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

May 2026 — Version 0.5

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 approximately 4 percentage points per year faster than labor productivity in the U.S. private business sector, and substantially more in some industries. 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 Julieta Caunedo, Sebastian Di Tella, Chang Hsieh, Erik Hurst, Pete Klenow, Danial Lashkari, Ben Moll, Michael Peters, Pascual Restrepo, Rob Shimer, Jón Steinson, Phil Trammell, and seminar participants at the CEPR Macro/Growth meeting, Chicago, 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 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 (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 substitutes. 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 require one heroic measurement: what is the share of labor costs for the tasks that get automated in each year? We measure measure this share via a set of API queries to frontier LLMs (OpenAI's GPT-5.5 and Anthropic's Claude Opus 4.7).

With our accounting framework, we derive several results:

1. If the automation process is continuous, task production switches from using labor to using capital when the costs are equal. Thus, the switching itself generates no productivity growth.

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

3. Historically, 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 — equal to -0.2% per year for the private business sector since 1950. In contrast, the average excess rate at which machines are getting better, $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$, is large — equal to 3.8% per year for the private business sector.

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, essentially all 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 automation. 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 cases: "A.I. as a continuation of historical patterns" and "A.I. as a break with the past." The first case is calibrated to the private business sector 1950–2023. The second case is modeled as "Moore's Law everywhere," by calibrating the aggregate economy to patterns from the computer sector, where automation was historically fast. Simulating the endogenous automation model forward in time we find:

1. *Growth accelerates due to automation.* Despite the stability of past economic growth, future growth accelerates as the automation process endogenously speeds up. There is a flywheel effect: automation leads to more new ideas which in turn leads to more automation.

2. *The growth explosion is slow due to weak links.* When A.I. is a continuation of broad historical patterns, growth rates reach only 2.6% by 2075. In our other calibration of "Moore's Law everywhere," the acceleration is substantially faster, and growth explodes — income becomes infinite in finite time. Nevertheless, even in this extreme calibration the explosion is slow: infinite income does not occur until around 2060. The flywheel effect accelerates economic growth, but growth is constrained by weak links. It is only when the last weak links are automated away that the explosion fully unfolds.

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. We suggest spending half your time on these sections, and then spending 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 measure 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 Caunedo and Keller to measure improvements to capital using the relative price of investment. Their main finding is that capital-embodied technical change 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 of how much 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. The payoff is that we provide a detailed accounting of the sources of TFP growth that informs our understanding of the future consequences of automation.

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. Lashkari, Li, Qiu, and Thompson (2026) also use a task-based model to quantify the future impact of automation. They use AI scaling laws to determine the fixed and variable costs of performing vision-related tasks with computers and forecast future automation and output given a path for the falling cost of compute.

2. Framework

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

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

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

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

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

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 shortly. 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]

[1] Nested CES specifications also permit this richness (Krusell, Ohanian, Ríos-Rull, and Violante, 2000).

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] Many of the most important conceptual insights of the paper are a direct result of the weak links production structure.

Aghion, Jones, and Jones (2019) referred to this feature as "bottlenecks," but we find the "weak links" interpretation to be more appropriate. "Bottlenecks" connotes a few specific impediments whereas every task is potentially a weak link. Put differently, in our CES specification with $\sigma < 1$, fixing a handful of bottlenecks may have small effects on output as the remaining weak links continue to limit production.

New Tasks? Our model features a fixed measure of tasks, but there are 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 task CES. Then new varieties 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

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

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 much capital and labor to hire and how to allocate those amounts across tasks in order to maximize profits, taking P_t, w_t, and r_t as given:[3]

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

subject to (1) and (2).

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 is reversed. Production uses the factor with the highest productivity per dollar cost. We therefore define the set of tasks using capital and labor as

$$\begin{aligned}
\Omega_{kt} &:= \{i \in [0,1] \mid \psi_{kit}/\psi_{\ell it} \geq r_t/w_t\}, \\
\Omega_{lt} &:= \{i \in [0,1] \mid \psi_{kit}/\psi_{\ell it} < r_t/w_t\}, \\
\beta_t &:= \|\Omega_{kt}\| \ \text{ and } \ 1-\beta_t := \|\Omega_{\ell t}\|,
\end{aligned}$$

where β_t is the measure of tasks that are produced with capital. Notice that any task j that is just at the margin of being automated — that is, a task at the boundary of the two sets — satisfies the automation condition

$$\frac{\psi_{kjt}}{\psi_{\ell jt}} = \frac{r_t}{w_t}. \tag{7}$$

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 $[0,1]$. In the special case in which $\frac{\psi_{kit}}{r_t} = \frac{\psi_{\ell it}}{w_t}$ occurs for only a

[3] We also apply this setup to the aggregate economy. In that case, equations (3) and (4) become resource constraints instead of definitions, and these constraints determine r_t and w_t in equilibrium, as usual. Everything else in this section still goes through: we are characterizing TFP growth for the competitive equilibrium.

Figure 1: Automation and Comparative Advantage: Examples

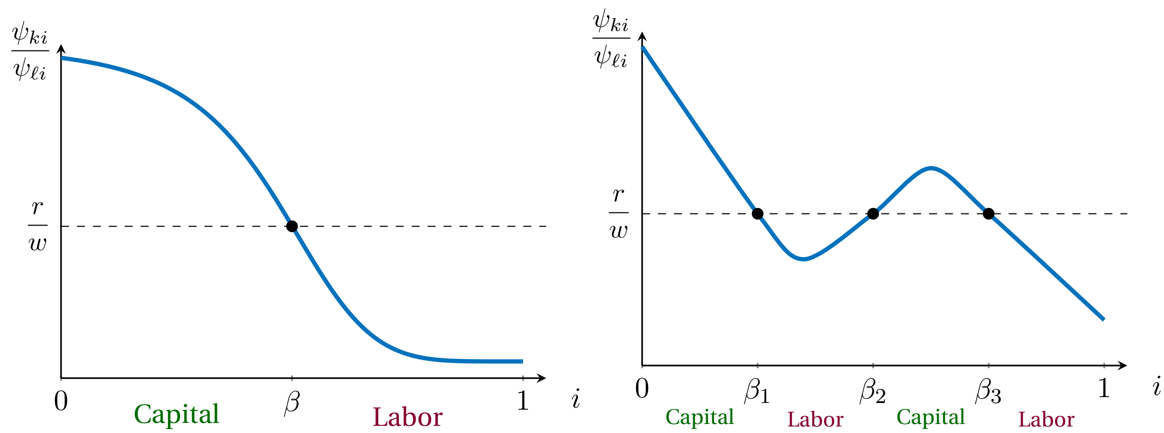

(a) Single automation point

(b) Multiple automation points

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 for 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.[4] 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 A CES-Like Production Function 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

[4]For any single point in time, we could always assign tasks to an index value such that there is a single crossing, but this would not ensure a single crossing for all times due to heterogeneous productivity growth.

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

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

Factor shares are

$$
\begin{aligned}
s_{Kt} &\equiv \frac{r_t K_t}{P_t Y_t} = \left(\frac{B_t K_t}{Y_t} \right)^{\frac{\sigma-1}{\sigma}} \quad \text{and} \quad s_{Kit} \equiv \frac{r_t K_{it}}{P_t Y_t} = \left(\frac{\psi_{kit} Z_t K_{it}}{Y_t} \right)^{\frac{\sigma-1}{\sigma}} \\
s_{Lt} &\equiv \frac{w_t L_t}{P_t Y_t} = \left(\frac{A_t L_t}{Y_t} \right)^{\frac{\sigma-1}{\sigma}} \quad \text{and} \quad s_{Lit} \equiv \frac{w_t L_{it}}{P_t Y_t} = \left(\frac{\psi_{\ell it} Z_t L_{it}}{Y_t} \right)^{\frac{\sigma-1}{\sigma}}.
\end{aligned}
$$

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.[5] 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 perspective 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. But 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,

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

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

[5] 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 $\omega_{\ell it} \equiv \frac{L_{it}}{L_t}$ is task i's labor cost share, $\omega_{kit} \equiv \frac{K_{it}}{K_t}$ 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 — weakening the links and therefore decreasing productivity when $\sigma < 1$ (Aghion et al., 2019).

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. A given amount of labor is being concentrated onto fewer tasks, so labor per task rises and weak links are strengthened, raising productivity.

Following standard logic, TFP growth is

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

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

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

Note: For the private business sector and agriculture, we use data for 1950, 1975, 2000, and 2021/3. For the BEA sectors, we use data at 10-year intervals.

lists the sources of the data for various sectors, chosen via introspection to be interesting.

The Elasticity of Substitution, σ. The elasticity of substitution across tasks, σ, is a key parameter. For our baseline, we assume $\sigma = 0.2$ and explore robustness to this choice. Two comments are warranted.

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 the 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) finds a mean estimate of 0.3. Recent papers support this view. Oberfield and Raval (2021) estimate values between 0.5 and 0.7, while Young (2025) finds estimates of 0.4-0.5. These are estimates of the elasticity of substitution between capital and labor which is an upper bound on σ. We therefore choose $\sigma = 0.2$, smaller than all these estimates. We will consider other values for robustness.

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. There

Figure 2: The Share of Factor Income Paid to Computers

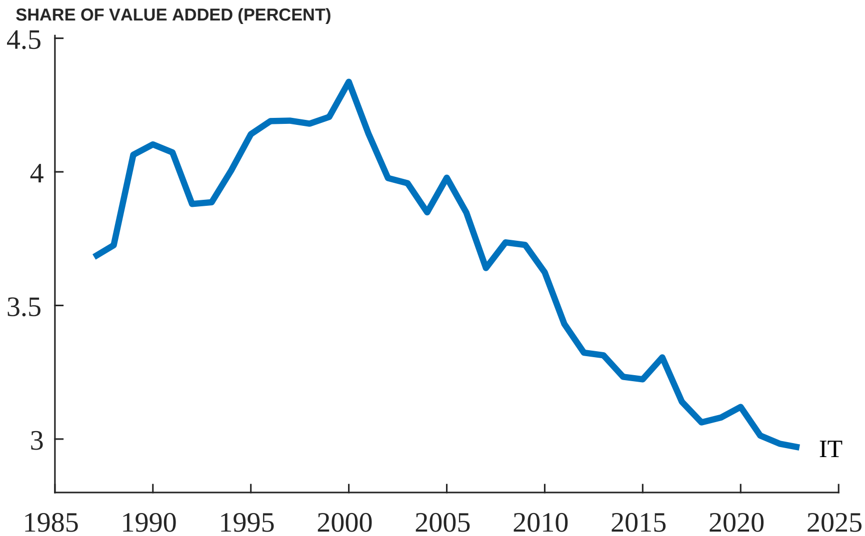

Note: The factor income share of information technology in the private business sector has declined over the past 25 years. *Source:* Bureau of Labor Statistics (2024).

are 100 million times more transistors on a computer chip today than 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.

Our model explains how the factor share can either rise or fall over time, reminiscent of a key result in B. Jones and X. Liu (2024). An increase in ψ_{kit} lowers the factor share (see Proposition 1). But an increase in β_t, the share of tasks that are automated, lowers B_t and hence raises the factor share. In the 1990s, the application of computers to more and more tasks may have raised its factor share while after 2000, the broad-based improvements in computers on all the tasks already using information technology could have lowered its factor share.

Measurement. 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. These measures are already quality adjusted in the sense of Jorgenson-Griliches embodied change.[6] Thus, our measures of productivity growth are the disembodied residuals after accounting for such quality-adjusted effective quantity growth in factors.

To compute aggregate factor-augmenting productivities, we assume there is perfect competition in markets so that factor payment shares equal their production function elasticities. With our reduced-form CES production function, this means that

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

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 from the TFP growth accounting equation (10). The calculation of the factor augmenting growth rates of A_t and B_t is therefore entirely straightforward.

Growth in TFP, A_t, and B_t. Table 2 reports, for a variety of sectors, average factor shares and the growth in total factor productivity, A_t, and B_t.

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, a conventional number. However, the growth rate of B_t, the capital augmenting component, was 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: $1.2 = .35 \cdot (-1.0) + .65 \cdot 2.3$. To the extent that factor shares are stable, this

[6] For labor, this includes adjusting for changes in the workforce's age, sex, and education composition and the associated conditional wages. For capital, this includes a combination of hedonic deflators and user-cost calculations of capital services.

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	4.3	0.57	0.43
Computers	12.8	9.8	14.9	0.41	0.59
Motor vehicles	1.7	0.4	2.7	0.43	0.57
Retail trade	1.7	−1.3	2.4	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.

calculation is invariant to the elasticity of substitution, σ. In particular, it is not driven by our calibration of $\sigma = 0.2$.

For agriculture, computers, and motor vehicles, the growth rate of B_t is positive, while for retail trade and software, 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 therefore be a sign of automation. The positive growth in B_t for agriculture and computers can either 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_{kit} growth concentrated on already-automated tasks so that the share of tasks that are newly automated is modest.

3. Automation 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 first decomposition:

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

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 Immediate Automation Effect. 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, using the definition of the weights in Proposition 2, at any point of automation, β, the automation effect depends on

$$\begin{aligned} s_{Lt}\omega_{\ell\beta t} - s_{Kt}\omega_{k\beta t} &= \frac{w_t L_t}{P_t Y_t}\cdot\frac{w_t L_{\beta t}}{w_t L_t} - \frac{r_t K_t}{P_t Y_t}\cdot\frac{r_t K_{\beta t}}{r_t K_t}, \\ &= \frac{w_t L_{\beta t}}{P_t Y_t} - \frac{r_t K_{\beta t}}{P_t Y_t} = 0. \end{aligned}$$

The automation term depends on the cost of doing the marginally-automated task with labor versus capital. Initially it is cheaper to do a task with labor. But as machines get better rapidly, it eventually becomes cheaper to do the task with capital. Automation occurs precisely when these costs are the same. Hence, the automation term is zero.

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

$$\begin{aligned} \widehat{TFP}_t &= s_{Kt}\hat{B}_t + s_{Lt}\hat{A}_t \\ &= \underbrace{s_{Kt}\hat{\psi}_{kt}}_{\text{Capital improvements}} + \underbrace{s_{Lt}\hat{\psi}_{\ell t}}_{\text{Labor improvements}} + \underbrace{\hat{Z}_t}_{\text{Other TFP growth}} \end{aligned} \tag{14}$$

Discussion. This proposition states that the direct switching effect of automation contributes nothing to TFP growth. So one of our key answers to the question "How much of growth is due to automation?" is zero! Below, we provide an alternative interpretation of the question that leads to a different answer, but this result makes clear that the answer depends on precisely what we mean by the contribution of 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 their 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 machines is 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 machines" is 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). In Section 4, we draw on this perspective to provide one of our primary measures of the effect of automation on growth.

3.1 The Effect of Automation on TFP Growth

Given that the automation "switching 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 result, it is helpful to substitute $s_{Lt} = 1 - s_{Kt}$ into the TFP decomposition in Proposition 3 to get

$$\widehat{TFP}_t = \underbrace{\hat{Z}_t + \hat{\psi}_{\ell t}}_{\text{Baseline TFP growth}} + \underbrace{s_{Kt}\,(\hat{\psi}_{kt} - \hat{\psi}_{\ell t})}_{\text{Automation boost from machines getting better}} \tag{15}$$

Equation (15) 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 the tasks using capital. Intuitively, the proper way to account for this is to multiply the boost by the cost share of tasks using capital — the factor income share s_{Kt}.

In the remainder of this section we show that the "Baseline TFP growth" term in equation (15) is small in most sectors. Therefore the majority of TFP growth is due to the "Automation boost" that comes from machines improving faster than humans.

3.2 Identification

To implement equation (15) 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. The key to identifying the remaining variables is the last term in equation (9), repeated here:

$$\hat{A}_t = \hat{Z}_t + \hat{\psi}_{\ell t} + \frac{1}{1-\sigma}\bar{\omega}_{\ell\beta t}\dot{\beta}_t. \tag{16}$$

Definition 1. *The automation rate, x_t, is defined as the share of total labor costs accounted for by newly-automated tasks before they were automated:*

$$x_t := \bar{\omega}_{\ell\beta t}\dot{\beta}_t = \sum_{m=1}^{M_t} \omega_{\ell\beta_m t}\dot{\beta}_{mt}. \tag{17}$$

To understand the automation rate x_t, begin with the case in which $M_t = 1$ so that there is only a single point of crossing, as in Figure 1a. The weights are employment shares: $\omega_{\ell\beta t} = \frac{L_{\beta t}}{L_t}$ (discussed more in Section B.1). That is, this weight captures the fraction of labor, and thus labor costs, used in the newly-automated tasks. The quantity of tasks being automated is given by the flow term $\dot{\beta}$. Thus, $\omega_{\ell\beta t}\dot{\beta}_t$ is the cost share per task being automated times the quantity of tasks being automated, which equals the total cost share of tasks being automated. With $M_t > 1$ there are multiple points of crossing, so the automation rate simply sums up across all the newly-automated tasks.

Plugging this definition into equation (16) and using the expression for TFP growth in (15) gives our next key result:

Proposition 4 (*Identifying* $\hat{\psi}_{kt}$, $\hat{\psi}_{\ell t}$, *and* $\hat{Z}_t$)**.** If we measure the automation rate x_t, then $\hat{Z}_t + \hat{\psi}_{\ell t}$ and $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ are identified:

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

$$\hat{\psi}_{kt} - \hat{\psi}_{\ell t} = \frac{1}{s_{Kt}} \left(\widehat{TFP}_t - (\hat{Z}_t + \hat{\psi}_{\ell t}) \right).$$

Proposition 4 shows that to identify the relative contribution of baseline productivity growth, $\hat{Z}_t + \hat{\psi}_{\ell t}$, and the automation boost from rapidly improving machines, $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$, we need to measure the automation rate, x_t.

3.3 Measuring the Automation Rate, x_t

This brings us to a crucial part of our analysis. We considered hiring a team of RAs 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 RAs' methods uniform across people, time, and tasks. It then occurred to us that this job is well-suited for state-of-the-art LLMs. At a high level, here is how we proceed, where the unit of analysis is a sector:

Task Categories. We cannot hope to measure a continuum of tasks. Instead, we partition the unit interval into 100 discrete task categories J_j, indexed by j.[7] We define the automation rate at the task category level analogously to the overall automation rate in (17):

$$x_t^j := \sum_{m \in J_j} \omega^j_{\ell \beta_m t} \dot{\beta}_{mt}, \tag{18}$$

where $\omega^j_{\ell\beta_m t} \equiv \frac{L_{\beta_m t}}{L_t^j}$ is the share of labor costs in category j used on newly automated task β_m (once again allowing for multiple points of automation). Then, the overall automation rate in (17) is simply the weighted sum of category-specific automation

[7]The fact that we find the automation rate to be positive for multiple task categories each period highlights the importance of allowing for multiple points of automation in the theory (Figure 1).

rates, where the weights, $\tilde{\omega}^j_{\ell t}$, are the shares of total labor costs accounted for by category j in each period:

$$x_t = \sum_j \tilde{\omega}^j_{\ell t} x^j_t. \tag{19}$$

We proceed to measure $\tilde{\omega}^j_{\ell t}$ and x^j_t to compute x_t using (18) and (19).

The Algorithm to Measure x_t using LLMs.

1. First, we ask the LLM to generate the list of task categories, such that each category is necessary for production. This is a one-time call to the LLM for each sector.

2. For each category and for each period, we then ask the LLM to estimate the fraction of tasks that were automated (i.e. performed by capital instead of by humans). We ask this for the final year, the initial year, and then for every period in between (decadal for the BEA sectors or 25 years for agriculture and private business). This gives us a measure of the fraction of tasks in category j that were automated, β^j_t.

3. We make a separate call to the LLM to measure $\tilde{\omega}^j_{\ell t}$, the share of total labor costs accounted for by category j in each year.

4. We repeat steps 2 and 3 a large number of times (currently 25) to average out the inherent randomness of the LLM's responses.

5. Finally, we compute the automation rate x_t from these β^j_t and $\tilde{\omega}^j_{\ell t}$. To implement this final step, we assume that labor cost shares within a task-category are equal, which implies that $\omega^j_{\ell \beta_m t} \equiv \frac{L_{\beta_m t}}{L^j_t} = \frac{1}{1-\beta^j_t}$. This, together with equations (18) and (19), delivers x_t.[8] In practice, we use the Tornqvist approximation to the continuous-time automation rate, as is common in the growth accounting literature.

Our methods are discussed in detail in Appendix A.

To provide some insight into the data underlying the automation rate, Table 3 shows the least and most automated task categories for the private business sector and for

[8]Across sectors, the automation rates x_t are quantitatively similar when computed using task-category labor cost shares $\tilde{\omega}^j_{\ell t}$ or uniform weights. This suggests that our within-category uniform assumption may not be quantitatively consequential.

Table 3: Least and Most Automated Task Categories

Least Automated Tasks	β	Most Automated Tasks	β
Private Business Sector, 2023			
Graft or bud fruit trees	0.159	Thresh harvested grain	0.993
Classroom instruction	0.207	Harvest small grains (reap/bind)	0.991
Administer vaccinations	0.219	Execute interbank funds transfer	0.989
Assist at surgery	0.220	Grain milling	0.989
Basic wound care and dressing	0.220	Textile spinning	0.989
Motor Vehicles, 2017			
Route and secure cables/hoses	0.174	Cylinder bore machining/honing	0.981
Install exterior trim and emblems	0.187	Phosphate conversion coating	0.979
Install headliner and pillar trims	0.195	Cylinder head machining	0.977
Lay carpets and sound-deadening mats	0.197	Transmission gear hobbing & heat treat	0.976
Install door hardware and locks	0.213	Body cleaning and degreasing	0.976

Note: β is the fraction of tasks in each task category that have been automated, pooled across 50 draws (25 from GPT-5.5 and 25 from Claude Opus 4.7).

motor vehicles in the final year of the sample. Among the least automated tasks are tasks related to grafting fruit trees, classroom instruction, administering vaccinations, and routing cables and hoses during vehicle assembly. Among the most automated tasks are threshing and harvesting grain, grain milling, executing interbank funds transfers, and cylinder bore machining in motor vehicle production.

Table 4 shows the automation rates by sector for our two main LLM models and for a pooled average. In the private business sector, the automation rate averages 1.84% per year for OpenAI's GPT-5.5 and 2.15% per year for Claude Opus 4.7. The numbers in parentheses are the standard deviations of the automation rate across 25 draws. There is generally more variation across models than within models. Looking across sectors, the automation rate is highest in computers — around 5% per year — and lowest in the private business sector, with retail trade and motor vehicles falling in between. See Table A.1 for results across all seven LLMs we queried. We decided to go with the average of the two state-of-the-art models, GPT-5.5 and Claude Opus 4.7, as our benchmark.

Figure 3 shows the automation rate x_t over time for the various sectors. The automation rate for the private business sector has risen over time, while the rate for agriculture

Table 4: Automation Rates by Sector and Model (percent per year)

Sector	GPT-5.5	Claude Opus 4.7	Average
Private business	1.84 (0.17)	2.15 (0.07)	2.00
Agriculture	2.30 (0.29)	3.18 (0.26)	2.74
Computers	4.71 (0.78)	5.63 (0.32)	5.17
Motor vehicles	3.02 (0.49)	4.39 (0.15)	3.70
Retail trade	4.16 (0.71)	4.75 (0.35)	4.45
Software	3.41 (0.44)	3.97 (0.31)	3.69

Note: Each cell shows the mean automation rate and, in parentheses, the standard deviation across 25 draws. The "Average" column reports the mean across the two models. There is typically little variation within model across draws and larger variation across models. Appendix Table A.1 shows results across all seven models we considered.

Figure 3: Automation Rates, x_t

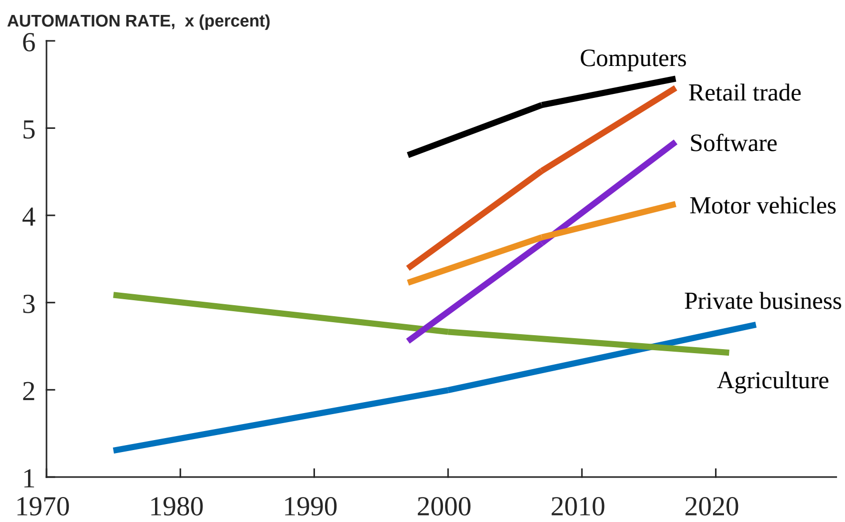

Note: The automation rate measures "What share of total labor costs is accounted for by tasks that get automated in each year?" It is an average annual rate over 25-year periods for the private business sector and for agriculture, and over decades for the other sectors.

Table 5: Top 5 Task Contributors to the Automation Rate

Task category	Labor cost share, $\bar{\omega}_j$	Automation rate, x_j	Contribution of task j	Share
Agriculture, 2000–2021 $(x = 2.4\%)$				
Feed livestock	0.054	0.037	0.0019	7.9%
Milk dairy animals	0.044	0.042	0.0018	7.6%
Hand–pick fruits/vegetables	0.144	0.008	0.0012	4.8%
Collect & handle eggs	0.020	0.048	0.0010	4.3%
Spray/dust pesticides	0.017	0.041	0.0007	2.8%
			Top 5 tasks account for	27.3%
Retail Trade, 2007–2017 $(x = 5.5\%)$				
Process non–cash payments	0.027	0.110	0.0029	5.3%
Compute totals and tax	0.017	0.131	0.0023	4.2%
Sales/labor/inventory reports	0.020	0.108	0.0021	3.9%
Gasoline–station courtesy service	0.027	0.067	0.0020	3.7%
Schedule and timekeeping	0.022	0.088	0.0019	3.5%
			Top 5 tasks account for	20.5%

Note: The automation rate x_j for task j is $-\Delta \log(1-\beta_j)/\Delta t$. The contribution of task j is $\bar{\omega}_j \cdot x_j$, where $\bar{\omega}_j$ is the Tornqvist (average) labor cost share. However, the contribution column is not literally the product of the previous two columns because we first take the product and then average across the draws. Share is the contribution as a fraction of the sector-level automation rate x.

is declining. For the BEA/BLS sectors, automation rates are generally increasing and higher than for agriculture or private business. The computer sector always has the fastest rate of automation.

Table 5 shows the top 5 task contributors to the automation rate for agriculture and retail trade. A high contribution to the automation rate requires both a high labor share working on the task and rapid automation within that task. For agriculture, the top 5 tasks account for 27.3% of the automation rate and include tasks related to livestock, produce harvesting, and pesticide spraying. Interestingly, "hand-picked fruits and vegetables" is a top-5 contributor, even with a relatively slow automation rate, because it uses so much labor. For retail trade, the top 5 tasks account for 20.5% of the automation rate. The key tasks here are processing payments, computing tax and totals, generating reports, scheduling, and the decline of full-service gasoline. Most of these are tasks where information technology and self-service have played a major role in automating retail trade.

It is fascinating to spot-check the automation that occurs on various tasks. For example, for "milk dairy animals," a key innovation is the rotary milking parlor introduced in the 2010s that allows 400-800 cows to be milked an hour with minimal staff (Therio.ai, 2026); see https://www.youtube.com/watch?v=P6OtnUmw37U. In the year 2000, there were only 800 milking robots worldwide but this number rose to around 100,000 by 2024 (Meihak, 2024).

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

We now implement Proposition 4 empirically to measure the contribution of $\hat{Z}_t + \hat{\psi}_{\ell t}$ and $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ to TFP growth. This is shown in Table 6.

$\hat{Z}_t + \hat{\psi}_{\ell t}$ is Small. The first three columns calculate $\hat{Z}_t + \hat{\psi}_{\ell t}$ as $\hat{A}_t - \frac{1}{1-\sigma} x_t$. Notice that with $\sigma = 0.2$, we have $\frac{1}{1-\sigma} = 1.25$, so this involves inflating the automation rate by an extra 25% before subtracting. The striking result is that $\hat{Z}_t + \hat{\psi}_{\ell t}$ is small or even negative in all sectors other than computers. These two sources account for little TFP growth in most sectors. This finding is somewhat surprising but perhaps not unreasonable. 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. Indeed, the National Assessment of Educational Progress (2023) long-term trend assessment for 13 year olds shows essentially zero cumulative progress in reading and mathematics scores since 1973. For the "other" TFP growth in $\hat{Z}_t$, it is possible that rising misallocation offsets increases in variety or quality; for example, see Bils, Klenow, and Ruane (2021). However, we also consider other sources of mismeasurement in Panels B and C, discussed shortly.

$\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ is Large. The last three columns of Table 6 calculate $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ as in Proposition 4 by subtracting $\hat{Z}_t + \hat{\psi}_{\ell t}$ from the TFP growth rate and scaling by the capital share. The key finding is that $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ is large. For the private business sector, the rate is 3.8% per year. Other sectors are even higher, including computers at 10.7% and an incredible 23%

Table 6: Measuring Baseline TFP Growth $\hat{Z}_t + \hat{\psi}_{\ell t}$ and the Automation Boost $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$

Sector	Growth rate of A_t	Automation rate, x_t	Baseline $\hat{Z}_t + \hat{\psi}_{\ell t}$	TFP Growth	Capital share, s_{Kt}	Automation boost, $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$
			Panel A: Main Results			
Private business	2.3	2.0	−0.2	1.2	0.35	3.8
Agriculture	4.3	2.7	0.9	3.3	0.57	4.4
Computers	14.9	5.2	8.4	12.8	0.41	10.7
Motor vehicles	2.7	3.7	−2.0	1.7	0.43	8.6
Retail trade	2.4	4.5	−3.1	1.7	0.20	23.3
Software	4.6	3.7	−0.1	1.8	0.47	4.0
			Panel B: Robustness to Mismeasurement of x_t			
Private business	2.3	1.8	0.0	1.2	0.35	3.3
Agriculture	4.3	2.7	0.9	3.3	0.57	4.4
Computers	14.9	5.2	8.4	12.8	0.41	10.7
Motor vehicles	2.7	2.1	0.0	1.7	0.43	4.1
Retail trade	2.4	1.9	0.0	1.7	0.20	8.2
Software	4.6	3.6	0.0	1.8	0.47	3.9
			Panel C: Robustness to Mismeasurement of TFP Growth			
Private business	2.5	2.0	0.0	1.4	0.35	3.8
Agriculture	4.3	2.7	0.9	3.3	0.57	4.4
Computers	14.9	5.2	8.4	12.8	0.41	10.7
Motor vehicles	4.6	3.7	0.0	3.7	0.43	8.6
Retail trade	5.6	4.5	0.0	4.8	0.20	23.3
Software	4.6	3.7	0.0	1.9	0.47	4.0

Note: Measured according to equation (15) and Proposition 4. Panel B replaces x_t by the implied rate need to raise $\hat{Z}_t + \hat{\psi}_{\ell t}$ to zero when it is negative in Panel A. Panel C instead raises $\hat{Z}_t$ by the amount needed to make the sum zero in the same four sectors; this raises $\hat{A}_t$ and $T\hat{F}P$ by the same amount, while $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ is invariant.

in retail trade (the latter is at least partly due to a low capital share). In other words, the average rate at which machines are improving is substantially higher than task-specific labor productivity growth. Machines get better much faster than people do. 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.

Measurement Error. An alternative interpretation of the negative $\hat{Z}_t + \hat{\psi}_{\ell t}$ in Panel A is that we have overstated the automation rate. Panel B of Table 6 asks what value of x_t would be required to drive $\hat{Z}_t + \hat{\psi}_{\ell t}$ to zero in sectors where Panel A measures it to be negative. The implied correction is essentially nil for private business and software, but requires cutting measured x_t by roughly half for motor vehicles ($3.7\% \to 2.1\%$ per year) and retail trade ($4.5\% \to 1.9\%$). Under this interpretation $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ is smaller — between 3% and 8% per year across the four adjusted sectors rather than 4% to 23% — because the negative $\hat{Z}_t + \hat{\psi}_{\ell t}$ term no longer must be overcome.

A second mismeasurement story turns this around: suppose TFP growth itself has been understated. It is well-known that TFP growth in manufacturing was negative between 2006 and 2024, which seems hard to believe; historically, this is the well-measured sector with rapid productivity growth.[9] Panel C of Table 6 raises $\hat{Z}_t$ by exactly the amount required to bring $\hat{Z}_t + \hat{\psi}_{\ell t}$ to zero in each of the four sectors where it is currently negative in Panel A: 0.2 percentage points for private business, 2.0 for motor vehicles, 3.1 for retail trade, and 0.1 for software. TFP growth is raised by this same amount, which is large for motor vehicles and retail trade. With this approach, the $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ term is invariant because both TFP growth and $\hat{Z}_t$ change by the same amount.

It is possible that some combination of these two sources of mismeasurement explain the outliers of motor vehicles and retail trade. Importantly, in either case — and therefore in linear combinations of the two explanations — the point that the automation boost term $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ is large and accounts for most growth is robust.

4. Counterfactual: Freezing Automation in the Initial Year

We now 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?

[9] See https://fred.stlouisfed.org/series/MFGPROD.

4.1 Theory

We allow task-level productivity to grow as measured in the historical data, but restrict further automation — these improved machines cannot be used on an increasing share of tasks.

Assumption 2 (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 $\hat{\psi}_{\ell t}$, and the "other" growth rate, $\hat{Z}_t$, are unchanged.

In a structural model with endogenous technological change, this need not be the case: research effort may shift when tasks can be automated. This counterfactual, however, is a useful tool to account for the historical contribution of automation to growth.

To implement this counterfactual, recall our key equation for TFP growth:

$$\widehat{TFP}_t = s_{Kt}(\hat{\psi}_{kt} - \hat{\psi}_{\ell t}) + \hat{Z}_t + \hat{\psi}_{\ell t}. \tag{20}$$

Under Assumption 2, 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

$$\widehat{TFP}_t - \widehat{TFP}_t^{cf} = \left(s_{Kt} - s_{Kt}^{cf}\right)(\hat{\psi}_{kt} - \hat{\psi}_{\ell t}). \tag{21}$$

Thus, we need to compute the counterfactual capital factor income share s_{Kt}^{cf}. 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. All else equal, $\dot{\beta}_t > 0$ lowers B_t and raises A_t, which raises the capital share since $\sigma < 1$. In the counterfactual, we set $\dot{\beta}_t^{cf} = 0$. Thus, the capital share in the counterfactual declines over time: the "machines getting better" force lowers the capital share and we've turned off the "more tasks get automated" force that historically kept the capital share from falling.

The evolution of the capital share in the counterfactual is

$$\frac{s_{Kt}^{cf}}{1-s_{Kt}^{cf}} = \frac{s_{Kt}}{1-s_{Kt}} \exp\left(-\int_0^t \frac{1}{s_{K\tau}} x_\tau d\tau\right). \tag{22}$$

The counterfactual capital share starts with the actual capital share and then "undoes" the contribution from automation.[10]

Putting all this together, we have the following proposition:

Proposition 5 (*Counterfactual contribution of automation*)**.** Under Assumptions 1 and 2, the lost TFP growth from "freezing" the set of automated tasks in some historical year satisfies

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

where s_{Kt}^{cf} is given by equation (22) and $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ is given in Proposition 4.

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 7 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 range from 26% to 66%. 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 $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ (from Table 6): capital productivity growth is very fast. Historically, the automation rate of $x_t = 2.0\%$ counteracted this rapid rate of capital productivity growth, to leave the capital share roughly constant.

[10]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 under the assumption of a fixed capital-output ratio, which changes the equation slightly.

Table 7: Counterfactual Contribution of Automation

Sector	Capital share $s_{K,T}$	$s_{K,T}^{cf}$	$\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$	Lost Growth $\widehat{TFP}_t - \widehat{TFP}_t^{cf}$	TFP Growth	Lost Growth Share of $\widehat{TFP}_t$
			Automation set frozen in 1950:			
Private business	0.420	0.000	3.8	1.3	1.2	108%
Agriculture	0.655	0.052	4.4	1.5	3.3	45%
			Automation set frozen in 1987:			
Computers	0.459	0.018	10.7	3.2	12.8	25%
Motor vehicles	0.524	0.073	8.6	2.2	1.7	126%
Retail trade	0.259	0.000	23.3	4.1	1.7	242%
Software	0.463	0.076	4.0	1.0	1.8	53%

Note: The counterfactual contribution of automation supposes the set of tasks that are automated is frozen in some initial year and is calculated according to Proposition 5. The "lost growth" column is not literally calculated from the previous columns because the calculation is done across multiple 25-year or 10-year periods rather than just in one long difference. Growth rates are percents per year, averaged over the relevant time period.

The fourth column in Table 7 implements Proposition 5 to measure the "lost growth" that comes from freezing automation in the initial year. For the private business sector, this lost growth is 1.3 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 more than all of measured TFP growth — the counterfactual path would have had *negative* TFP growth because $\hat{Z}_t + \hat{\psi}_{\ell t}$ is negative.

Across the other sectors of the economy, the missing growth ranges from 1.0 percent per year (software) to 4.1 percent per year (retail trade). The last column expresses this as a share of TFP growth. Across the BEA sectors, freezing automation in 1987 would have cost 25% of growth in the computer sector, 53% in software, and more than all of measured TFP growth in motor vehicles and retail trade.

This "freeze automation" exercise reinforces an important lesson. Machines improving rapidly on a fixed set of tasks is not sufficient to sustain growth. Even with infinite quantities of a fixed set of inputs, output is not infinite because production is constrained by the remaining weak links. Historically, long-run growth occurred because

Table 8: Robustness of Automation Results

Sector	$\sigma = 0.1$	$\sigma = 0.3$	$\sigma = 0.5$	$x_t \to 0.8\, x_t$	GPT-5.5	Claude Opus 4.7
			“Other” TFP growth: $\hat{Z}_t + \hat{\psi}_{\ell t}$			
Private business	0.1	−0.5	−1.6	0.3	0.0	−0.4
Agriculture	1.2	0.4	−0.9	1.6	1.4	0.3
Computers	9.0	7.7	5.2	9.7	9.0	7.8
Motor vehicles	−1.6	−2.4	−3.9	−1.0	−1.1	−2.8
Retail trade	−2.6	−3.8	−6.1	−2.0	−2.8	−3.5
Software	0.4	−0.7	−2.6	0.9	0.3	−0.4
			Excess rate at which machines getting better: $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$			
Private business	3.1	4.8	7.8	2.4	3.3	4.4
Agriculture	3.8	5.1	7.5	3.2	3.4	5.3
Computers	9.3	12.6	18.4	7.6	9.3	12.1
Motor vehicles	7.8	9.7	13.0	6.4	6.6	10.6
Retail trade	20.6	26.7	37.5	17.9	21.7	24.9
Software	3.0	5.3	9.4	2.1	3.3	4.7
			Lost TFP growth – freeze automation (percent share)			
Private business	91	129	190	70	96	121
Agriculture	40	52	73	30	33	59
Computers	21	29	43	16	21	29
Motor vehicles	114	141	190	83	86	168
Retail trade	214	278	393	179	220	264
Software	39	71	128	24	40	66

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 $\sigma = 0.2$. The fourth column ($x_t \to 0.8\, x_t$) holds σ at the baseline but scales the measured automation rate down to 80% of its value. The last two columns use the model-specific automation rate from GPT-5.5 and Claude Opus 4.7 individually.

we found ways to rapidly improve the productivity of machines while also increasing the set of tasks that benefited from this rapid growth, strengthening more of our weak links.

4.3 Robustness

So far, all results use our calibrated value of $\sigma = 0.2$ and automation rates obtained by averaging measurements from GPT-5.5 and Claude Opus 4.7. In Table 8, we explore sensitivity to these values in order to illustrate the robustness of the main findings.

Two results from the robustness exercise stand out. First, given our measured automation rates, values of σ above 0.3 lead to implausibly negative values of $\hat{Z}_t + \hat{\psi}_{\ell t}$ and very large values of $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$. Second, we consider an alternative where the true automation rate is 80% of the measured rate. This leads to plausible results for $\hat{Z}_t + \hat{\psi}_{\ell t}$ while confirming the result that the automation boost $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ is substantial. Essentially, TFP growth rates are relatively modest in the U.S. economy, so some combination of low elasticities of substitution and automation rates that are not too large is needed to give plausible values for $\hat{Z}_t + \hat{\psi}_{\ell t}$. As long as automation rates are not too small, the automation boost $\hat{\psi}_{kt} - \hat{\psi}_{\ell t}$ remains substantial.

5. The Future of A.I.

A.I. is the latest form of an automation process that has been ongoing for centuries. With our historical accounting in hand, we now consider the potential consequences of A.I. for future economic growth. We develop an endogenous growth model with endogenous automation and calibrate it based on our historical evidence. A key feature of the model is that the production of ideas can also be automated. A.I. that raises the productivity of idea production obviously offers a key channel for raising economic growth.

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?

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 is constrained by the weakest links.

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.

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

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

$$
s_{jt} \equiv \frac{P_{jt} Y_{jt}}{P_t Y_t} = \left(\frac{Y_{jt}}{Y_t} \right)^{\frac{\sigma-1}{\sigma}}. \tag{23}
$$

Infinite Automation on a Fixed Set of Tasks. 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, $Y_{cf} = Y_{\varnothing t}$. Now divide both sides by initial GDP, Y_t, and use equation (23):

$$
\begin{aligned}
\frac{Y_{cf}}{Y_t} &= \frac{Y_{\varnothing t}}{Y_t} = s_{\varnothing t}^{\frac{\sigma}{\sigma-1}}, \\
&= \left(\frac{1}{1 - s_{\infty t}} \right)^{\frac{\sigma}{1-\sigma}} \approx 1 + \frac{\sigma}{1-\sigma} s_{\infty t}.
\end{aligned} \tag{24}
$$

where the approximation is valid when the initial cost share of the infinitely-automated tasks is small. When $\sigma = 1/2$, this approximation tells us that the percent gain in output from automating the tasks with infinite productivity is simply equal to the 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 task 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.[11]

[11] Formulas related to (24) appear in Aghion, Jones, and Jones (2019) and B. Jones and X. Liu (2024), but the results there are 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.

Automating Software and Cognitive Labor. Given the advances in LLMs at coding, software is generally thought to be one of the first industries that will be largely automated by A.I. The share of software in GDP is around 2%.[12] 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 $\sigma = 0.2$. We've already seen a related version of this in the past 50 years: each of us has essentially infinitely more computing power on our desks than economists did in the 1970s, but we are not infinitely more productive. In both cases, the remaining weak links constrain output.

More speculatively, transformative A.I. could 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. What would happen if half of this were fully automated with infinite productivity? With $\sigma = 0.2$, equation (24) gives a gain of $(1/(1-1/3))^{0.2/0.8} \approx 1.11$. Infinitely automating 1/3 of GDP would only raise GDP by 11%, a number that seems quite small. The logic is again one of weak links. The economy is constrained by the other two-thirds of GDP that is not automated.

5.2 Long-Run Growth with Infinite Automation

These static warm-up exercises freeze the set of tasks that are infinitely automated. Perhaps infinitely productive software would be used to perform more tasks than current software does. We now consider what happens in the long run when the set of automated tasks increases.

Figure 4 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 $\hat{A}_t = \hat{\psi}_{\ell t}$. The infinite

[12] 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 4: Three Types of Automation in the Long Run

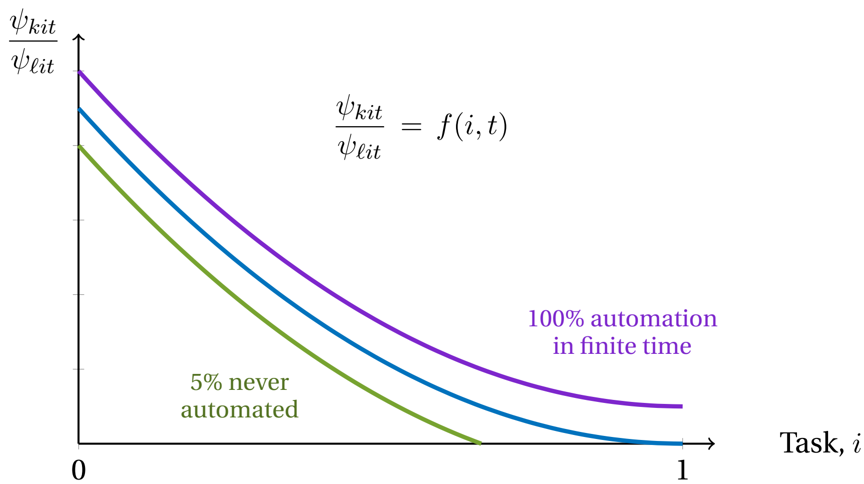

Note: The figure shows three possibilities for automation in the long run: (a) In the green line, some fraction of tasks can never be automated, so that $f(\beta, \cdot) = 0$ for $\beta > \bar{\beta}$. (b) In the purple line, $f(1, \cdot) > 0$ so that 100% of tasks are automated in finite time. (c) In 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.

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 output growth is explosive.

Finally, the blue line provides an intermediate case in which the productivity of capital on the last task is zero. Labor will therefore always be used in at least this task, even as the share of tasks using labor vanishes arbitrarily close to 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.

Table 9: 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.

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.

We now turn to our main analysis of the future and 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 start with the automation model of goods production that we've already developed, which is useful for calibrating the parameters. Second, we introduce an idea production function that allows us to endogenize ψ_{kit} and $\psi_{\ell it}$, and thus the automation process itself. Third, the model is a "lab equipment" version of an 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 9.

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:

$$f(i) = \frac{(1-i)^{\mu}}{1+\mu_0(1-i)^{\mu}} + \bar{f}. \tag{25}$$

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 4.

This functional form is monotonically decreasing, so we get a "single crossing" in the automation condition $\frac{\psi_{kit}}{\psi_{\ell it}} = Q_t^{(\theta_k-\theta_\ell)} f(i) = \frac{r_t}{w_t}$. 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 10.

The parameters μ_0 and μ, which control 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.[13]

We consider two calibrations that differ in the key parameter θ_k, which governs how fast machines improve as ideas accumulate. The first calibration sets the initial automation boost to match its historical rate in the private business sector, $\hat{\psi}_k - \hat{\psi}_\ell = 3.8\%$ per year. We view this as "A.I. as a continuation of historical patterns."

One might justifiably argue that A.I. is different in important ways, and this concern motivates our second calibration. We think of this case as "Moore's Law Everywhere": the initial automation boost throughout the entire economy is set to the rapid historical rate of the computer sector, $\hat{\psi}_k - \hat{\psi}_\ell = 10.7\%$ per year. This is a case in which A.I. is a distinct break from the past, and we see it as an upper bound on how fast automation

[13] With our functional forms, 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 $\xi(\beta_t) = \left(\int_0^{\beta_t} f(i)^{\sigma-1} di\right)^{\frac{1}{\sigma-1}}$.

Table 10: 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	.28, .45, .63, .75	1950, 1975, 2000, 2023
Task TFP growth, $\hat{\psi}_k - \hat{\psi}_\ell$	3.8% or 10.7%	Pvt Business or Computers
Labor task TFP growth, $\hat{\psi}_\ell$	0.5%	Assume humans improve slowly
	Chosen to match moments in the data / first half of paper	
ψ_{kit} idea elasticity, θ_k	4.34 or 11.23	Pvt Business or Computers
$\psi_{\ell it}$ idea elasticity, θ_ℓ	0.5	$\hat{\psi}_{\ell t} = 0.5\%$ per year
"Other" TFP growth, $\hat{Z}_t$	-0.7%	From $\hat{A}_t$ and automation rate. Decays to zero
$f(i)$ parameters: μ, μ_0	3.48, 3.61	To match capital share, s_{Kt}, given β_t, 1950 – 2023
Initial ideas, Q_0	0.75 or 0.90	Pvt Business or Computers
Initial "other TFP", Z_0	22.3 or 20.4	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	0.2	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

might accelerate the economy in the near term.

Both calibrations share the same $\hat{\psi}_{\ell 0} = 0.5\%$ per year, chosen via introspection to be consistent with the low value of $\hat{Z} + \hat{\psi}_\ell$ in the data.[14] These calibrations only pin down the initial period's productivity growth rates; the future path of $\hat{\psi}_{kt}$, $\hat{\psi}_{\ell t}$, and x_t are determined endogenously.

As anticipated earlier in Figure 4, we consider three values of the fraction of tasks that are never automated: 3%, 0%, and -3%. Our general approach is to discipline future automation by calibrating the parameters of the comparative advantage function $f(i)$ to match historical data on β_t and s_{Kt} and then extrapolating. However, this is least likely to be informative of automation in the distant long run. To be agnostic about the end-state of automation, 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. We begin with the calibration based on the private business sector. Figure 5a 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 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 5b 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 near infinite effective capital on the automated tasks, so production settles down to $Y_t = A_t L_t$ and growth is limited by the rate at which people get better. The baseline case with $\bar{f} = 0$ has the share of tasks using labor vanishing to zero, but only as the wage goes to infinity. The automation rate x_t rises over time, which causes growth to accelerate even though the capital share remains stable.[15]

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

[15] The power functions in $f(i)$ lead to a stable capital share. 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 5: The Future if AI = Private Business

(a) The Capital Share

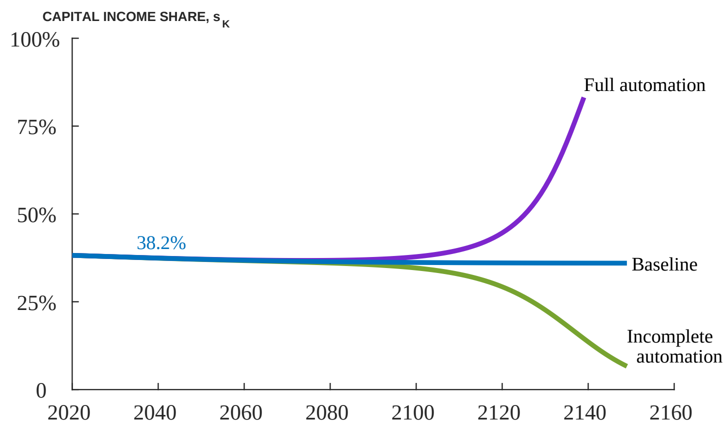

Depending on the ultimate nature of automation, the capital share can rise to 100%, fall to zero, or remain stable at its current value.

(b) Economic Growth

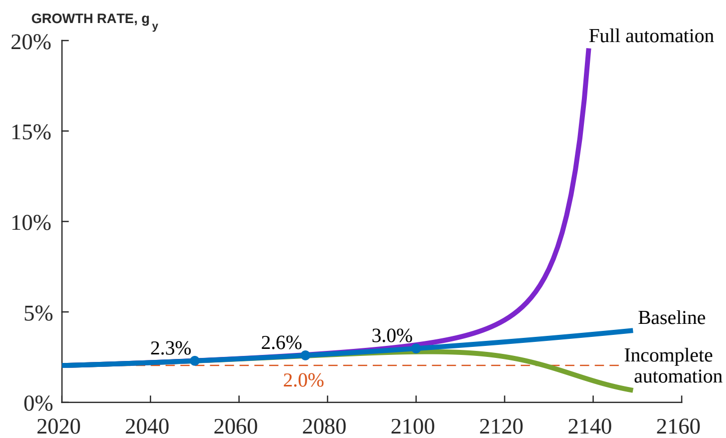

If the capital share reaches 100% in finite time, then growth explodes. If it falls to zero, growth falls to a low rate based on $\hat{\psi}_\ell = 0.5\%$.

(c) GDP per Person

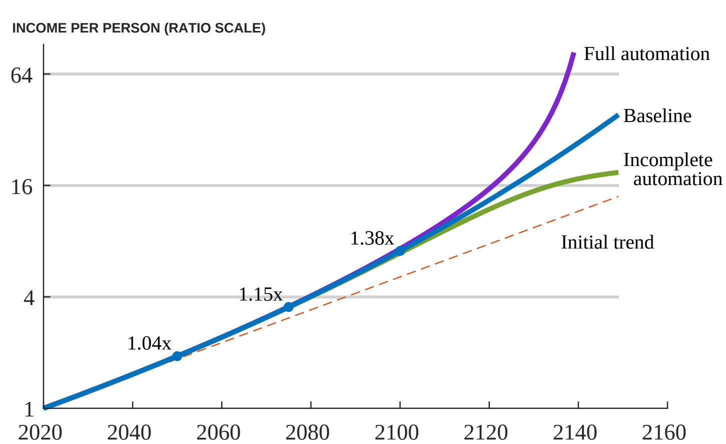

Even though the futures are eventually very different, the paths are indistinguishable for the next 50 years. The labels on the dots report the factor gain over the initial trend line.

Another important finding in Figure 5b is that even though the growth paths are eventually very different, the paths are indistinguishable for the next 50 years. A rising automation rate means that growth rates rise over time in all three cases in similar ways. The incomplete automation effects only become visible as we approach that constraint in the distant future.

To show how this accelerating growth plays out, we plot GDP per person in Figure 5c on a logarithmic scale. The red dashed line extrapolates the initial growth trend at a constant rate. The growth acceleration is then apparent in the rising slope of GDP per person for the three scenarios.

Despite the accelerating growth, the effects of A.I. on GDP per person are remarkably small for the next 30 years: incomes are only 4% higher in 2050 than they would be if growth remained constant rather than accelerating. The acceleration is slow because the economy remains constrained by weak links — the shrinking set of tasks for which human labor is still essential.

A.I. as "Moore's Law Everywhere." We next consider the calibration based on the computer sector, in which the initial automation rate is substantially faster. Figure 6 shows the results. The difference between the two calibrations is straightforward: in the "Moore's Law Everywhere" calibration, machines improve much more rapidly from the start ($\hat{\psi}_{kt} - \hat{\psi}_{\ell t} = 10.7\%$ per year instead of 3.8%). These rapid improvements drive a faster acceleration of both automation and growth.

Growth is already faster at the start of the simulation: 4.7% per year rather than 2% per year, as shown in Figure 6b. By 2030, annual growth exceeds 7% and by 2040 it exceeds 13%. The acceleration is so fast that GDP per person becomes infinite before 2060, unless some essential tasks must always be performed by labor. Nevertheless, the role of weak links in this calibration is still crucial: even with "Moore's Law Everywhere," weak links slow the explosion considerably relative to some of the most aggressive timelines that have been proposed, e.g., the "A.I. 2027" scenario in Kokotajlo et al. (2025).

Figure 6: The Future if AI = "Moore's Law Everywhere"

(a) The Capital Share

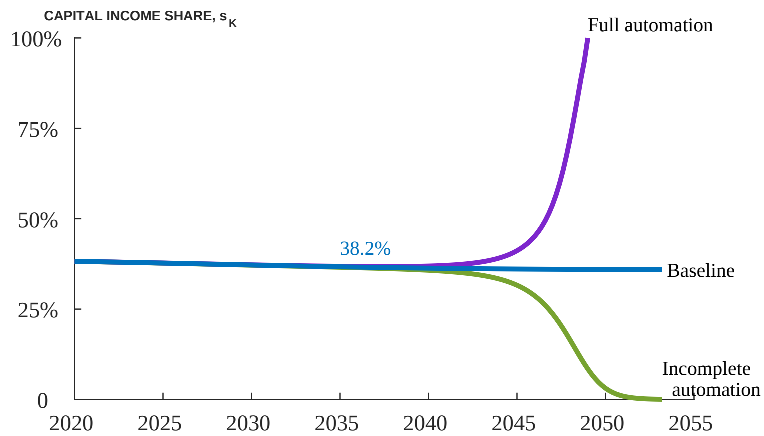

Same three cases as the private business calibration, but with faster automation.

(b) Economic Growth

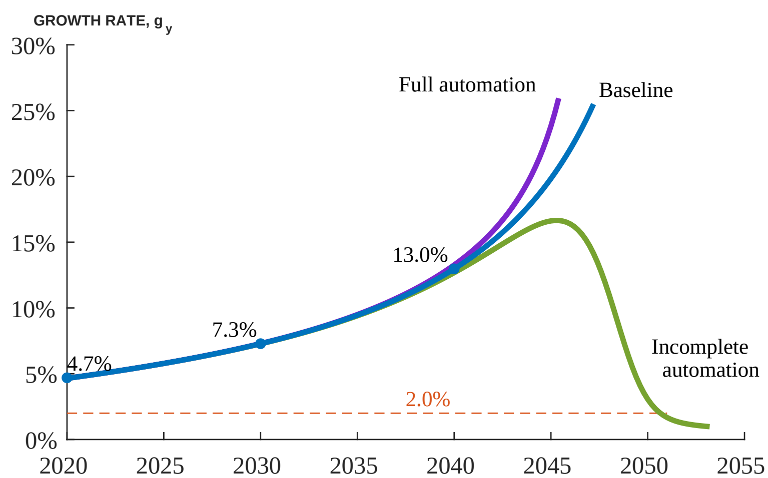

With the faster automation rate calibrated to the computer sector, growth accelerates much more rapidly.

(c) GDP per Person

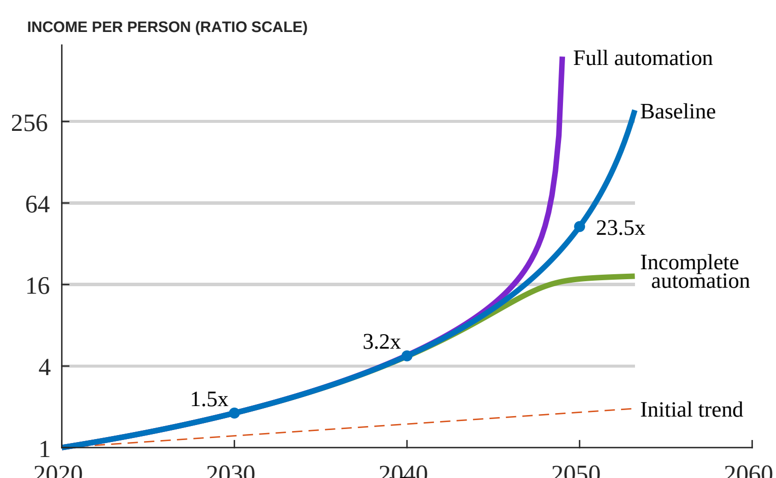

The computer-sector calibration leads to infinite income by 2060 unless some essential tasks can only be performed by labor.

5.4 Intuition for Slowly Accelerating Growth

When does growth explode, with infinite output in finite time? How long does it take for growth to explode? The answers depend on the degree of dynamic increasing returns.

It is helpful to focus on the baseline case in which $\bar{f} = 0$, where the capital share s_{Kt} stabilizes at some interior value s_K^* rather than going off to 100% or 0%. This is the case in which one might expect a standard BGP. Appendix C shows that as $\beta_t \to 1$ the capital share settles to $s_K^* = \frac{1}{\mu(1-\sigma)}$. For $\mu = 3.48$ and $\sigma = 0.2$, this gives $s_K^* \approx 0.36$.

The key measure of dynamic increasing returns in this model is Φ, defined as

$$\Phi \equiv \frac{\lambda}{1-\phi}\left(\theta_\ell + \frac{s_K^*}{1-s_K^*}\theta_k\right). \tag{26}$$

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

$$g_y = \frac{\Phi n}{1-\Phi}. \tag{27}$$

See Appendix C for the derivation.

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}.

If the degree of dynamic increasing returns is exactly unity ($\Phi = 1$), semi-endogenous growth turns into fully endogenous growth. This is the knife-edge condition that generates a BGP with endogenous growth when population is constant ($n = 0$). Of course, with positive population growth, economic growth explodes, as suggested by equation (27). If $\Phi > 1$, then equation (27) has no positive solution. This is the case "beyond endogenous growth" in which there is explosive economic growth, even with zero population growth.

Inserting our calibrated parameter values from Table 10 into equation (26) gives strikingly different values for the two calibrations: $\Phi_{pbs} = 0.98$ for the private business sector calibration and $\Phi_{comp} = 2.27$ for Moore's Law Everywhere. For the private business calibration, growth is semi-endogenous, but the long-run growth rate is $g_y = \frac{\Phi n}{1-\Phi} = 46.7\%$ per year. Because Φ is close to one, the transition to this long-run growth rate takes centuries.

We now consider two additional questions: What in the data causes Φ to be so large? And what determines the speed of the growth acceleration and explosion?

Connecting Φ to the Data. The mapping from Φ to moments in the data is:[16]

$$\Phi \approx \frac{\hat{\psi}_{\ell t} + \frac{1}{1-\sigma} x_0}{g_{Y0}}.$$

Here one can clearly see the role played by weak links (σ) and the initial automation rate (x_0). The more important are weak links — i.e., the lower is σ — the lower is the degree of dynamic increasing returns and hence the harder it is for growth to accelerate and explode. Conversely, the faster is the initial automation rate, the higher is the degree of dynamic increasing returns.

For our two calibrations using this formula, we get:

$$\Phi_{pbs} \approx \frac{0.5\% + 1.25 \times 2.0\%}{3\%} \approx 1.00,$$
$$\Phi_{comp} \approx \frac{0.5\% + 1.25 \times 5.2\%}{3\%} \approx 2.33,$$

which compare well to the exact values of $\Phi_{pbs} = 0.98$ and $\Phi_{comp} = 2.27$.

Why Are the Growth Explosions Slow? 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

[16] The approximation in this formula is that $\hat{B}_0 - \hat{Z}_0 \approx 0$.

goes to infinity is given by

$$t_\infty = \frac{1}{(\Phi - 1)\bar{g}}. \tag{28}$$

For the private business calibration, $\Phi = 0.98$ sits just below unity, placing the model on the semi-endogenous side of the knife edge $\Phi = 1$. But as an example suppose Φ were 1.10 instead. In that case, the formula in (28) gives $t_\infty = \frac{1}{0.1 \times 0.02} = 500$ years. Explosions can be slow. For Moore's Law Everywhere, the initial growth rate is 4.7% and $\Phi = 2.27$, giving $t_\infty = \frac{1}{1.27 \times 0.047} \approx 17$ years.

Of course, our dynamic system is actually two-dimensional rather than one-dimensional, so the above analysis is only a rough approximation. Nevertheless, it gives a sense of how and why explosions can be slow.

Summary. The growth model with endogenous automation of both goods and ideas features two key forces. First is the positive feedback between automation and innovation, a flywheel that accelerates growth. Second is the presence 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. The slow acceleration of growth is therefore a direct consequence of the presence of weak links. However, unless we impose that there are some essential tasks that must always be performed by labor, the flywheel effect ultimately dominates. Once the weak links are automated away, growth can accelerate substantially and even explode.

We think of our two calibrations — a "continuation of historical patterns" and "Moore's Law Everywhere" — as two extremes. In reality, the transition to A.I. is likely to involve a mix of both. Nevertheless, the basic story of how automation leads to accelerating growth while being slowed by weak links appears to be a robust feature of our analysis.

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, retail trade, 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 3.8% per year for the U.S. economy. The key benefit of automation is that we switch from using slowly-improving labor to using rapidly-improving machines on a growing share of tasks. Growth is limited by how quickly we strengthen our weak links.

Looking to the future, we develop an endogenous growth model in which the production of both goods and ideas is endogenously automated. We consider two calibrations based on our historical evidence: one in which A.I. continues the historical patterns of automation and one in which A.I. is a distinct break with the past, featuring "Moore's Law everywhere." 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.

One final insight from weak links is worth noting. In a weak link model, the benefits of automation are slow to accrue because the economy is constrained by the remaining weak links. But the same is not necessarily true with respect to A.I. safety. In particular, as Kremer (1993) originally noted, weak links make the economy fragile on the downside. Destroying one of the links in the chain can potentially cause a catastrophic collapse. We already noted that A.I. automating software engineering and computer use — which seems to be well on its way — may have small aggregate effects when the set of tasks performed by software is held fixed. However, the downside risks of this limited automation need not be small. If A.I. can use a computer as well as the best humans, then it presumably has the capability to hack the electricity grid or the banking system and to communicate with biological labs. An A.I. that is manipulated by a malicious actor could then cause substantial damage. In a world of weak links, the benefits of A.I. may arrive slowly while the dangers can arrive quickly.

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.

Bils, Mark, Peter J. Klenow, and Cian Ruane, "Misallocation or Mismeasurement?," *Journal of Monetary Economics*, 2021, *124*, S39–S56. The Real Interest Rate and the MarginalProduct of Capital in the XXIst CenturyOctober 15-16, 2020.

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.

Kokotajlo, Daniel, Scott Alexander, Thomas Larsen, Eli Lifland, and Romeo Dean, "AI 2027," *ai-2027.com*, 2025.

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.

Lashkari, Danial, Wensu Li, Christina Qiu, and Neil Thompson, "AI and Scale: A Quantitative Task-Based Theory of Automation," *Working Paper*, 2026.

Meihak, Dan, "Successfully transition into milking automation," *AgProud.com*, May 2024. https://www.agproud.com/articles/59564-successfully-transition-into-milking-automation.

National Assessment of Educational Progress, "Long-Term Trend Reading and Mathematics Assessments," Technical Report, U.S. Department of Education, Institute of Education Sciences, National Center for Education Statistics 2023.

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.

Therio.ai, "What is a Rotary Milking Parlor," 2026. https://therio.ai/dairy-directory/what-is-a-rotary-milking-parlor accessed April 22, 2026.

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.

APPENDIX

A. Instructions for the LLMs to measure β_t and $\omega_{\ell it}$

We have a three-step process for measuring the task-level inputs the growth accounting requires. First, for each sector we ask the LLM to construct a detailed list of 100 specific tasks that are essential over the relevant time frame. Second, using that task list, we ask the LLM to estimate the automation rate for each task in each benchmark year; this gives β_{it}, which aggregates to β_t. Third, we ask the LLM to estimate each task's share of total sector labor costs in each year; this gives $\omega_{\ell it}$. We repeat the second and third steps across several LLMs — GPT-5.5, Claude Opus 4.7, and others — and pool the results, as reported in Table A.1. The prompts below use motor vehicles as the running example; for other sectors we substitute sector-specific wording (e.g. "fraction of farms" for agriculture).

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. Below is a list of 100
   tasks that are essential in this sector over this time frame.
   Please do the following:
```

```
1. Consider each task in the list below.
2. Consider the year 1957. Indicate whether the task 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 plants that have automated the task in
   constructing your estimate.
3. Repeat Step 2 for 2017, the final year of our exercise, for
   that task.
4. Repeat Step 2 for the intermediate years 1967, 1977, 1987,
   1997, and 2007 for that task.
5. Repeat Steps 1 to 4 for each task in the list.

**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.

**TASK LIST**
[The 100-task list from the spreadsheet is inlined here, one task
    per line in the format:
 <id>. [<Category>] <Task Name> -- <Detailed description>]

**OUTPUT FORMAT**
Report the automation percentage (0 to 100) for each task and
   each year.
Return your results as a JSON object with the following structure
   :
{"tasks": [{"task_id": 1, "automation": {"1957": 5.0, "2017":
   75.0, ...}}, ...]}
Each task must have a "task_id" (integer) and an "automation"
   object with keys for each year: 1957, 1967, 1977, 1987, 1997,
   2007, 2017. Values are percentages from 0 to 100. Return ONLY
   the JSON object.
```

Finally, we ask the LLM to estimate each task's share of total sector labor costs. We split this into two stages to reduce the combinatorial burden on the model: Stage 1 asks for the *category-level* share of labor costs across a handful of broad task categories; Stage 2 asks for the *within-category* share of each task. We then multiply the normalized category share by the normalized within-category share to obtain $\omega_{\ell it}$ for every task. We repeat this pair of prompts for each benchmark year. Here is the Stage 1 prompt:

```
**INSTRUCTIONS FOR LABOR COST SHARES BY CATEGORY**
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 need to estimate what percentage of total labor costs (wages,
    salaries, and benefits paid to workers) in motor vehicles is
    accounted for by each category of tasks in the year 1957.

Here are the categories of tasks, with the number of individual
    tasks in each:
1. Stamping & Body Fabrication (10 tasks)
2. Paint & Sealing (10 tasks)
3. General Assembly (20 tasks)
...
11. Management/IE/H&S (4 tasks)

For each category, estimate the percentage of total labor costs
    in motor vehicles devoted to that category in 1957. The
    percentages should ideally sum to 100, but don't worry about
    exact precision -- just give your best estimate for each
    category independently.

Important: estimate labor costs (wages/salaries/benefits), not
    output value or capital costs.

**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.
```

```
(b) Include assembly line workers, quality inspectors,
    maintenance, office staff, and engineering.
(c) Estimate labor costs (wages, salaries, benefits paid to
    workers), not output value or capital costs.

**OUTPUT FORMAT**
Return a JSON object using the numeric category IDs shown above:
{"categories": [{"cat_id": 1, "cost_pct": 15.0}, {"cat_id": 2, "
    cost_pct": 10.0}, ...]}
Return ONLY the JSON object.
```

And here is the Stage 2 prompt, repeated once per category (shown here for "General Assembly"):

```
**INSTRUCTIONS FOR LABOR COST SHARES WITHIN A CATEGORY**
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.

Consider the category "General Assembly" within motor vehicles in
     the year 1957.

Below are the 20 tasks in this category. For each task, estimate
    what percentage of this category's labor costs is devoted to
    that task. The percentages should ideally sum to 100, but don'
    t worry about exact precision -- just give your best estimate
    independently.

Important: estimate labor costs (wages/salaries/benefits), not
    output value or capital costs.

**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.
(b) Include assembly line workers, quality inspectors,
```

Table A.1: Automation Rates by Sector and Model — All Models (%)

Sector	GPT-5.5	GPT-5.4 Thinking	GPT-5.4 No thinking	Claude Opus 4.7	Claude Opus 4.6	Claude Sonnet 4.6	Gemini 3 Flash
Private business	1.84 (0.17)	2.24 (0.20)	2.07 (0.10)	2.15 (0.07)	1.37 (0.05)	1.24 (0.08)	4.01 (0.37)
Agriculture	2.30 (0.29)	2.50 (0.30)	2.40 (0.20)	3.18 (0.26)	2.49 (0.07)	1.62 (0.10)	5.13 (1.14)
Computers	4.71 (0.78)	4.79 (0.93)	6.12 (0.68)	5.63 (0.32)	4.78 (0.35)	3.79 (0.31)	10.67 (2.76)
Motor vehicles	3.02 (0.49)	4.25 (0.40)	4.57 (0.26)	4.39 (0.15)	3.46 (0.08)	2.90 (0.27)	9.13 (1.73)
Retail trade	4.16 (0.71)	5.40 (1.16)	6.58 (0.79)	4.75 (0.35)	2.81 (0.20)	2.50 (0.29)	9.39 (1.71)
Software	3.41 (0.44)	3.15 (0.49)	2.99 (0.38)	3.97 (0.31)	1.76 (0.18)	1.96 (0.17)	7.17 (1.19)

Note: Each cell shows the mean automation rate and standard deviation (in parentheses) across 25 draws. There is typically little variation within model across draws and larger variation across models. Gemini 3 Flash has fewer draws for some sectors. The benchmark Average pool used in the body of the paper combines the GPT-5.5 and Claude Opus 4.7 columns.

```
    maintenance, office staff, and engineering.
(c) Estimate labor costs (wages, salaries, benefits paid to
    workers), not output value or capital costs.

**TASK LIST**
[The tasks belonging to "General Assembly" are inlined here, one
    per line.]

**OUTPUT FORMAT**
Return a JSON object:
{"tasks": [{"task_id": 21, "cost_pct": 8.5}, {"task_id": 22, "
    cost_pct": 12.0}, ...]}
Each task must have a "task_id" (integer matching the IDs in the
    list) and "cost_pct" (a non-negative number). Return ONLY the
    JSON object.
```

A.1 Task Category Automation Rates

We use the Tornqvist approximation to the continuous-time automation rate, as is common in the growth accounting literature:

$$x_{t,t+1} = \sum_j \tilde{\omega}^j_{\ell t,t+1} x^j_{t,t+1}$$

where $\tilde{\omega}^j_{\ell t,t+1}$ is the average of the labor cost shares for category j in the years t and $t+1$, and $x^j_{t,t+1} = -\Delta \log(1-\beta^j_{t+1})/\Delta t$ is the annualized automation rate for category j between t and $t+1$.

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

$$L_{it} = (\psi_{\ell it} Z_t)^{\sigma-1} \left(\frac{w_t}{P_t}\right)^{-\sigma} Y_t. \tag{29}$$

Integrating this equation over all tasks yields

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

Taking ratios of these last two equations yields

$$\frac{L_{it}}{L_t} = \frac{w_t L_{it}}{w_t L_t} = \frac{\psi^{\sigma-1}_{\ell it}}{\int_{\Omega_{\ell t}} \psi^{\sigma-1}_{\ell it}\, di} \equiv \omega_{\ell it}. \tag{30}$$

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

$$\frac{K_{it}}{K_t} = \frac{r_t K_{it}}{r_t K_t} = \frac{\psi^{\sigma-1}_{kit}}{\int_{\Omega_{kt}} \psi^{\sigma-1}_{kit}\, di} \equiv \omega_{kit}. \tag{31}$$

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:

$$\bar{\omega}_{k\beta t} \equiv \sum_{m=1}^{M_t} \omega_{k\beta_m t} \frac{|\dot{\beta}_{mt}|}{\dot{\beta}_t} \quad \text{and} \quad \bar{\omega}_{\ell\beta t} \equiv \sum_{m=1}^{M_t} \omega_{\ell\beta_m t} \frac{|\dot{\beta}_{mt}|}{\dot{\beta}_t} \tag{32}$$

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 $\bar{\omega}_{\ell\beta t} = \omega_{\ell\beta t}$.

To be completed.

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

$$Y_t = Z_t \left(\beta_t \left(\frac{\psi_{kt} K_t}{\beta_t} \right)^{\frac{\sigma-1}{\sigma}} + (1 - \beta_t) \left(\frac{\psi_{\ell t} L_t}{1 - \beta_t} \right)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \tag{33}$$

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 (33) as $Y_t = F(B_t K_t, A_t L_t)$, where B_t collects the first two β_t terms:

$$B_t = Z_t \left(\frac{1}{\beta_t}\right)^{\frac{1}{1-\sigma}} \psi_{kt} \tag{34}$$

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

$$A_t = Z_t \left(\frac{1}{1-\beta_t}\right)^{\frac{1}{1-\sigma}} \psi_{\ell t} \tag{35}$$

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

$$g_y = \left(\theta_\ell + \frac{s_K^*}{1-s_K^*}\theta_k\right) g_Q \tag{36}$$

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:

$$g_y = \frac{\Phi n}{1-\Phi} \quad \text{where} \quad \Phi \equiv \frac{\lambda}{1-\phi}\left(\theta_\ell + \frac{s_K^*}{1-s_K^*}\theta_k\right) \tag{37}$$

Connecting Φ to the Data. Add this derivation as well.

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