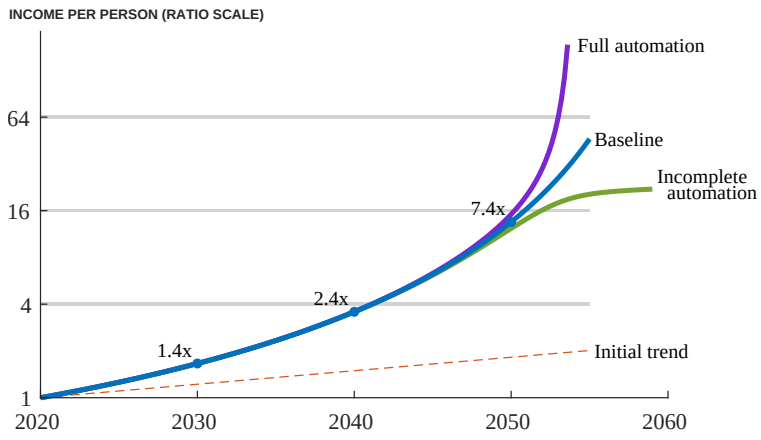
Same three cases as the private business calibration, but with faster automation.

(b) Economic Growth



With the faster automation rate calibrated to the computer sector, growth accelerates much more rapidly.

(c) GDP per Person



The computer-sector calibration leads to infinite income by 2060 unless some essential tasks can only be performed by labor.

gets multiplied by s_K^* reflecting the cost share of the tasks that have been automated and by $1/(1-s_K^*)$ to incorporate the dynamic feedback that comes from $K_t \rightarrow Y_t \rightarrow K_t$. But the bottom line is that when parameter values are such that $\Phi < 1$, this semi-endogenous growth model with automation features a BGP in which the growth rate is proportional to the rate of population growth and where the factor of proportionality is increasing in Φ.

If the degree of dynamic increasing returns is exactly unity ($\Phi = 1$), the semi-endogenous growth turns into fully endogenous growth. This is the knife-edge condition that generates endogenous growth when population is constant ($n = 0$). Of course, with positive population growth, economic growth explodes, as suggested by equation (29).

If $\Phi > 1$, then equation (29) has no positive solution. This is the case "beyond endogenous growth" in which we get explosive growth, even with zero population growth.

Plugging in our calibrated parameter values from Table 7 into equation (28) gives strikingly different values for the two calibrations: $\Phi_{pbs} = 0.95$ for the private business sector and $\Phi_{comp} = 2.06$ for the computer sector. The private business calibration features semi-endogenous growth with a long-run growth rate of $g_y = \frac{0.95 \times 0.01}{1-0.95} = 19\%$ per year. This is a very high growth rate, but it is not explosive — the economy is transiting to this long-run rate from its initial 2% growth. The computer calibration, in contrast, features explosive growth with Φ well above unity.

Speed of transition dynamics. For the computer calibration with $\Phi = 2.06$, further intuition comes from a one-dimensional version of a system that exhibits explosive growth. Consider the differential equation $\dot{X}_t = \bar{g} X_t^{\Phi}$ where $\Phi > 1$. The growth rate of X_t then satisfies $\hat{X}_t = \bar{g} X_t^{\Phi-1}$, so that with $\Phi > 1$, the growth rate is increasing in the level of X_t. Hence the explosion.

This differential equation can be integrated to yield:

$$X_t = \left(\frac{1}{X_0^{1-\Phi} - (\Phi - 1)\bar{g}t} \right)^{\frac{1}{\Phi-1}}$$

This solution has an asymptote where X_t goes to infinity in finite time. Setting $X_0 = 1$ as our initial condition — as we do with y_{2020} in our simulations — the date t_∞ at which

X_t goes to infinity is given by

$$t_\infty = \frac{1}{(\Phi - 1)\bar{g}} \tag{30}$$

For the computer calibration, the initial growth rate is 4.2% and $\Phi = 2.06$, giving

$$t_\infty = \frac{1}{1.06 \times 0.042} = 23 \text{ years}$$

For the private business calibration with $\Phi = 0.95$, growth does not explode but instead transits slowly toward the long-run growth rate of 19%. The weak links — captured by the low $\sigma = 0.3$ — are the key force slowing this transition. The economy starts at 2% growth and takes a very long time to approach 19%.

The role of weak links is somewhat hidden by the way we've written Φ in equation (28). The expression depends on the capital share s_K^* but recall that $s_K^* = \frac{1}{\mu(1-\sigma)}$. The higher is σ, the higher is s_K^*, and as s_K^* approaches one, Φ goes to infinity. A lower σ reduces Φ: weak links tame explosive growth. With $\sigma = 0.3$, the private business calibration features $\Phi < 1$ even though the computer calibration with its faster automation rate gives $\Phi > 1$. The magnitude of weak links is therefore central.

6. Conclusion

How much of past economic growth is due to automation, and what does this imply about the effects of A.I. and automation in the coming decades?

We perform growth accounting using a task-based model for agriculture, motor vehicles, computers, software, and for the aggregate U.S. economy. Historically, TFP growth is largely due to improvements in the productivity with which capital performs tasks. We estimate that the task-specific growth rate of capital productivity exceeds labor task productivity growth by at least 3 percentage points per year across all our sectors. The key benefit of automation is that we switch from using slowly-improving labor to using rapidly-improving capital. Growth is limited by how quickly we strengthen more of our weak links.

Looking to the future, we develop an endogenous growth model in which the production of both goods and ideas gets endogenously automated. We consider two cal-

ibrations based on our historical evidence: one matching the private business sector and one matching the faster-automating computer sector. In both cases, automation leads economic growth to accelerate. A key reason the acceleration is slow is the prominence of weak links. Even when most tasks are automated by rapidly improving capital, output is constrained by the tasks performed by slowly-improving labor.

APPENDIX

A. Instructions for the LLM to measure β_t

We have a two-step process for measuring β_t in each sector. First, we ask the LLM to construct a detailed list of 150 specific tasks that are essential in the sector in the United States over the past century. Second, we ask the LLM to indicate whether each task was automated or not in each year from 1920 to 2020. For partially automated tasks, we ask the LLM to break them into subtasks.

Here is the instruction we gave to the model to create the list of tasks:

```
**INSTRUCTIONS FOR CREATING TASK LIST**
Consider the motor vehicles sector of the U.S. economy for the
   past 75 years since 1950. I am writing an economics research
   paper at the PhD level on automation. Please do the following:
1. Construct a detailed list of 100 specific tasks that are
   essential and important in the motor vehicles industry in the
   United States over that time frame.

**CLARIFICATIONS**
(a) All tasks must in principle be able to be performed by people
   . Over time, some tasks may have been automated, which means
   they are performed without human involvement. But it is
   crucial that the tasks be things that could be performed by
   labor historically. Check carefully to ensure this is the case
   .
(b) On the other hand, do not neglect tasks that are fully
   mechanized now. For example, "Exterior painting" is surely an
   important task historically that is now fully automated.
   Prioritize tasks that are both required and essential and
   economically significant in some way, at least historically.
(c) All tasks should be things that were accomplished in motor
   vehicle production in the year 1950 as well as today. Make
   sure that all tasks were present more than 75 years ago.
(d) Please only use excellent, reputable resources that are
   specific to motor vehicles and empirical in nature. Do not use
```

models or theory papers from macroeconomics or growth economics.
(e) Examples of tasks might include "Engine assembly" or "Tire attachment" or "Windshield installation." Also, "Management of the factory" is one possible high-level task that may have subtasks; we certainly want to consider management as one of the important categories of tasks.

DELIVERABLES
DELIVERABLE 1: Provide a short narrative summary of the results.

DELIVERABLE 2: Provide an Excel file containing the detailed results.
- The first sheet should be called "Overview". It should contain the date, the prompt, and the narrative summary.
- The second sheet should be called "Task Data". Report your task results in the form of a table with the tasks as rows. Some entries that explain the "category" of the task in each row would be helpful, with one column for the high-level category and another column for the detailed task description.
- The third sheet should be called "Sources". Document all sources used in a standard academic reference style. Include hyperlinks.

Here is the instruction we gave to the model to measure β_t:

INSTRUCTIONS FOR AUTOMATION RATES
Consider the motor vehicles sector of the U.S. economy for the past 75 years since 1950. I am writing an economics research paper at the PhD level on automation. I have uploaded a file containing a list of 100 tasks that are essential in motor vehicle production over this time frame. Please do the following:
1. Consider a task (each row of the "Task Data" sheet in the spreadsheet).
2. Consider the year 1957. Indicate whether the tasks was automated or not in 1957; for partially automated, use a

```
   fraction such as 10% or 35% or 85%. Take into account the
   fraction of farms that have automated the task in constructing
    your estimate.
4. Repeat Step 2 for 1967, 1977, 1987, 1997, 2007, and 2017 for
   that task and record your answer in a new column for each time
    period.
5. Repeat Steps 1 to 4 for each task in the spreadsheet.

**CLARIFICATIONS**
(a) Please only use excellent, reputable resources that are
   specific to motor vehicles and empirical in nature. Do not use
    models or theory papers from macroeconomics or growth
   economics.

**DELIVERABLES**
DELIVERABLE 1: Provide a narrative summary of the results. At the
    end of this summary be sure to report what fraction of tasks,
    equally weighted across the 100 tasks (i.e. taking averages
   across all cells), were automated in each year.

DELIVERABLE 2: Provide an Excel file containing the results. You
   should build on the Excel file that has been provided.
- The first sheet should be called "Overview". It should contain
   the date, the prompt, and the narrative summary.
- The second sheet should be called "Task Data". Report your task
    results in the form of a table with the tasks as rows and the
    years as columns. Most of the entries in the table should be
   numbers such as 0%, 10%, 35%, 100%, etc. (in numerical format,
    of course).
- The third sheet should be called "Sources". Document all
   sources used in a standard academic reference style. Include
   hyperlinks.
```

B. Proof of Proposition 2

B.1 Key Weights are Cost Shares

Notice that from the share equations in Proposition 1, the FOC for the representative firm's problem to allocate labor is

$$L_{it} = (\psi_{\ell it} Z_t)^{\sigma-1} \left(\frac{w_t}{P_t}\right)^{-\sigma} Y_t. \tag{31}$$

Integrating this equation over all tasks yields

$$L_t = \int_{\Omega_{\ell t}} L_{it}\, di = \int_{\Omega_{\ell t}} \psi_{\ell it}^{\sigma-1}\, di \cdot Z_t^{\sigma-1} \left(\frac{w_t}{P_t}\right)^{-\sigma} Y_t.$$

Taking ratios of these last two equations yields

$$\frac{L_{it}}{L_t} = \frac{w_t L_{it}}{w_t L_t} = \frac{\psi_{\ell it}^{\sigma-1}}{\int_{\Omega_{\ell t}} \psi_{\ell it}^{\sigma-1}\, di} \equiv \omega_{\ell it}. \tag{32}$$

By a similar argument, the same type of expression holds for capital:

$$\frac{K_{it}}{K_t} = \frac{r_t K_{it}}{r_t K_t} = \frac{\psi_{kit}^{\sigma-1}}{\int_{\Omega_{kt}} \psi_{kit}^{\sigma-1}\, di} \equiv \omega_{kit}. \tag{33}$$

That is, the key weights that will show up in our aggregation of growth rates are equal to the cost shares of the relevant tasks.

To handle the "multiple points of automation" possibility, it is useful to define the average of the weights across all points of automation:

$$\bar{\omega}_{k\beta t} \equiv \sum_{m=1}^{M_t} \omega_{k\beta_m t} \frac{|\dot{\beta}_{mt}|}{\dot{\beta}_t} \quad \text{and} \quad \bar{\omega}_{\ell\beta t} \equiv \sum_{m=1}^{M_t} \omega_{\ell\beta_m t} \frac{|\dot{\beta}_{mt}|}{\dot{\beta}_t} \tag{34}$$

where $\dot{\beta}_t = \sum_m |\dot{\beta}_{mt}|$ is the total flow of automation that occurs across the different automation points.

B.2 Proof

To see a simple version of the proposition, consider the case in which there is only a single point of automation, β_t, at which $\psi_{k\beta t}/\psi_{\ell\beta t} = r_t/w_t$. In that case, the sets are just the intervals $[0, \beta_t]$ and $[\beta_t, 1]$, and the derivatives in the proposition are easy to compute

using Leibniz's rule. In that case, $\bar{\omega}_{k\beta t} = \omega_{k\beta t}$ and $\bar{\omega}_{\ell\beta t} = \omega_{\ell\beta t}$.

ZZZ

B.3 Intuition for Automation and $F(BK, AL)$: Homogeneous ψ's

For intuition, it is helpful to consider an example in which there is almost no heterogeneity in the ψ's. In particular, suppose $\psi_{kit} = \psi_{kt}$ for $i \in [0, \beta_t]$ while $\psi_{kit} = 0$ for $i \in [\beta_t, 1]$. That is, only the tasks up to β_t can use capital. But all tasks can use labor: $\psi_{\ell it} = \psi_{\ell t}$ for all i. Furthermore, suppose $\psi_{kt}/\psi_{\ell t} > r_t/w_t$: if you can use capital then it is profitable to automate.

In this case, the production function in (1) becomes

$$Y_t = Z_t \left(\beta_t \left(\frac{\psi_{kt} K_t}{\beta_t} \right)^{\frac{\sigma-1}{\sigma}} + (1-\beta_t) \left(\frac{\psi_{\ell t} L_t}{1-\beta_t} \right)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \tag{35}$$

Several insights can be gleaned from this special case. First, notice that β_t and $1-\beta_t$ enter the CES reduced-form production function in two ways. Consider the K_t term. The first β_t functions as a share parameter and captures the fact that capital is used in the fraction β_t of tasks. The second way β_t enters is through the K_t/β_t term. In this case, the capital K_t is spread across β_t tasks, so the capital per task is K_t/β_t; that is, capital per task gets smaller as we spread capital over more tasks. The net of these two effects is shown by writing (35) as $Y_t = F(B_t K_t, A_t L_t)$, where B_t collects the first two β_t terms:

$$B_t = Z_t \left(\frac{1}{\beta_t} \right)^{\frac{1}{1-\sigma}} \psi_{kt} \tag{36}$$

Since $\sigma < 1$ so that tasks are complements, an increase in β_t *reduces* B_t. That is, an increase in automation is *capital depleting* rather than capital augmenting. Better computers — a higher ψ_{kt} — are indeed capital augmenting. But when a given amount of capital is spread across a larger number of tasks because of automation, one effect is that this is capital depleting.

This is only one effect because there is a related effect working through A_t. The

same logic reveals that

$$A_t = Z_t \left(\frac{1}{1-\beta_t} \right)^{\frac{1}{1-\sigma}} \psi_{\ell t} \tag{37}$$

In other words, an increase in β_t is *labor augmenting*. The total labor L_t is concentrated on fewer tasks, so labor per task increases.

With homogeneous ψ's, there are two effects of automation that work in different directions. Automation is simultaneously capital depleting and labor augmenting. That is, it is a twist of the production function, a point emphasized by Aghion, Jones, and Jones (2019).

We now return to the general case with heterogeneous ψ's. As we see next, there are then two additional effects from an increase in β_t that need to be considered. In our full model, β_t is not an independent exogenous variable, but rather the set of tasks that are automated changes because ψ_{kit} and $\psi_{\ell it}$ change.

C. Characterizing explosive growth when $f(1) = 0$

The dynamics of the idea-driven growth model when $f(1) = 0$ are interesting. When $f(1) > 0$ the model looks like $Y_t = B_t K_t$ and growth explodes, while when $f(1) < 0$, the model eventually looks like $Y_t = A_t L_t$ and growth slows to the rate at which people improve, $\hat{\psi}_{\ell t}$. But what happens in between?

As was clear in the graphs and as we show at the end of this section, the capital share s_{Kt} stabilizes at some value s_K^* rather than going off to 100% or 0%. (In fact, as we show at the end of this section, $s_K^* = \frac{1}{\mu(1-\sigma)}$.) So the question is: how can growth explode when the capital share is constant, and what are the conditions under which that occurs?

C.1 The conditions for semi-endogenous growth

The easiest way to see the answer to these questions is to characterize the condition on parameter values such that the model exhibits a BGP with semi-endogenous growth. Then, if the degree of increasing returns is even larger, then growth will explode. We now develop this characterization.

Step 1. First, we study the growth rate of y_t. From the basic labor share equation for CES, $Y = s_L^{\frac{\sigma}{\sigma-1}} AL$, and since factor shares are constant, $g_y = g_A$. Also, $\hat{A} = \theta_\ell \hat{Q} + \frac{1}{1-\sigma} x$, which implies

$$g_y = \theta_\ell g_Q + \frac{1}{1-\sigma} x$$

Step 2. Now we need x. Recall that $f(\beta) = \frac{r}{w} Q^{-\theta}$ where $\theta \equiv \theta_k - \theta_\ell$. Focus on the case in which $f(i) \equiv (1-i)^\mu$ (noting that as i gets close to 1, this is valid even for our richer specification with $\mu_0 \neq 0$). In addition, from the factor share equations, $\frac{w}{r} = \frac{s_L}{s_K} \frac{K}{L}$. Putting all this together and taking logs and derivatives with constant factor shares gives

$$\mu g_{1-\beta} = -\theta g_Q - g_k$$

where $k \equiv K/L$. The automation rate is $x = -g_{1-\beta}$. Also, along a BGP, $g_k = g_y$. Therefore,

$$x = \frac{1}{\mu}(\theta g_Q + g_y)$$

Step 3. Combining Steps 1 and 2 gives

$$g_y = \theta_\ell g_Q + \frac{1}{\mu(1-\sigma)}(\theta g_Q + g_y)$$

It is now convenient to use a result shown at the end of this section and already anticipated above: $s_K^* = \frac{1}{\mu(1-\sigma)}$. Making this substitution, recalling $\theta \equiv \theta_k - \theta_\ell$, and rewriting the previous equation gives

$$g_y = \left(\theta_\ell + \frac{s_K^*}{1-s_K^*}\theta_k\right) g_Q \tag{38}$$

Step 4. Now we need a separate equation for g_Q. From the idea production function, $g_Q \propto \frac{Y_t^\lambda}{Q_t^{1-\phi}}$. Along a BGP,

$$g_Q = \frac{\lambda}{1-\phi} g_Y = \frac{\lambda}{1-\phi}(g_y + n)$$

Combining these last two equations gives the expression for the semi-endogenous growth rate:

$$g_y = \frac{\Phi n}{1-\Phi} \quad \text{where} \quad \Phi \equiv \frac{\lambda}{1-\phi}\left(\theta_\ell + \frac{s_K^*}{1-s_K^*}\theta_k\right) \tag{39}$$

C.2 The Capital Share when $f(1) = 0$

In the dynamic model, the capital share satisfies

$$\frac{s_{Kt}}{1-s_{Kt}} = \left(\frac{f(\beta_t)}{\xi(\beta_t)}\right)^{1-\sigma} \cdot \frac{1}{1-\beta_t}$$

where $\xi(\beta_t) = \left(\int_0^{\beta_t} f(i)^{\sigma-1} di\right)^{\frac{1}{\sigma-1}}$.

Consider the basic functional form $f(i) = (1-i)^\mu$ which implies

$$\xi(\beta_t)^{\sigma-1} = \frac{1}{\mu(1-\sigma)-1}\left[\left(\frac{1}{1-\beta_t}\right)^{\mu(1-\sigma)-1} - 1\right]$$

Combining these equations gives

$$\begin{aligned}
\frac{s_{Kt}}{1-s_{Kt}} &= \frac{1}{\mu(1-\sigma)-1}\left(\frac{1}{1-\beta_t}\right)^{1-\mu(1-\sigma)}\left[\left(\frac{1}{1-\beta_t}\right)^{\mu(1-\sigma)-1} - 1\right] \\
&= \frac{1}{\mu(1-\sigma)-1}\left[1 - \left(\frac{1}{1-\beta_t}\right)^{1-\mu(1-\sigma)}\right] \\
&\to \frac{1}{\mu(1-\sigma)-1} \quad \text{as } \beta_t \to 1 \text{ if } \mu(1-\sigma) > 1
\end{aligned}$$

which in turn implies that $s_{Kt} \to \frac{1}{\mu(1-\sigma)}$ as $\beta_t \to 1$. The condition $\mu(1-\sigma) > 1$ also implies the capital share settles down to an interior point between 0 and 1.

The functional form we use in the dynamic model is slightly richer, i.e., $f(i) = \frac{(1-i)^\mu}{1+\mu_0(1-i)^\mu}$. But the two functional forms are asymptotically equivalent as $i \to 1$, so that the result holds with the richer functional form as well.

D. Deriving the Counterfactual in Proposition 5

References

Acemoglu, Daron, "The Simple Macroeconomics of AI," Working Paper 32487, National Bureau of Economic Research May 2024.

__ 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.

__ and __ , "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 2020, *128* (6), 2188–2244.

__ and __ , "Tasks, Automation, and the Rise in U.S. Wage Inequality," *Econometrica*, 2022, *90* (5), 1973–2016.

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

__ , 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, 2019, pp. 237–282.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb, University of Chicago Press, 2019.

Aschenbrenner, Leopold, "Situational Awareness: The Decade Ahead," June 2024.

Brynjolfsson, Erik, Anton Korinek, and Ajay K. Agrawal, "A Research Agenda for the Economics of Transformative AI," Working Paper 34256, National Bureau of Economic Research September 2025.

Bubeck, Sebastien, Christian Coester, Ronen Eldan, Timothy Gowers, Yin Tat Lee, Alexandru Lupsasca, Mehtaab Sawhney, Robert Scherrer, Mark Sellke, Brian K. Spears, Derya Unutmaz, Kevin Weil, Steven Yin, and Nikita Zhivotovskiy, "Early science acceleration experiments with GPT-5," 2025.

Bureau of Labor Statistics, "Annual capital details: information processing equipment for major sectors and major industries," December 4 2024.

Caunedo, Julieta and Elisa Keller, "Capital Embodied Structural Change," May 2024. unpublished manuscript.

__ , David Jaume, and Elisa Keller, "Occupational Exposure to Capital-Embodied Technical Change," *American Economic Review*, June 2023, *113* (6), 1642–85.

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

__ , Basil Halperin, Thomas Houlden, and Anton Korinek, "Is Automating AI Research Enough for a Growth Explosion?," 2025. unpublished manuscript.

Epoch AI, "GATE — AI and Automation Scenario Explorer," 2025.

Erdil, Ege and Tamay Besiroglu, "Explosive growth from AI automation: A review of the arguments," 2023.

Farboodi, Maryam, Andrew Koh, and Bryant Xia, "Data-Driven Automation," December 2025. unpublished manuscript.

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

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell, "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, June 1997, *87* (3), 342–362.

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

Herrendorf, Berthold, Richard Rogerson, and Ákos Valentinyi, "Structural Change in Investment and Consumption—A Unified Analysis," *The Review of Economic Studies*, 03 2020, *88* (3), 1311–1346.

Jones, Benjamin, "Artificial Intelligence in Research and Development," NBER Working Paper 34312 October 2025.

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

Jones, Charles I., "Intermediate Goods and Weak Links in the Theory of Economic Development," *American Economic Journal: Macroeconomics*, April 2011, *3* (2), 1–28.

Korinek, Anton and Donghyun Suh, "Scenarios for the Transition to AGI," Working Paper 32255, National Bureau of Economic Research March 2024.

Kremer, Michael, "The O-Ring Theory of Economic Development," *Quarterly Journal of Economics*, August 1993, *108* (4), 551–576.

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.

Moravec, Hans, *Mind Children: The Future of Robot and Human Intelligence*, Harvard University Press, 1988.

Oberfield, Ezra and Devesh Raval, "Micro Data and Macro Technology," *Econometrica*, 2021, *89* (2), 703–732.

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

Trammell, Philip and Anton Korinek, "Economic Growth under Transformative AI," 2020. GPI Working Paper No. 8-2020.

U.S. Bureau of Labor Statistics, "Multifactor Productivity," Technical Report 2025. https://www.bls.gov/mfp/.

Wang, Sun Ling, Eric Njuki, Roberto Mosheim, and Richard Nehring, "Agricultural Productivity in the United States," November 2024. US Department of Agriculture, Economic Research Service, https://www.ers.usda.gov/data-products/agricultural-productivity-in-the-united-states.

Young, Alwyn, "Cointegrated Factor Augmenting Productivity," February 2025. unpublished manuscript.

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