Artificial Intelligence and Macroeconomic Dynamics: Growth, Pricing, and Distribution

Helen Popper* Christoph Schult†

DRAFT — September 29, 2025

Abstract

This paper builds a simple general equilibrium model in which an AI producer is a monopolist who both learns by doing and uses AI recursively as an input. These mechanisms link today's scale to tomorrow's costs, so pricing is dynamic: the firm sets a price below the static monopoly benchmark to expand capacity and speed learning. Final goods are produced by monopolistic competitors using both specialized and nonspecialized labor and with constant returns to scale each period. This setting is rich enough to generate two distinct implementation patterns–learn first or scale first, and the premium to specialized labor can rise or fall.

We first use Cobb–Douglas technologies to solve for a generalized balanced growth path that pins down the condition for stable, non-explosive growth. On this path, AI output grows faster than final output, the relative price of AI falls persistently, real wages rise with overall output, and the specialized–to–nonspecialized wage ratio is flat.

We then analyze CES producdtion and derive a closed-form effective demand elasticity for AI that combines input substitution in production with final-goods market substitution across varieties. Finally, simulations allow us to explore the dynamics. When final-goods inputs are complements, adoption is learning-first and capital-light before scaling; when they are substitutes, adoption is scale-first and the two-phase pattern attenuates. Greater substitutability in AI production amplifies these patterns without changing their sign. On the distribution side, the specialized–to–nonspecialized wage premium rises typically rises; however, it falls when inputs are substitutes in AI production but complements in final goods production.

*Santa Clara University, Email: hpopper@scu.edu

†Halle Institute for Economic Research, Email: Christoph.Schult@iwh-halle.de

1 Overview

This paper describes a simple, closed-economy model of a macroeconomy that includes AI production. In this model, AI is an intermediate good produced by a monopolist that uses its own output as an input and learns by doing. So, the AI price path and profits are determined endogenously.

The combination of learning by doing and AI's use in its own production allows for a rich characterization of observed AI development cycles. Automated design, coding and data processing—once instantiated—can be reused to refine or extend themselves; and existing models and software libraries are used to generate better subsequent models. While learning by doing reflects the accumulation of knowledge through repeated production, recursive self use represents a strategic decision: the producer's internal use of AI enhances future capabilities but comes with an opportunity cost in terms of foregone sales.

The dual mechanism underscores both the cumulative nature of experience-driven improvements and the allocation decision inherent in recursive use. The result is that the AI producer internalizes the mechanism via a dynamic term in its pricing rule, which early on pushes the price below a standard, static benchmark price.

We introduce final-goods producers as monopolistic competitors; and both AI and final goods are produced using a constant-returns-to-scale technology in each period. We show that this set-up implies that learning by doing dampens the AI monopolist's markup, the production of AI accelerates, and final-goods production rises. We also distinguish between specialized and nonspecialized labor; and we characterize the behavior of their relative wages and of profits.

After introducing the model and some general results, we focus on the special case of Cobb-Douglas production. This allows us to derive a generalized balanced growth path (GBGP) analytically. On the GBGP, we describe the long-term growth rates of AI, of final production, and of the relative prices of AI, of energy, and of wages. We show that AI production grows faster than final output, and the price of AI falls continuously toward zero. In this simple case, the wage of specialized labor relative to nonspecialized labor remains constant. The GBGP gives a closed-form long-run benchmark. Specifically, it pins down: the condition for non-explosive growth with constant shares, the relative growth of AI vs. final output, and the trend in the AI price.

Next, we move from Cobb–Douglas to a CES specification and simulate

the model over a range of elasticities. The distribution of income evolves over time, and its dynamics depend on these elasticities. The AI monopolist's forward-looking pricing produces an initial sharp price decline. The AI price decline generates a profit dip and a fall in total AI capital when inputs are complements in final-goods production. These effects are mitigated—though not eliminated—by substitutability within AI production; and they vanish or even change sign when inputs are substitutes in final-goods production.

These simulations reveal an early-adoption pattern that depends on input substitutability in final-goods production. With recursive self-use and forward-looking pricing of AI as an intermediate input, complements yield a learning-first, capital-light phase; while substitutes yield a scale-driven, capital-deepening phase. And, throughout, greater substitutability in final-goods production raises the specialized–to–nonspecialized wage premium; with complements, the premium is lower.[1]

The general equilibrium approach of our paper builds on Acemoglu and Restrepo's seminal series of task-based models of automation (e.g., Acemoglu and Restrepo 2018, 2020). Those papers characterize balanced growth paths driven by the reallocation and creation of tasks. Our approach also draws on Korinek and Stiglitz's (2019) general equilibrium model in which AI can be treated as a perfect substitute for labor. We complement these approaches by endogenizing the dynamics of AI production and pricing, and by characterizing the associated implications for capital, profits and the wage of specialized labor relative to that of nonspecialized labor.[2]

In separating these two types of labor–specialized and nonspecialized, our work is in the spirit of Benzell and Brynjolfsson's (2019) distinction between scarce 'genius' labor and ordinary labor. Benzell and Brynjolfsson develop a general equilibrium model in which the replicability of digital technologies raises the returns to scarce "genius" labor while compressing

[1]Related ongoing work extends the model to allow for heterogeneous households. The heterogeneous approach allows for the characterization of the distribution of welfare, not just income. In addition, the formulation lends itself to an evaluation of the effects of a universal basic income (UBI) while incorporating participation utility from work (alongside disutility from hours) and the role of entrepreneurial contributions.

[2]Both Acemoglu–Restrepo and Korinek–Stiglitz treat AI as a replacement for human labor. We abstract from replacement mechanisms that can lower real wages and instead analyze how recursive self-use and learning shape AI pricing, adoption, profits, and the skilled–nonskilled wage gap.

ordinary wages and the interest rate. Our contribution extends this insight by embedding the distinction between labor types in a dynamic framework with endogenous AI production; this allows us to characterize not only distributional outcomes but also the long-run evolution of AI production, prices, and profits.[3]

Our treatment of AI as an intermediate input with recursive use connects directly to other recent work differentiating AI from earlier technologies. Aghion, Jones, and Jones (2017) model AI as a factor in idea production that can substitute for researchers; Brynjolfsson and Mitchell (2017) emphasize AI as a general-purpose technology; Cockburn, Henderson, and Stern (2019) describe AI as an "invention of a method of invention" with broad spillovers; and Veldkamp and Chung (2024) model data as a factor of production with increasing returns. Our approach offers a way to capture aspects of these ideas within a unified macroeconomic framework.

2 The Model Setup

2.1 AI Production

We treat AI as an intermediate good produced by a monopolist. The monopolist benefits from learning by doing and produces AI (A_t) using inputs of capital ($K_{A,t}$), specialized labor (${L^s}_{A,t}$), energy ($E_{A,t}$), and recursively using AI as an input, ($A_{A,t}$):

$$A_t = (A_{t-1})^{\phi} f(K_{A,t}, L^s_{A,t}, E_{A,t}, A_{A,t}) \tag{2.1}$$

where $\phi \in (0, 1)$ is a parameter that governs learning by doing based on the producer's prior output, AI is used recursively as an input into AI, and t denotes the time period. We assume that f itself is constant returns to scale in each period.

Note also that while production exhibits constant returns to scale in each period (the distribution parameters sum to one), it exhibits dynamically increasing returns to scale since $\phi > 0$. At the same time, the restriction that $\phi < 1$ prevents explosive dynamics and keeps the model tractable.[4]

[3]Like their work, ours also reflects (though to a lesser extent) the issues raised in Brynjolfsson, Chandar, and Chen's (2025) empirical study of generative AI's labor-market impacts.

[4]Nordhaus (2021), among others, explores the possibility of such explosive dynamics.

2.1.1 AI Producer's Problem

The AI monopolist maximizes the present discounted value of future profits:

$$\max_{\{A_t, K_{A,t}, L^s_{A,t}, E_{A,t}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \Pi_{A,t}, \tag{2.2}$$

subject to the AI production function, Equation 2.1, and period profits are:

$$\Pi_{A,t} = p_{A,t} A_t - w^s_t L^s_{A,t} - r_t K_{A,t} - p_{E,t} E_{A,t} - p_{A,t} A_{A,t},$$

where $p_{A,t},\ w^s_t,\ p_{E,t}$ are the period-t prices of AI, specialized labor, and energy; and r_t is the interest rate.

The AI monopolist's Bellman equation is:

$$\begin{aligned} V(A_{t-1}) = \max_{A_t} \Big\{ p_{A,t}\, A_t \ - \ C_t(A_t) \ + \ \beta\, V(A_t) \Big\} \\ \text{s.t.} \quad p_{A,t} = P_t(A_t), \end{aligned} \tag{2.3}$$

where $C_t(A_t)$ is the cost-minimizing expenditure to produce A_t using $\{K_{A,t}, L^s_{A,t}, E_{A,t}, A_{A,t}\}$ at factor prices $(r_t, w^s_t, p_{E,t}, p_{A,t})$.[5]

The first order condition is then:

$$p_{A,t}\Big(1 - \tfrac{1}{\epsilon_t}\Big) \ = \ MC_{A,t} \ - \ \beta\, V'(A_t), \tag{2.4}$$

where marginal cost is $MC_{A,t} \equiv C'_t(A_t)$; and ϵ_t denotes the effective (positive) price elasticity of demand for AI faced by the monopolist.[6]

2.1.2 Dynamic Pricing with Self-Use and Learning.

Accounting for both the benefits of learning-by-doing and the recursive use of AI in its own production, the monopolist sets prices dynamically to maximize the present value of profits.

[5] We also impose a participation constraint: $V_{A,0} = \sum_{t \geq 0} \beta^t \Pi_{A,t} \geq 0$. Note also that the recursive use of $A_{A,t}$ entails foregone revenue at the market price $p_{A,t}$.

[6] Since AI is an intermediate good, the effective elasticity reflects both the final-goods producer's input substitution and the household's demand elasticity. Under CES production (derived in the appendix), it is a convex combination of the two.

Specifically, solving the first order condition, (2.4), for price gives the dynamic monopoly pricing rule:

$$p_{A,t} = \frac{MC_{A,t} - \beta\, V'(A_t)}{1 - \frac{1}{\epsilon_t}}. \tag{2.5}$$

The term $\beta\, V'(A_t)$ captures the value of turning today's output into future capability: part from experience (learning-by-doing) and part from self-use (own output used as an input). Because $\beta\, V'(A_t) > 0$, the optimal price is strictly below the standard static monopoly price: $\frac{MC_{A,t} - \beta\, V'(A_t)}{1 - \frac{1}{\epsilon_t}} < \frac{MC_{A,t}}{1 - \frac{1}{\epsilon_t}}$.

That is, Equation 2.5 implies that to maximize long-run profits, the AI monopolist lowers prices today to expand production, accelerate learning, and reduce future costs.

Equation 2.5 also has two other implications: the price drop is higher when ϕ and α_A are higher; and the self-use share $\frac{A_{A,t}}{A_t}$ rises early, then declines as external sales scale.

2.2 Final-Goods Production

We assume a continuum of differentiated final goods indexed by $i \in [0, 1]$.[7] Each variety $Y_t(i)$ is produced by a firm with constant returns to scale:

$$Y_t(i) = Z_t h(K_{Y,t}(i), L^s_{Y,t}(i), L^n_{Y,t}(i), E_{Y,t}(i), A_{Y,t}(i)), \tag{2.6}$$

where Z_t denotes total factor productivity in the final-goods sector and is assumed to grow exogenously at the rate g_z.[8]

Aggregate final output is a CES aggregate over the differentiated varieties:

$$Y_t = \left(\int_0^1 Y_t(i)^{\frac{\varepsilon - 1}{\varepsilon}}\, di \right)^{\frac{\varepsilon}{\varepsilon - 1}}, \tag{2.7}$$

[7]The introduction of a variety of final goods allows for a finite elasticity for each good, so combined with input substitutability in final-goods production, the AI firm can charge a markup over marginal cost. Without this assumption, the AI firm faces an infinite elasticity, and no markup is possible.

[8]To ensure that an elasticity of substitution, σ_Y, is well defined, we assume that the production function is twice continuously differentiable and quasi-concave with positive input use.

where $\varepsilon > 1$ is the household's elasticity of substitution between the differentiated final goods, ε.

This gives rise to demand for each variety:

$$Y_t(i) = \left(\frac{p_t(i)}{P_t}\right)^{-\varepsilon} Y_t, \tag{2.8}$$

where $p_t(i)$ is the price of variety i, and P_t is the final-goods price index:

$$P_t = \left(\int_0^1 p_t(i)^{1-\varepsilon}\, di\right)^{\frac{1}{1-\varepsilon}} \tag{2.9}$$

We define the final good as the numeraire: $P_t = 1$. Under symmetry all firms choose the same price in equilibrium, so $p_t(i) = 1\ \forall i$, and the aggregate price index remains constant.

Each final-goods firm chooses its price and input quantities to maximize profits (subject to the production function and CES demand), where profits for variety i in period t are given by:

$$\begin{aligned}\Pi_{Y,t}(i) = {} & p_t(i)Y_t(i) \\ & - w_t^s L_{Y,t}^s(i) - w_t^n L_{Y,t}^n(i) \\ & - r_t K_{Y,t}(i) - p_{E,t}E_{Y,t}(i) - p_{A,t}A_{Y,t}(i) - f,\end{aligned} \tag{2.10}$$

where f is a fixed cost; and we assume free entry.

The pricing first-order condition implies:

$$p_t(i) = \mu_Y \cdot MC_t(i), \quad \mu_Y = \frac{\varepsilon}{\varepsilon - 1}. \tag{2.11}$$

That is, final goods are priced at a constant markup over marginal cost; and the markup, as is standard in such models, depends on the household's elasticity of substitution between the differentiated final goods.

Note that the AI monopolist, in turn, will face an effective elasticity of demand (ϵ) that reflects two distinct margins: the product-market substitutability across final-goods varieties (ε) discussed here, and the final good producer's technological substitutability of AI with other inputs. (Below, in section 4, we discuss the closed form solution, derived in the Appendix, for the case where production exhibits CES.)

2.3 The Household

The representative household consumes a continuum of differentiated final-goods varieties $C_t(i)$ for $i \in [0, 1]$, which aggregate into a composite good

C_t according to a CES aggregator:

$$C_t = \left(\int_0^1 C_t(i)^{\frac{\varepsilon-1}{\varepsilon}} \, di \right)^{\frac{\varepsilon}{\varepsilon-1}}, \quad \varepsilon > 1. \tag{2.12}$$

The household derives utility from this composite good and supplies labor inelastically, and lifetime utility is given by a standard CRRA function:

$$\max \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\sigma_c} - 1}{1 - \sigma_c}, \quad \sigma_c > 0, \sigma_c \neq 1. \tag{2.13}$$

The household earns income from supplying both types of labor, renting capital, and receiving monopoly profits from the AI firm.[9] It allocates income between consumption of final goods and saving:[10]

$$\int_0^1 p_t(i)[C_t(i) + S_t(i)]di = w_t^s L_t^s + w_t^n L_t^n + r_t K_t + \Pi_{A,t}.$$

Since we assume symmetry across goods, we can write: $P_t = p_t(i)$, $C_t = C_t(i)$, and $S_t = S_t(i)$; and, this simplifies to:

$$C_t + S_t = w_t^s L_t^s + w_t^n L_t^n + r_t K_t + \Pi_{A,t} \tag{2.14}$$

The household chooses consumption and saving subject to this constraint and the intertemporal accumulation of capital. This yields the standard Euler equation:

$$\frac{C_{t+1}}{C_t} = [\beta(1 + r_{t+1})]^{\frac{1}{\sigma_c}}. \tag{2.15}$$

2.4 Factor Markets

We assume that specialized labor, nonspecialized labor, and energy are all fixed, that capital accumulates endogenously, and that markets clear. So,

[9] Note that we assume symmetry and free entry in the final-goods sector, so final-goods firms make zero profits in equilibrium, as in Dixit and Stiglitz (1977).

[10] We assume the household does not buy energy directly.

we have the following conditions:

Labor Markets

$$L^s_{Y,t} + L^s_{A,t} = \bar{L}^s$$
$$L^n_t = \bar{L}^n$$

AI Market

$$A_{Y,t} + A_{A,t} = A_t$$

Energy Market

$$E_{Y,t} + E_{A,t} = \bar{E}$$

Interest Rate

$$r_t = MPK_Y - \delta = MPK_A - \delta.$$

Capital Accumulation[11]

$$K_{Y,t+1} = (1-\delta)K_{Y,t} + I_{Y,t}$$
$$K_{A,t+1} = (1-\delta)K_{A,t} + I_{A,t}$$
$$S_t = I_{Y,t} + I_{A,t}$$

Note that S_t reflects household saving out of income, which includes labor income, capital income, and monopoly profits.

3 A Generalized Balanced Growth Path

We derive a generalized balanced growth path (GBGP) that serves as the paper's long-run benchmark. It tells us exactly when self-use plus learning sustain stable growth with constant factor shares; and it specifies the long-run growth rates of the economy's key variables.

In order to analytically characterize this path, we specify both production of AI and of final goods to be Cobb-Douglas at the firm level.

So, AI production becomes:

$$A_t = A^{\phi}_{t-1} K^{\alpha_K}_{A,t} (L^s_{A,t})^{\alpha_{L^s}} E^{\alpha_E}_{A,t} A^{\alpha_A}_{A,t}, \tag{3.1}$$

where the αs represent the distribution parameters, with α_A capturing the recursive use of AI; and, as before, ϕ represents the extent of learning by doing.

[11] We assume that the unconsumed final goods can be converted into capital for use in the production either of AI or of any of the varieties of final goods.

Final-goods production becomes:

$$Y_t = Z_t K_{Y,t}^{\omega_K} (L_{Y,t}^s)^{\omega_{L^s}} (L_{Y,t}^n)^{\omega_{L^n}} E_{Y,t}^{\omega_E} A_{Y,t}^{\omega_A} \tag{3.2}$$

where the ωs represent the distribution parameters in final production.

In each case, we assume the distribution parameters sum to one: $\alpha_K + \alpha_{L^s} + \alpha_E + \alpha_A = 1$ and $\omega_K + \omega_{L^s} + \omega_{L^n} + \omega_E + \omega_A = 1$. And, we let Z_t grow at the rate g_Z. as in Section 2.2, we have symmetry in equilibrium; So we can treat one representative variety as determining aggregate output. Hence, the growth path of the economy can be characterized using a representative firm's Cobb-Douglas production function, Equation 3.2.

3.1 Consumption Growth, the Interest Rate, and Capital–Output Ratios in the GBGP

We begin with the household's consumption growth rate. Since the GBGP requires constant consumption growth, the Euler equation implies a constant real interest rate.

Specifically, the Euler equation (Equation 2.15, reproduced here) relates the growth rate of consumption to the interest rate:

$$\frac{C_{t+1}}{C_t} = [\beta(1 + r_{t+1})]^{\frac{1}{\sigma_c}} .$$

The GBGP requires constant consumption growth. Since the representative household satisfies a standard Euler equation, the growth rate of consumption is determined by the intertemporal preferences of the household. This pins down the real interest rate. So the Euler equation becomes:

$$\frac{C_{t+1}}{C_t} = [\beta(1 + r^*)]^{\frac{1}{\sigma_c}} , \tag{3.3}$$

where r^* is the constant interest rate in the GBGP.

Taking logs of both sides and rearranging gives the constant growth rate of consumption:

$$c^* = \frac{1}{\sigma_c} \log (\beta(1 + r^*)) \tag{3.4}$$

This expression determines the growth rate of the consumption as a function of preference parameters and the equilibrium real interest rate.

Turning to the interest rate, we note that each final good variety is produced by a monopolistically competitive firm using Cobb-Douglas technology, and although firms set prices with a markup, they minimize cost conditional on output. As a result, input demands satisfy standard marginal productivity conditions. Specifically, cost minimization implies:

$$r^* = \omega_K \frac{Y_t(i)}{K_{Y,t}(i)} - \delta = \alpha_K \frac{A_t}{K_{A,t}} - \delta.$$

Under symmetry, $Y_t(i) = Y_t$ and $K_{Y,t}(i) = K_{Y,t}$, and this simplifies to:

$$r^* = \omega_K \frac{Y_t}{K_{Y,t}} - \delta = \alpha_K \frac{A_t}{K_{A,t}} - \delta.$$

So the marginal product of capital is constant and equal across the two sectors. This implies constant capital–output ratios, and the following growth rate equalities:

$$g^* \equiv g_Y = g_{K_Y}, \tag{3.5}$$

$$g_A = g_{K_A}, \tag{3.6}$$

where g^* is the growth rate of final-goods output, g_{K_Y} is the growth rate of capital in the final-goods sector, and g_A and g_{K_A} are the growth rates of AI output and of capital used in AI production, respectively.

3.2 AI Production in GBGP

The growth rate of AI production is determined by learning-by-doing, recursive AI use, and capital accumulation. Taking logs and growth rates of the AI production function (Equation 3.1) gives:

$$g_A = \phi g_A + \alpha_K g_{K_A} + \alpha_{L^s} g_{L^s} + \alpha_E g_{E_A} + \alpha_A g_{A_A}, \tag{3.7}$$

where, as above, $g_A = g_{K_A}$ denote the growth rates of AI production and of the capital used in its production; and g_{L^s}, g_{E_A}, and g_{A_A} denote the growth rates of the inputs of specialized labor, energy and AI itself into AI production.

We first pare this down by noting that under the Cobb-Douglas structure, cost shares remain constant (and in this case, equal to the distribution parameters). Because the total supplies of labor and of energy are fixed, their prices must rise at the rate of output to preserve these cost shares.

So, the wage of specialized and nonspecialized labor both grow at the rate of output in the GBGP:

$$g_{w^s} = g_{w^n} = g^* \tag{3.8}$$
$$g_{p_E} = g^*. \tag{3.9}$$

And, the wage of specialized labor relative to nonspecialized labor remains constant.

Correspondingly, the quantities allocated to each sector remain constant over time. That is, their growth rates are zero:

$$g_{L_A^s} = g_{L_Y^s} = 0, \tag{3.10}$$
$$g_{E_A} = g_{E_Y} = 0. \tag{3.11}$$

Then, with $g_{K_A} = g_A$ (following Equation 3.6), Equation 3.7 becomes: $g_A = \phi g_A + \alpha_K g_A + \alpha_A g_{A_A}$. This gives us the following expression for the growth of AI production:

$$g_A = \kappa g_{A_A}. \tag{3.12}$$

where:

$$\kappa = \frac{\alpha_A}{1 - \phi - \alpha_K}.$$

This expression relates the growth rate of AI production, g_A, to the growth rate of the recursive AI input, g_{A_A}. And the coefficient, κ, captures the extent to which recursive use amplifies AI output growth through accumulated learning by doing and accumulated capital. The larger the coefficient, the more recursive AI use is amplified in AI production over time.

For consistency with this GBGP, we constrain κ to equal one. Feasibility alone requires that $\kappa \geq 1$. That is, the recursive use of AI must ultimately grow no faster than the rate of AI production. Otherwise, the input of

AI eventually would exceed all of the AI produced. More generally, $\kappa \neq 1$ implies a contradiction with the constant interest rate implied by the constant growth rate of consumption required by a generalized balanced growth path. That is, we impose: $\phi = 1 - \alpha_A - \alpha_K$. So:

$$g_{A_A} = g_A. \tag{3.13}$$

That is, in the GBGP, the recursive use of AI grows at the same rate as the production of AI itself.

With $\kappa = 1$, we can read off the long-run objects in closed form: AI grows faster than final output, the AI price falls every period, real wages grow with output, and the specialized/nonspecialized wage ratio is flat in this benchmark.

3.3 AI Prices and Profits in the GBGP

As above, each final good variety is produced by a monopolistically competitive firm that sets price as a constant markup over marginal cost. Cost minimization implies that the marginal product of AI input equals its marginal cost divided by the final-goods producer's markup. Thus, the price of AI is:

$$p_{A,t} = \mu_Y \cdot \omega_A \cdot \frac{Y_t}{A_{Y,t}} \tag{3.14}$$

Taking logs and differentiating with respect to time yields the growth rate of the price of AI:

$$g_{p_A} = g_Y - g_{A_Y}$$

Combining this with Equation 3.15, we obtain

$$g_{p_A} = \left(\frac{\omega_A}{1 - \omega_K} - 1 \right) g_{A_Y}.$$

Since constant returns to scale implies $\omega_A < 1 - \omega_K$, we have $g_{p_A} < 0$: that is, the price of AI falls continuously in the GBGP.

Because factor shares and markups are constant on the GBGP, per-period AI profits grow at the same rate as final output. Hence the AI profit share is constant and is pinned down by the AI markup, the AI share in final-goods costs, and the steady self-use share.

3.4 Final-Goods Production Growth in GBGP

Since each final good variety is produced using a Cobb-Douglas technology, and—as above—symmetry allows us to treat the growth rate of a representative firm's output as determining the growth rate of aggregate output, its growth rate is:

$$g_Y = \omega_K g_{K_Y} + \omega_{L^s} g_{L^s} + \omega_{L^n} g_{L^n} + \omega_E g_E + \omega_A g_{A_Y}, \tag{3.15}$$

As above, capital grows at the same rate as output under the constant GBGP interest rate. And since labor and energy are fixed in supply, cost minimization requires their allocation across sectors to remain constant over time. So, $g_{L^s} = g_{L^n} = g_E = 0$; and Equation 3.15 becomes $g^* = \omega_K g^* + \omega_A g_{A_Y}$, or:

$$g^* \equiv g_Y = \frac{\omega_A}{1 - \omega_K}\, g_{A_Y} \tag{3.16}$$

The growth rate of final output is proportional to the growth of AI used in its production, where the proportionality factor depends simply on the distribution parameters for capital and AI, ω_A and ω_K; and—since constant returns to scale impose that their sum is less than one—the proportionality factor is less than one. That is, final output grows more slowly that the AI input.

To summarize: in this GBGP, AI production grows faster than final output, so the price of AI falls continuously. Real wages of both specialized and nonspecialized labor rise in step with overall output, while their relative wage remains constant. Finally, the real interest rate is pinned down by household preferences.

In the Cobb–Douglas benchmark with fixed labor and energy and no exogenous Z_t growth, the economy converges to a constant K/Y ratio and a constant level of output. Consumption growth likewise converges to zero, so the Euler equation implies $r_t \to r^\star = \beta^{-1} - 1$. Hence, if the pre-AI steady state also featured zero growth under the same β, the real interest rate returns to its original level asymptotically.

This interplay makes the Cobb–Douglas case a useful benchmark, clarifying how learning and recursion sustain balanced growth, and setting the stage for the more flexible CES setting.

4 CES Production: Simulation

In this section, we relax the unit elasticity assumption and specify both AI and final-goods production as CES. Keeping period-by-period constant returns and the market structure intact yields a closed-form expression for the effective demand elasticity faced by the AI monopolist, which we use to construct the simulations that follow. We first lay out the CES setup, then we turn to the simulations.

4.1 CES Specification

We begin by specifying the CES production function for AI.

$$A_t = (A_{t-1})^{\phi}\left[\alpha_K^{\frac{1}{\sigma_A}} K_{A,t}^{\rho_A} + \alpha_{L^s}^{\frac{1}{\sigma_A}} (L_{A,t}^s)^{\rho_A} + \alpha_E^{\frac{1}{\sigma_A}} E_{A,t}^{\rho_A} + \alpha_A^{\frac{1}{\sigma_A}} (A_{A,t})^{\rho_A}\right]^{\frac{1}{\rho_A}}. \quad (4.1)$$

As before: $\phi \in (0,1)$ is a parameter that governs learning by doing based on the producer's prior output; $\alpha_A \in (0,1)$ is the distribution parameter of AI used recursively as an input into AI; from constant returns to scale, $\alpha_K + \alpha_{L^s} + \alpha_E + \alpha_A = 1$; and $\sigma_A = \frac{1}{1-\rho_A}$ governs the degree of input substitutability.[12]

In final-goods production, we again assume a continuum of differentiated final goods indexed by $i \in [0,1]$. Each variety $Y_t(i)$ is now produced by a firm using a CES technology, with a fixed cost (f) and free entry:

$$Y_t(i) = Z_t\left[\omega_K^{\frac{1}{\sigma_Y}} K_{Y,t}^{\rho_Y} + \omega_{L^s}^{\frac{1}{\sigma_Y}} (L_{Y,t}^s)^{\rho_Y} + \omega_{L^n}^{\frac{1}{\sigma_Y}} (L_{Y,t}^n)^{\rho_Y} + \omega_E^{\frac{1}{\sigma_Y}} E_{Y,t}^{\rho_Y} + \omega_A^{\frac{1}{\sigma_Y}} (A_{Y,t})^{\rho_Y}\right]^{\frac{1}{\rho_Y}} \quad (4.2)$$

where $\omega_K + \omega_{L^s} + \omega_{L^n} + \omega_E + \omega_A = 1$; Z_t denotes total factor productivity in the final-goods sector and is assumed to grow exogenously at the rate g_z; and and ρ_Y governs the degree of input substitutability.

Profits for variety i in period t are then given by:

[12] In the numerical work below, we generalize the production function to allow for the learning by doing to change over time.

$$\begin{aligned}\Pi_{Y,t}(i) = p_t(i)Y_t(i) \\ - w_t^s L_{Y,t}^s(i) - w_t^n L_{Y,t}^n(i) \\ - r_t K_{Y,t}(i) - p_{E,t}E_{Y,t}(i) - p_{A,t}A_{Y,t}(i) - f\end{aligned} \tag{4.3}$$

Combining the resulting final-goods producers' production decisions with the household's consumption preferences (unchanged from Equations 2.12 and 2.13) gives us the monopolist's effective elasticity of demand.

As mentioned in section 2, the effective elasticity has two parts. First, because final-goods producers use AI in CES production with elasticity σ_Y, their AI input demand, holding their output fixed, has elasticity σ_Y. At the same time, the final-goods producers face CES demand from households with elasticity ε, so their own output also adjusts in response to the final-goods prices they set. Since the AI monopolist internalizes both the input demand response and the induced change in output of final goods, it faces an effective elasticity, ϵ, that reflects both underlying elasticities.[13] Specifically,

$$\epsilon \;=\; \sigma_Y(1 - s_{A,t}) \;+\; \varepsilon\, s_{A,t}, \tag{4.4}$$

where $s_{A,t}$ is the AI cost share. This is the effective elasticity in the optimal AI pricing rule in Section 2.1.2. It (along with $s_{A,t}$) is derived in the appendix.

With the effective demand elasticity in hand, we can now turn to the parameterization and numerical simulations.

4.2 Simulation

The CES formulation nests the Cobb–Douglas case as the special case with $\sigma_A = \sigma_Y = 1$. Here, we simulate the special case to examine its dynamics. Then, we depart from unit elasticity to allow us to examine how input substitutability affects the AI production, pricing, and income distribution.

[13]The monopolist sets a single price based solely on final-goods producers' demand. While the monopolist also uses AI recursively in its own production (with market clearing requiring $A_t = A_{Y,t} + A_{A,t}$), this internal use represents optimal input choice at the given price rather than independent demand that affects pricing decisions.

4.2.1 Parameterization

The parameter values for all the simulations are chosen to be broadly consistent with standard macroeconomic practice and with the stylized facts of the AI sector. Preferences are set to yield plausible consumption–savings behavior; production parameters are drawn from factor income shares; and elasticities are varied over a range to capture cases of complementarity and substitutability. Key parameter values are given in Table 1. Full parameter values are reported in Appendix B.

Table 1: Key Parameters Used in Simulations

Symbol	Description	Values
ϕ	Learning	0.150
α_A	Recursive AI distribution parameter	0.425
α_K	Capital distribution parameter in AI production	0.425
σ_A	CES elasticity, AI production	0.900, 1.000, 1.100
ω_A	AI distribution parameter in final-goods production	0.025
ω_K	Capital distribution parameter in final-goods production	0.350
σ_Y	CES elasticity, final-goods production	0.100–2.000
β	Discount factor	0.980
σ_c	CRRA coefficient (risk aversion)	2.000
ε	Elasticity of substitution across varieties	6.000

The distribution parameters are chosen to reflect suggestive evidence regarding the shares of specialized labor. The share of specialized labor (with the kind of qualifications at the major AI labs in the U.S. and elsewhere) in the workforce appears to be something less than about 0.01; and workers with specialized AI development knowledge currently are said to earn roughly five times what educated non-specialist workers earn, and about ten times that of nonspecialists overall.[14]

4.2.2 Dynamic Simulation Results

We begin the simulations with the Cobb-Douglas specification, $\sigma_A = \sigma_Y = 1$. These are shown in figures 1a-1c, at the end of the paper.

Figure 1a illustrates the results for the AI production. As shown in the top row of Figure 1a, AI production accelerates early then tapers off, and the use of energy and specialized labor spikes before reverting toward its initial starting level. In the second row, we can see that the price

[14]See, for example, OECD (2023); U.S. Bureau of Labor Statistics (2024); Insider (2025); H1B Data (2025); PYMNTS (2025).

of AI drops and the recursive use of AI rises early before tapering off as production tapers off.

Figure 1b gives the corresponding results for final goods production. As shown in the top row, final production rises more gradually than AI production, and—since total energy and specialized labor are fixed–their use in final goods production counterbalances their use in AI production: first falling sharply before recovering substantially. As shown in the second row, unspecialized labor use is constant (since it is in fixed supply and not used in AI production), but the use of AI rises rapidly–with the production of AI seen in Figure 1a.

Figure 1c illustrates how prices behave.As shown in the top row, real wages for both specialized and nonspecialized labor rise; and the relative specialized-to-nonspecialized labor spikes before returning toward its original level. The second row shows a gradual, but permanent rise in the price of energy. We also see a temporary rise in the interest rate that corresponds to the build up of capital in both sectors. Profit, like wages also rise.

Figures 2a-2c depict the results when inputs are complements in AI production ($\sigma_A = 0.9$), with a range of values of for σ_Y. As shown in the top row of figure 2a, AI production initially rises rapidly, then tapers off; however it rises more when inputs are highly substitutable in final-goods production. The use of energy and specialized labor follow a corresponding pattern. As shown in the second row (and as in the Cobb-Douglas case), the price of AI drops initially, then levels off, and AI's recursive use rises with AI production.

Figure 2b gives the results for final-goods production when inputs are complements in AI production. While final goods production rises more when its inputs are highly substitutable, the use of energy and specialized labor again again counterbalance their use in AI: they decline before tapering off— all the more so as σ_Y rises.

Figure 2c gives the results for prices. As shown in the top row, real wages rise, as expected final; but the specialized-to-nonspecialized wage premium rises with the degree of input substitutability in final goods. The second row, shows that energy prices, the interest rate, and AI profits, like the wage premium, also rise with final goods input substitutability. The bottom row shows that the AI markup and the corresponding price move inversely with the final goods' input substitutability.

Figures 3a-3c give the results when inputs in AI production are substitutes, with figure 3a giving the results for AI production. As shown

in the top row and in the first panel and last panels in the bottom row, all aspects of AI production are higher when final goods inputs are substitutable, while the price–shown in the middle panel of the bottom row–falls correspondingly more.

Final goods production, shown in Figure 3b, rises most when its inputs are substitutable. Its use of AI rises correspondingly; and, again, the use of energy and specialized labor counterbalance their use in AI.

Figure 3c gives the pricing results. As in the case of AI input complements, real wages–shown in the top row, both rise. And, while the specialized-to-nonspecialized wage premium rises when σ_Y is greater than one, it falls when the elasticity equals one or less, and it is lowest when the inputs in final goods production are complements. As shown in the second row, profits follow the same pattern as specialized labor: notably, profits fall when final goods inputs are complements. The price of AI, shown in the last row, falls in all cases, but it falls most when final good inputs are substitutes.

Across cases, higher substitutability in final-goods production strengthens pass-through, lowers the AI markup and price, expands quantity, and raises AI profits. And, when final-goods inputs are complements, the opposite holds—markups are higher but quantity and external sales respond weakly, so AI profits are lower. When AI inputs are more substitutable, these patterns are amplified but not reversed. The specialized–to–nonspecialized wage premium rises when $\sigma_Y > 1$ and is lowest when $\sigma_Y < 1$.

With complements in final-goods production, adoption is learning-first and capital-light: the AI producer still lowers price to scale and learn, but a larger share of output is allocated to self-use, so external sales and profits grow more slowly.

5 Conclusion

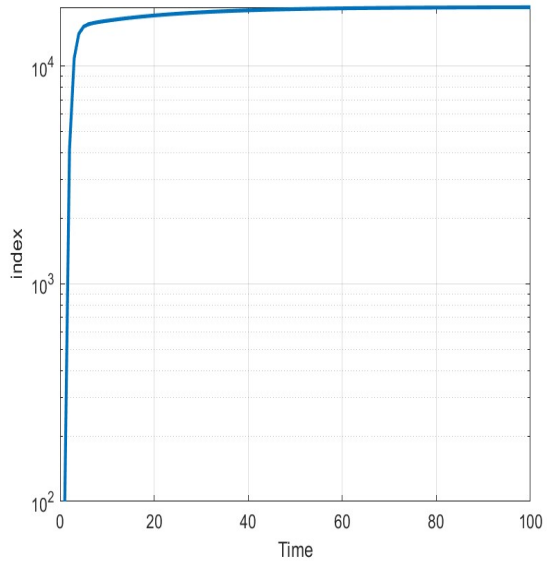
This paper establishes four results. First, recursive self-use of AI, together with learning by doing, reshapes monopoly pricing and scale: early price cuts raise future productivity, and self-use amplifies the returns to learning. Second, with CES production we derive a closed-form effective elasticity that links markups and the AI price path to the endogenously evolving AI cost share. Third, under Cobb–Douglas we solve for a GBGP in closed form. The solution pins down the conditions for stable, non-explosive growth, yields long-run growth rates as functions of primitives,

and provides the benchmark for the simulations.

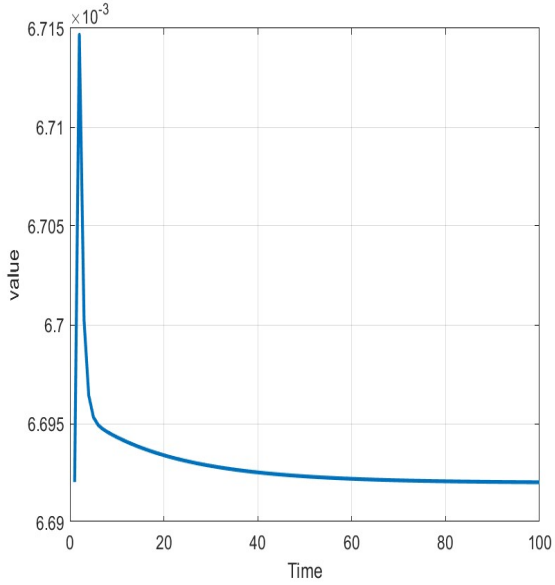
Finally, the CES simulations connect adoption and distribution to elasticities: when final-goods inputs are complements, AI adoption is learning-first and capital-light before scaling; when inputs are substitutes, adoption is scale-first and the two-phase pattern attenuates. On the distribution side, the specialized–to–nonspecialized wage premium is lowest when final-goods inputs are complements and rises when inputs are substitutes. Greater substitutability in AI production amplifies these patterns without changing their sign.

Figure 1a: AI Sector: $\sigma_A = 1$ and $\sigma_A = 1$

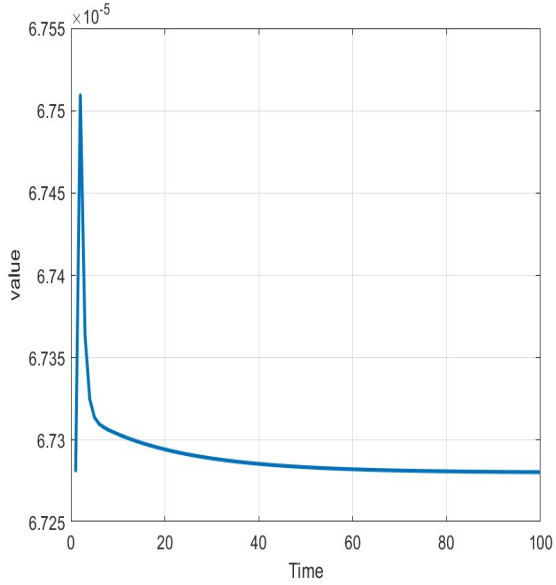
AI Production



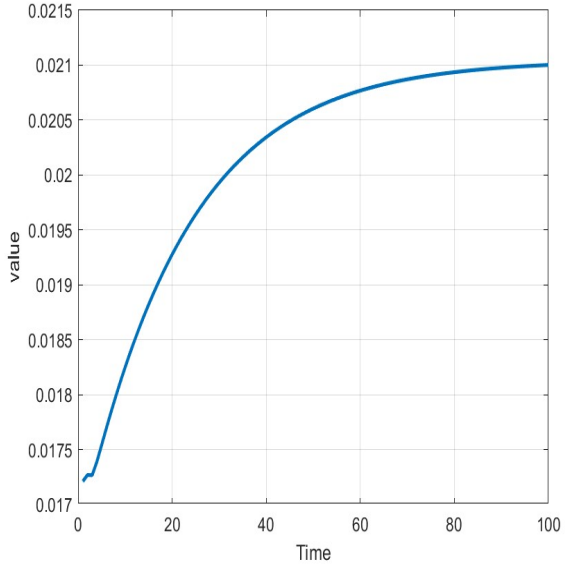
AI Energy use



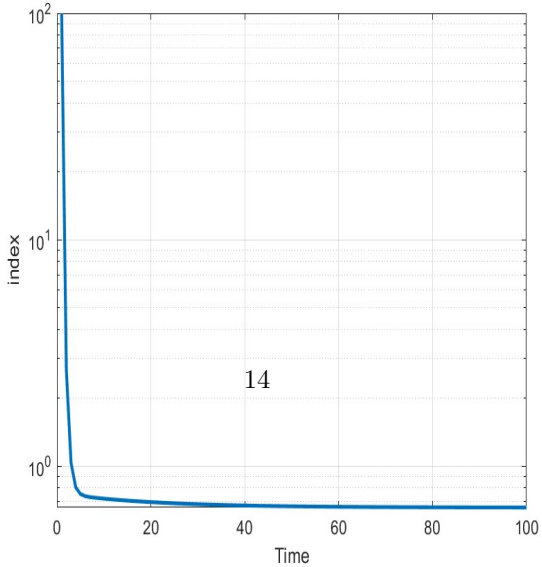
L^s



AI Capital



AI Price



AI in AI

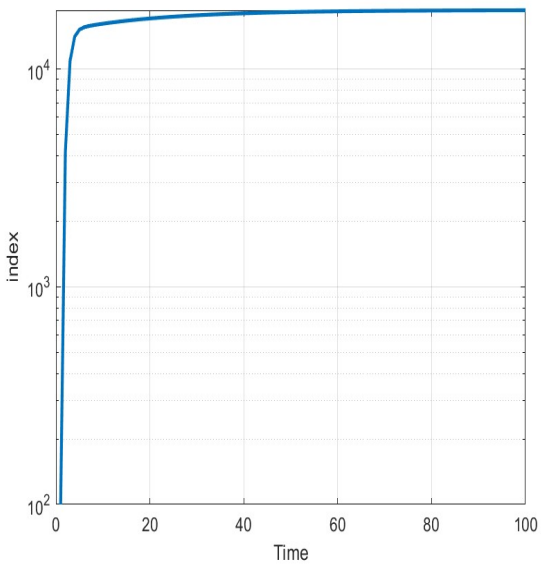
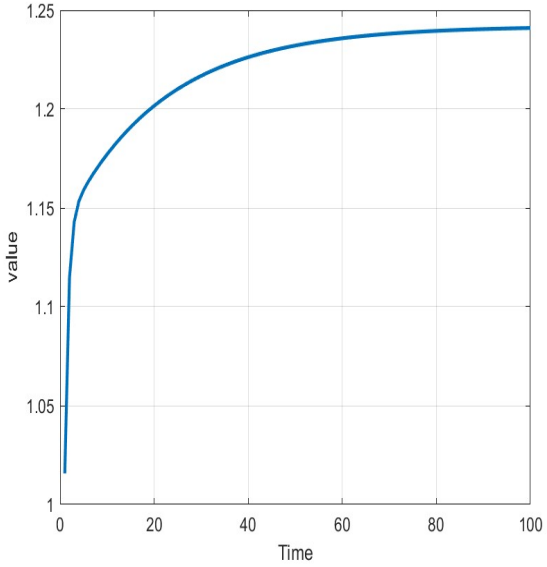
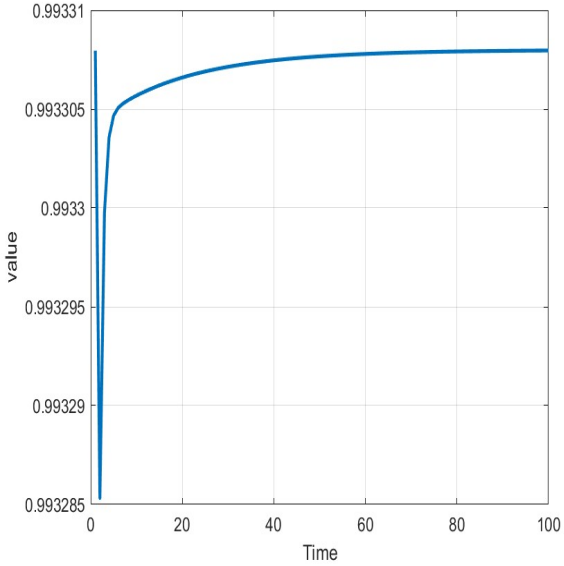


Figure 1b: Final goods: $\sigma_A = 1$ and $\sigma_Y = 1$

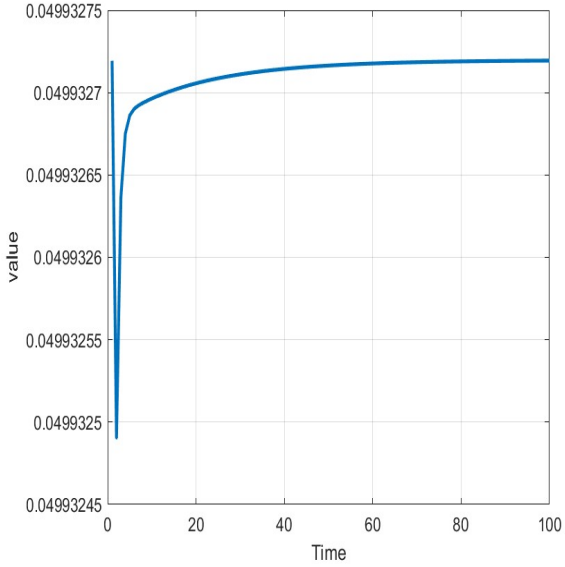
Final Production



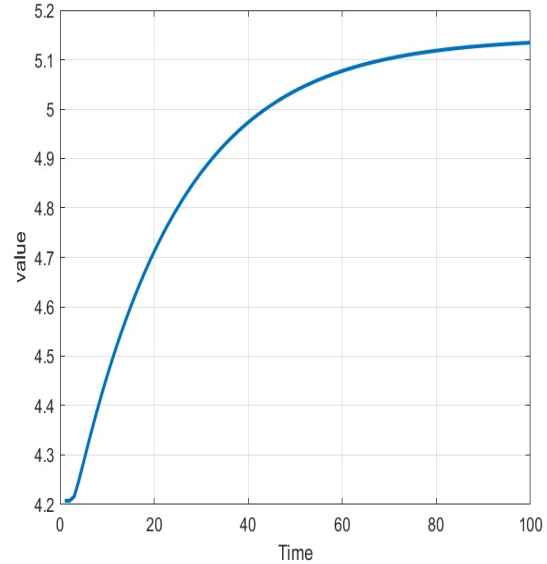
Energy use



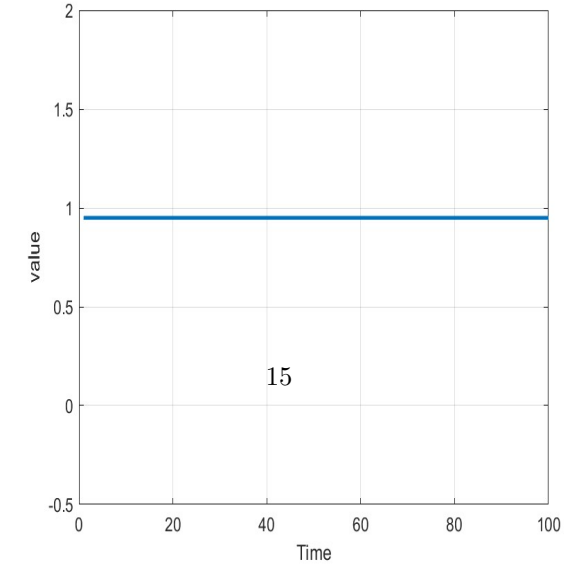
L^s



Capital



L^N



AI in final

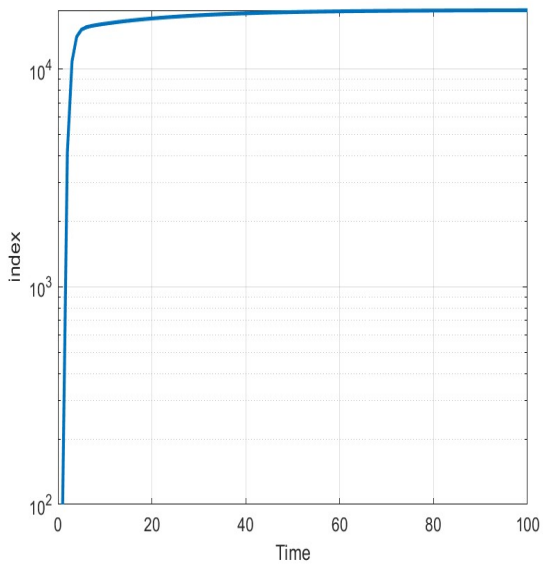
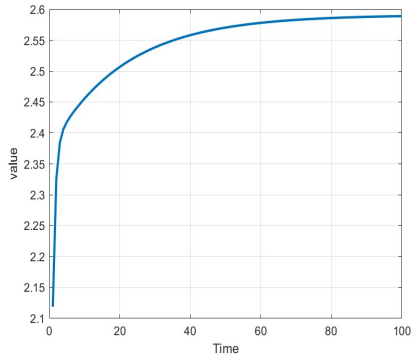
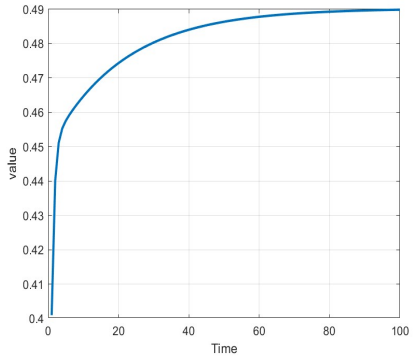


Figure 1c: Results for relative prices: $\sigma_A = 1$ and $\sigma_Y = 1$

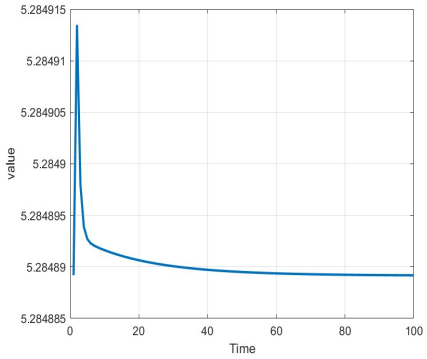
real wage: L^s



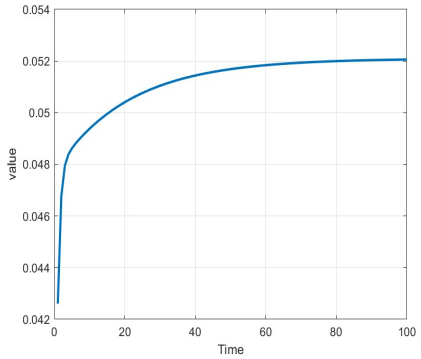
real wage: L^n



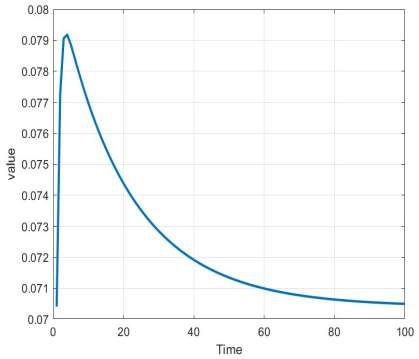
wage premium $(\frac{w}{w^u})$



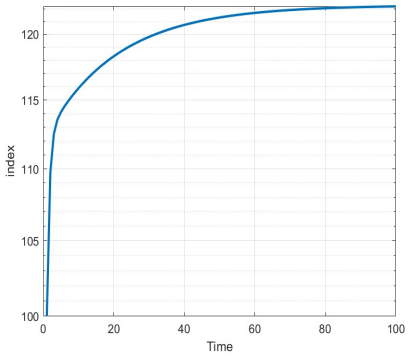
price for energy



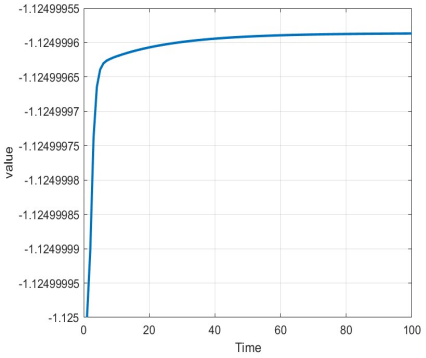
interest rate



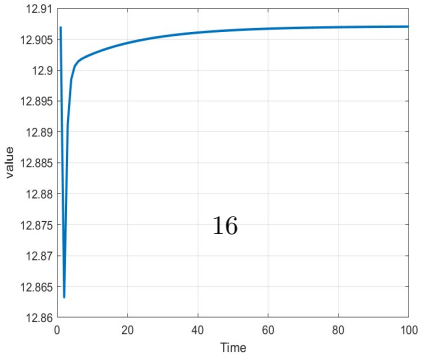
profits AI



AI price elasticity $(\epsilon_{A,t})$



AI mark-up



Price of AI

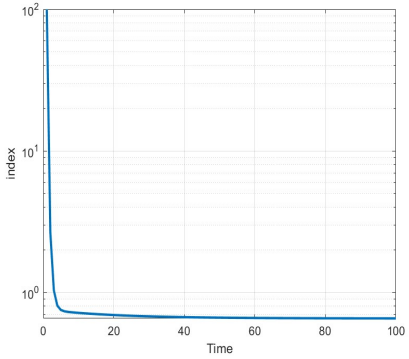


Figure 2a: AI Sector: $\sigma_A = 0.9$

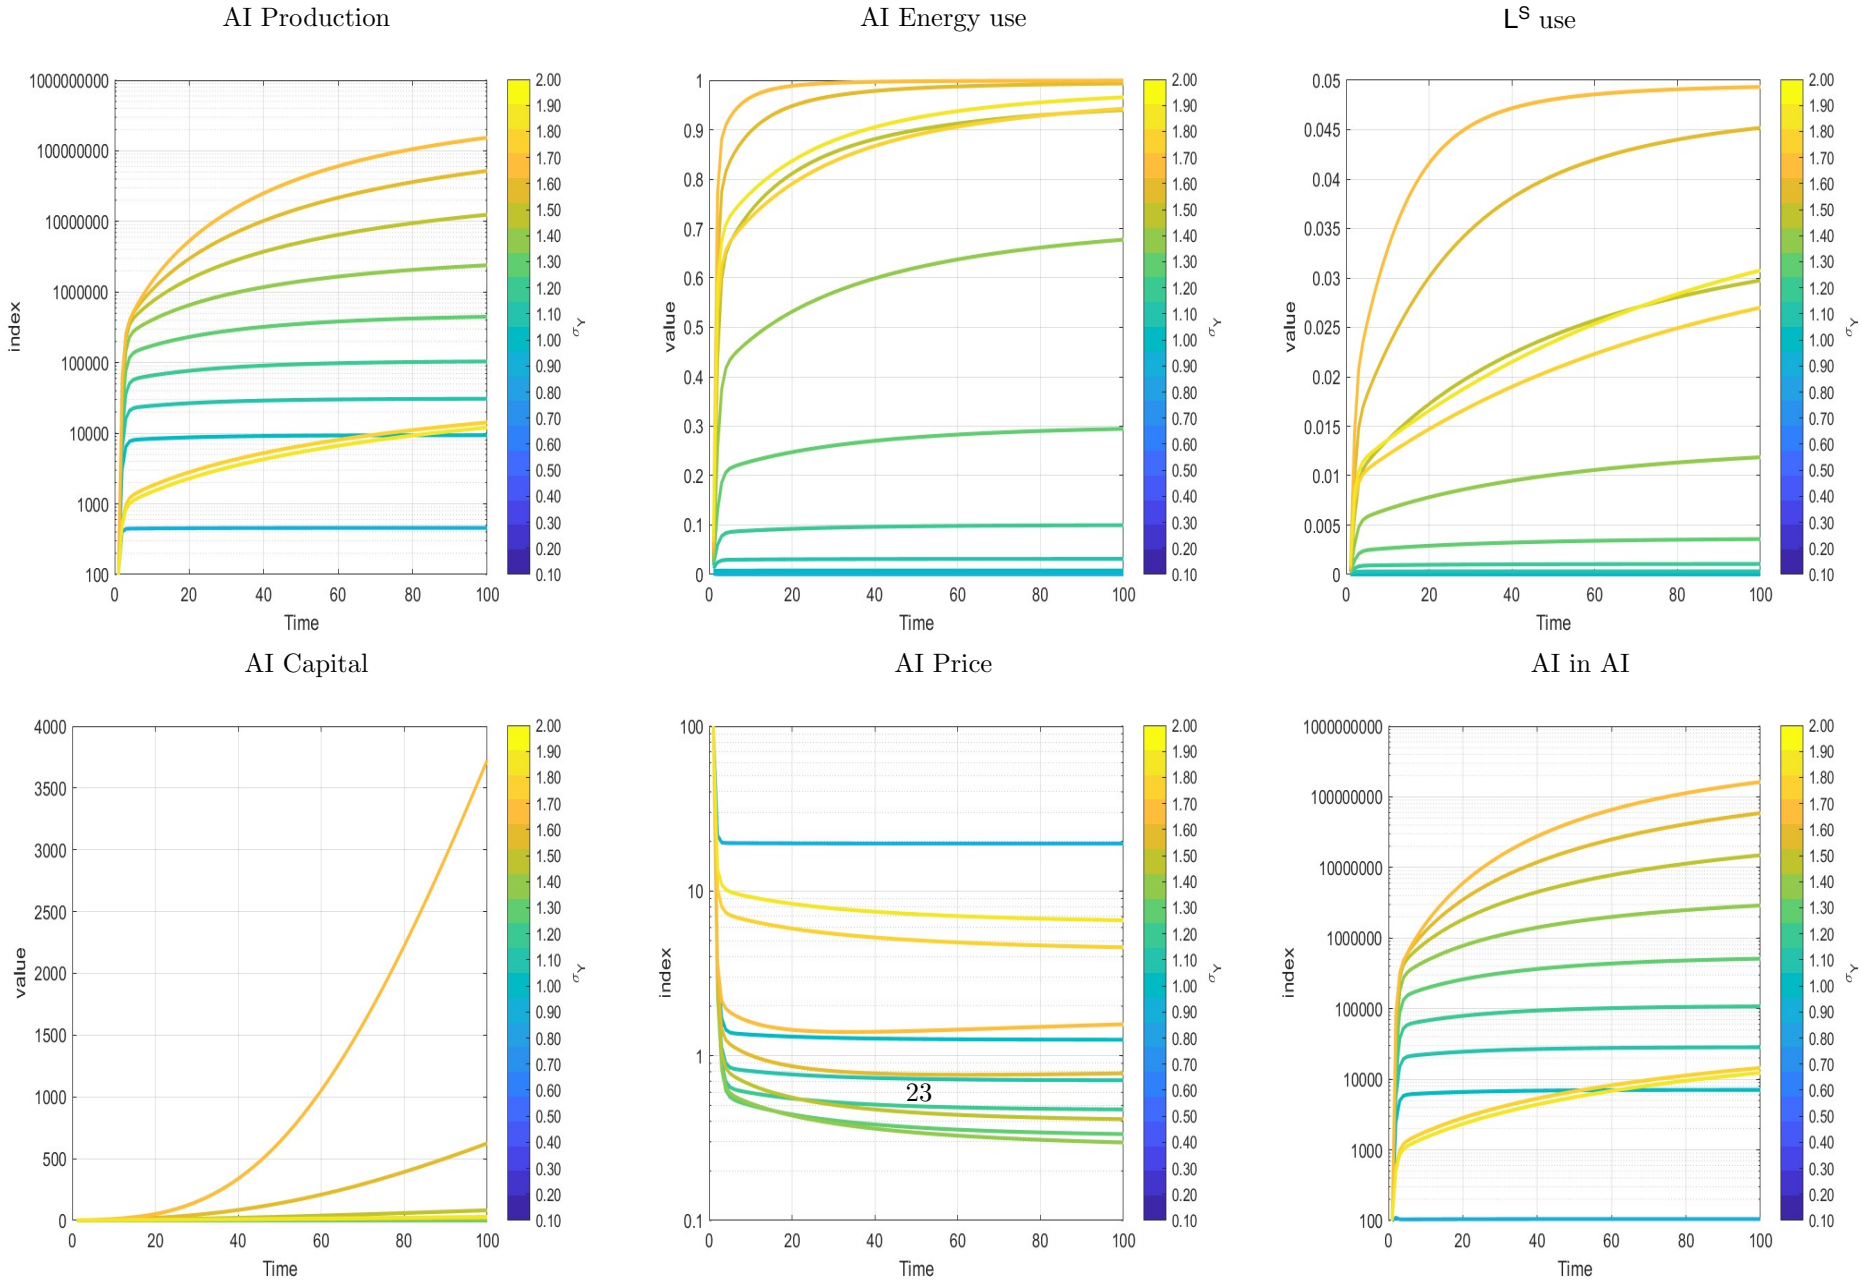


Figure 2b: Final goods: $\sigma_A = 0.9$

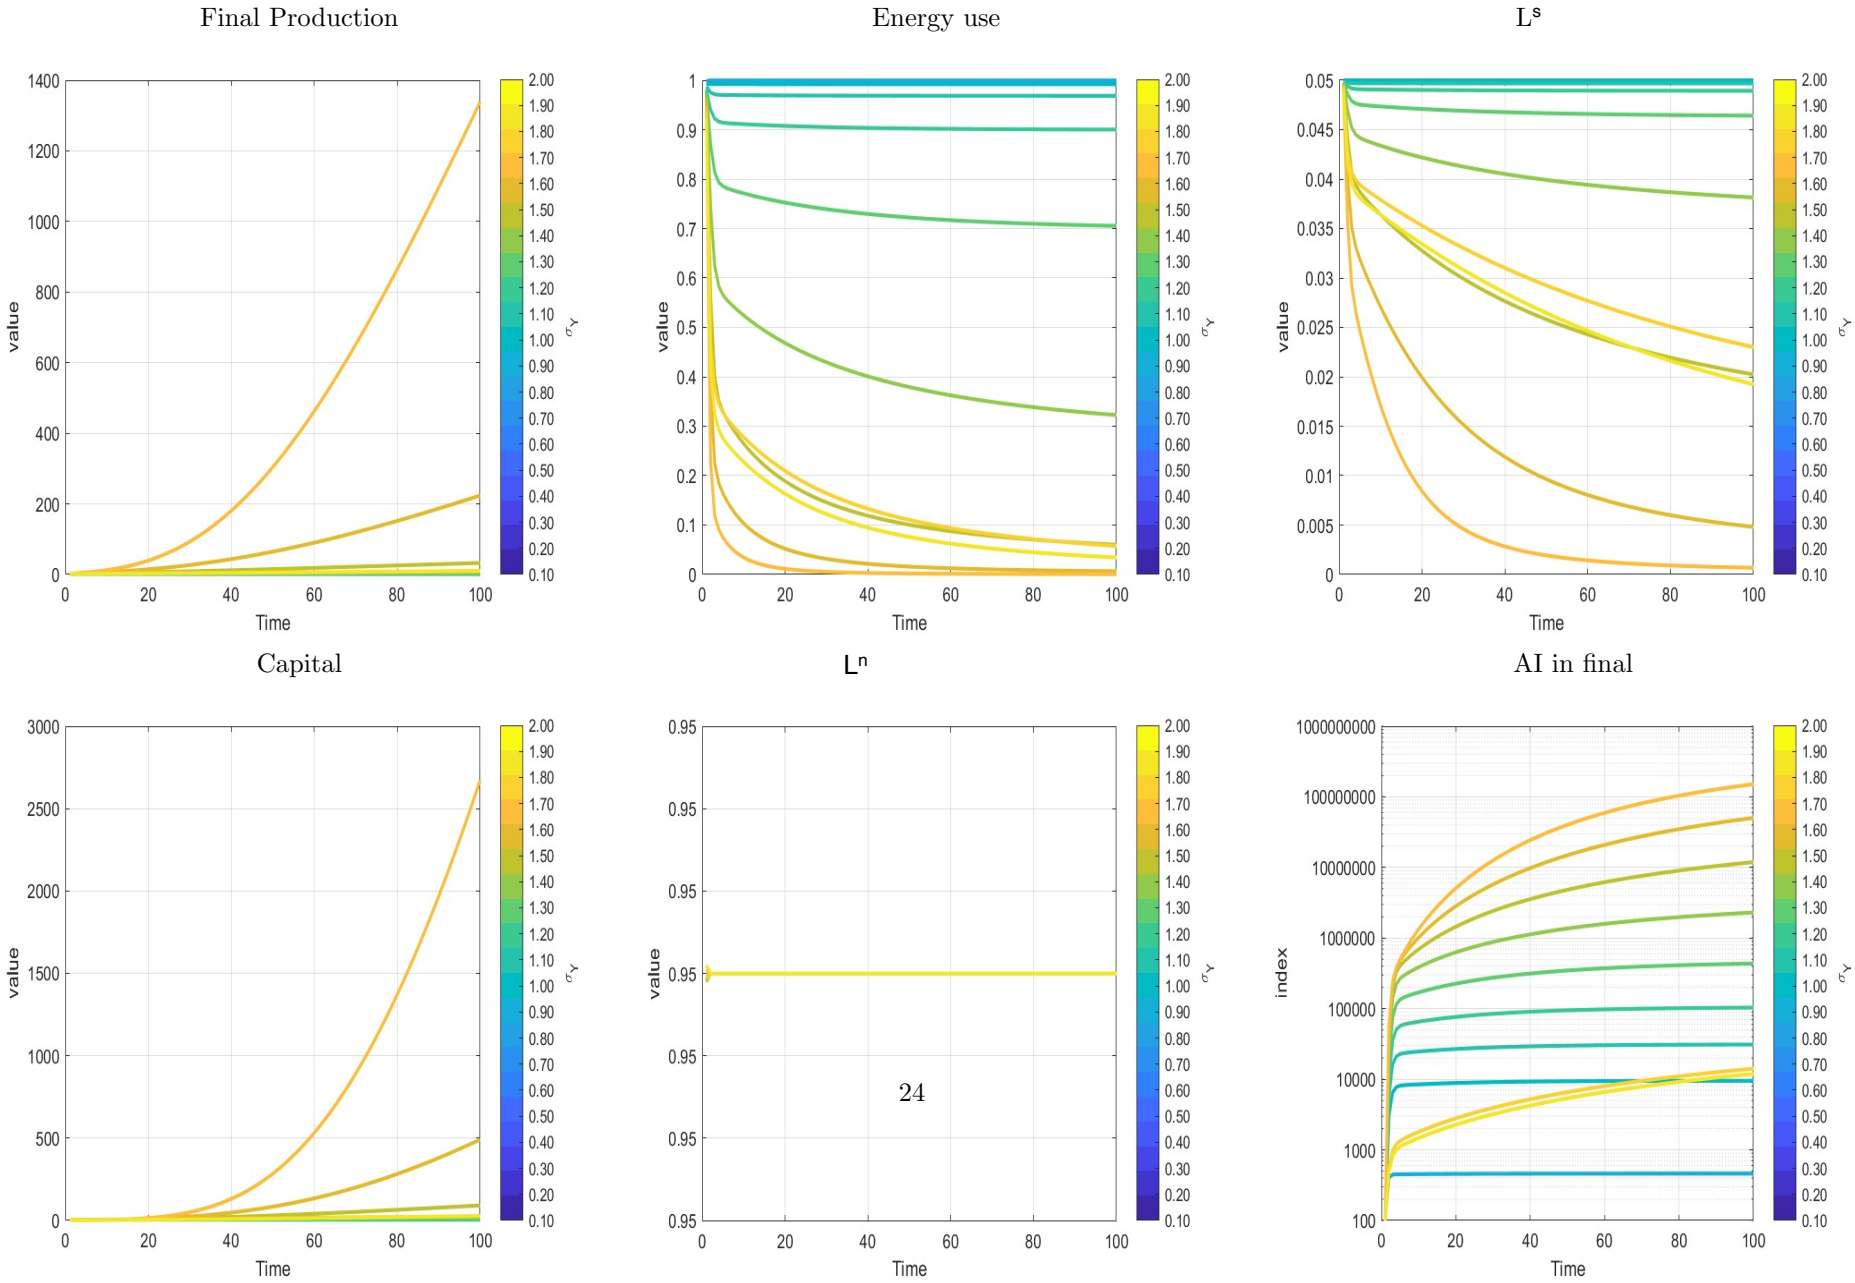
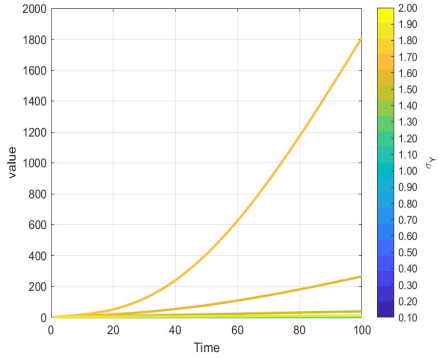
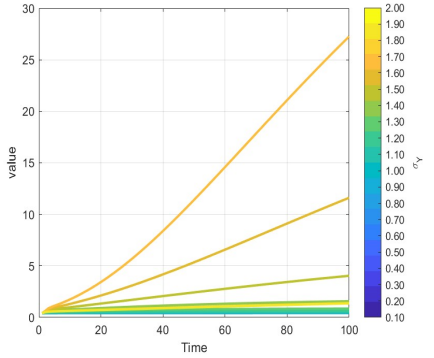


Figure 2c: Relative prices: $\sigma_A = 0.9$

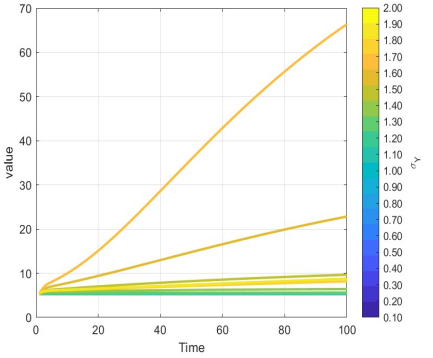
real wage: L^s



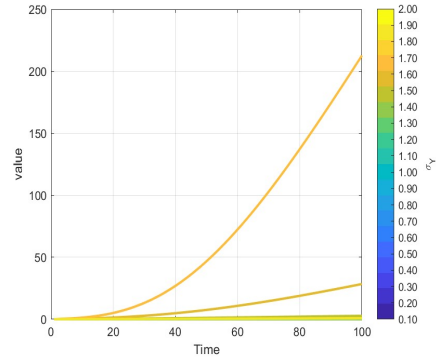
real wage: L^n



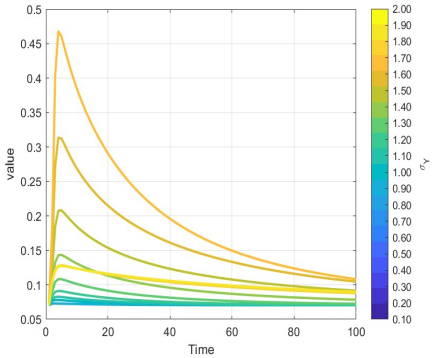
wage premium $(\frac{w}{w^u})$



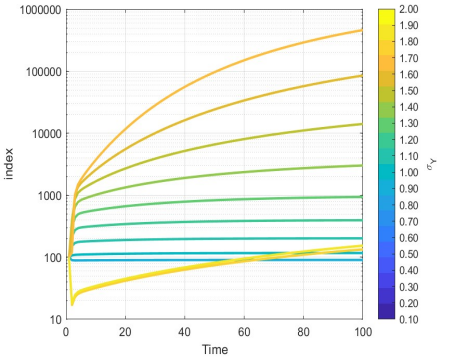
price for energy



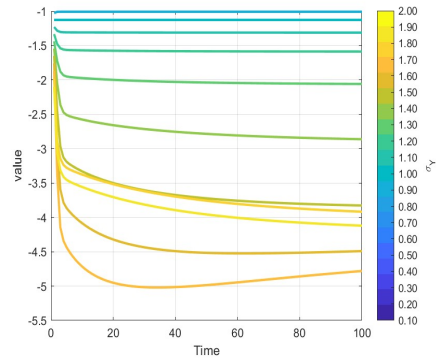
interest rate



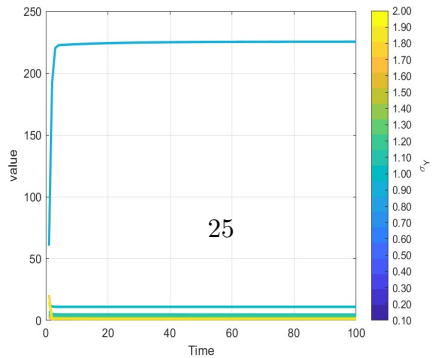
profits AI



AI price elasticity $(\epsilon_{A,t})$



AI mark-up



Price of AI

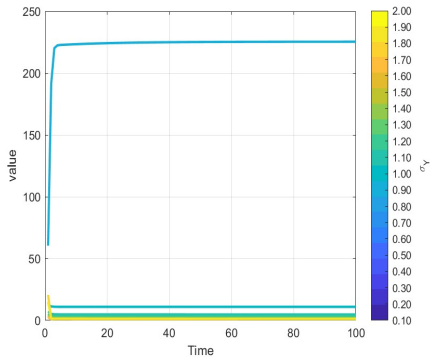
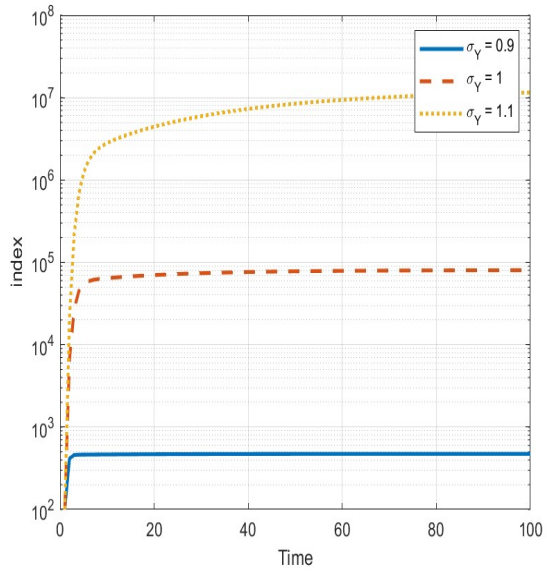
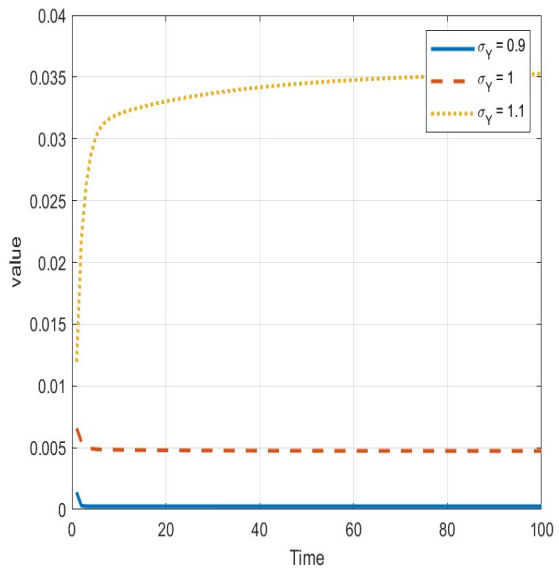


Figure 3a: AI Sector: $\sigma_A = 1.1$

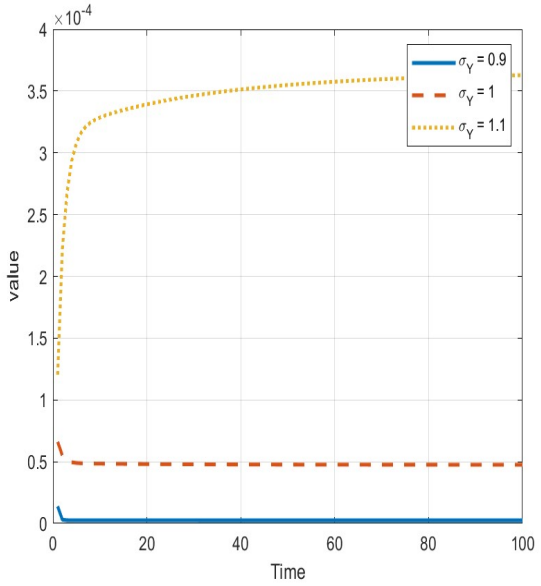
AI Production



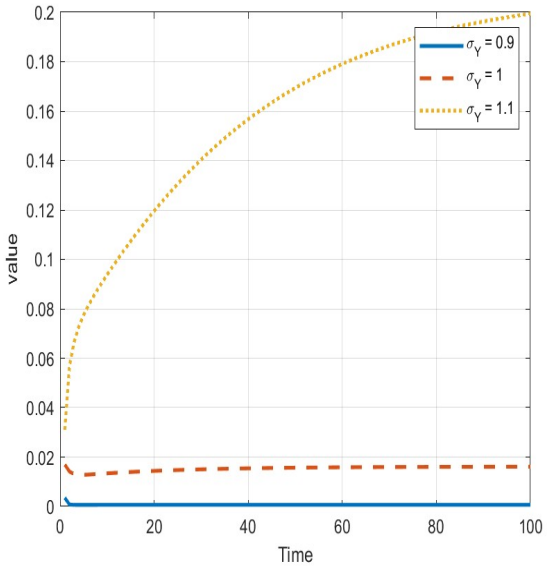
AI Energy use



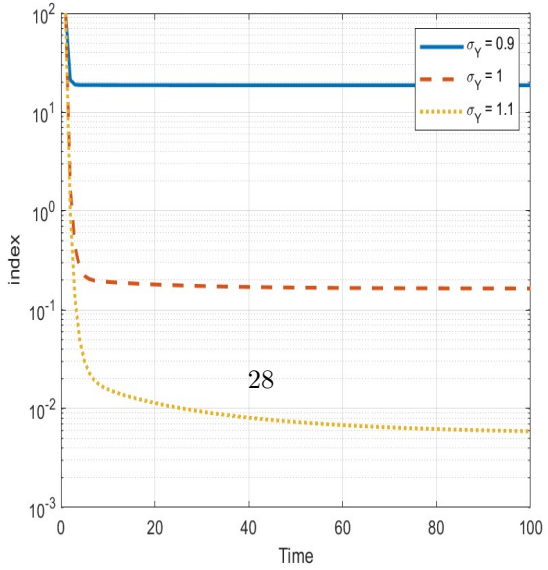
L^S use



AI Capital



AI Price



AI in AI

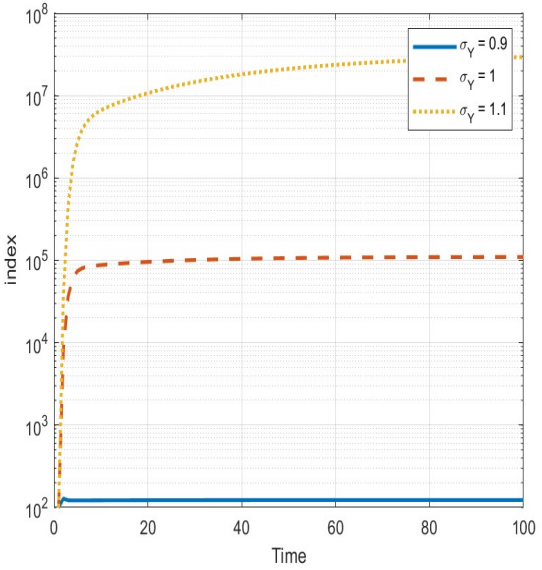
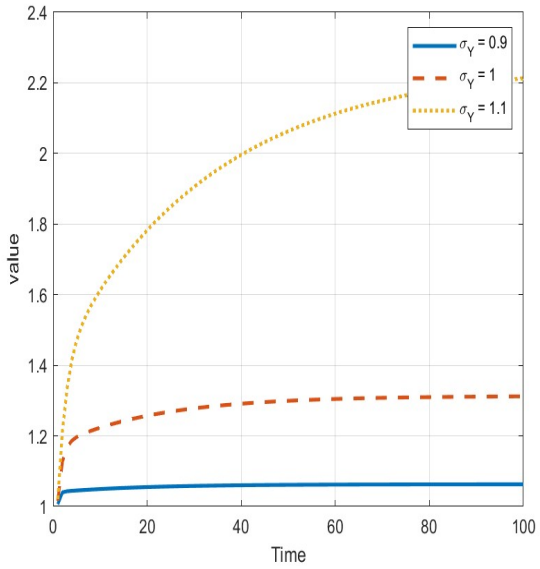
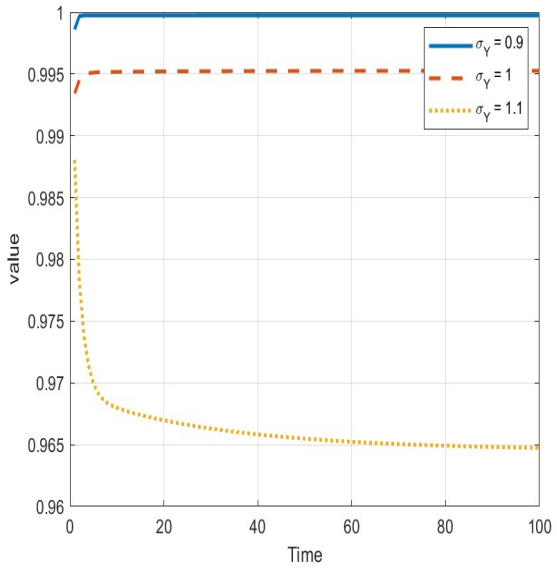


Figure 3b: Final goods: $\sigma_A = 1.1$

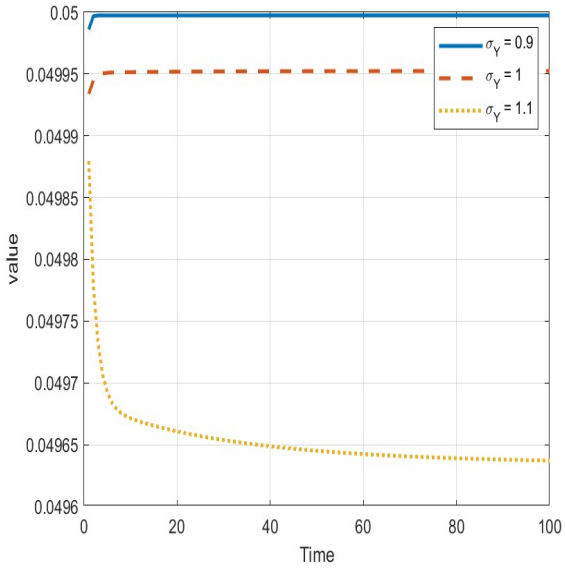
Final Production



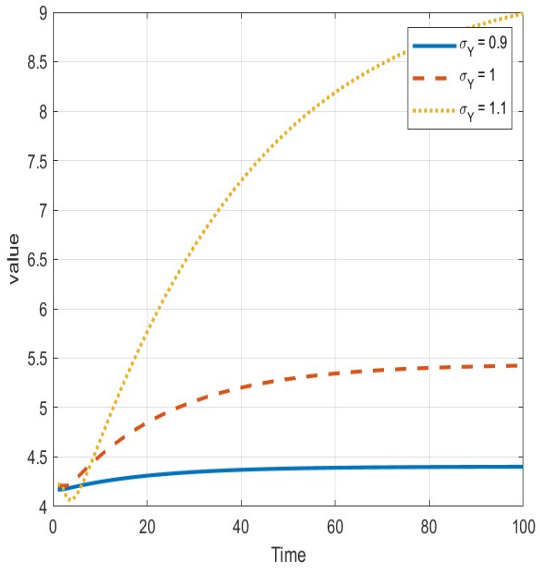
Energy use



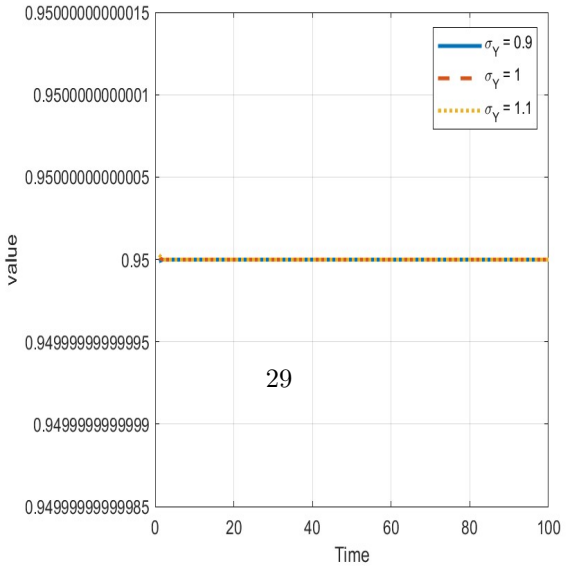
L^s



Capital



L^n



AI in final

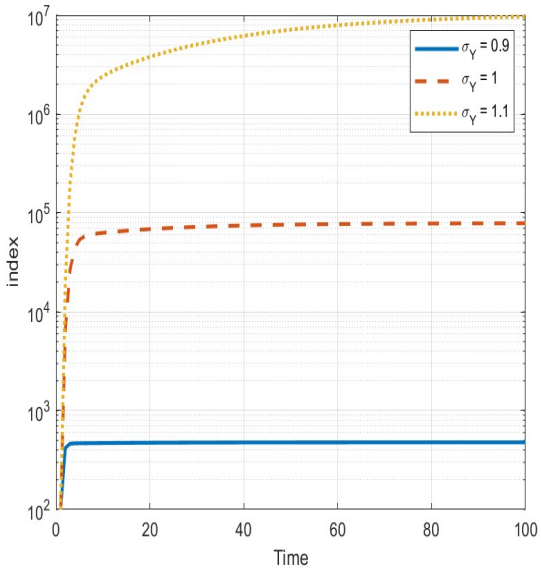
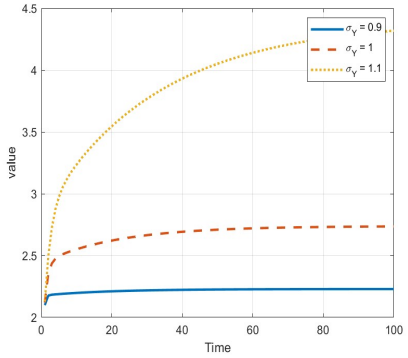
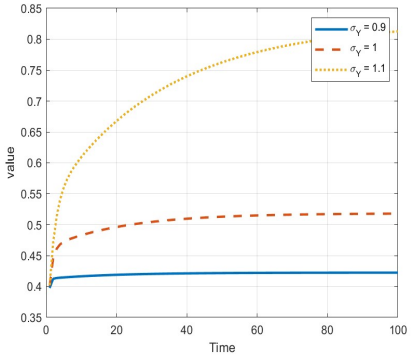


Figure 3c: Relative prices: $\sigma_A = 1.1$

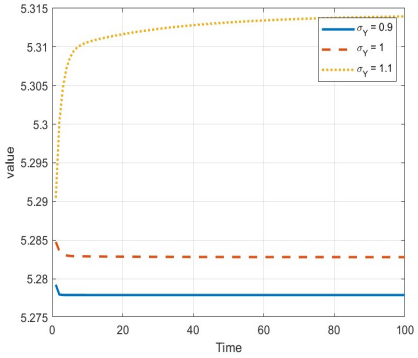
real wage:L^s



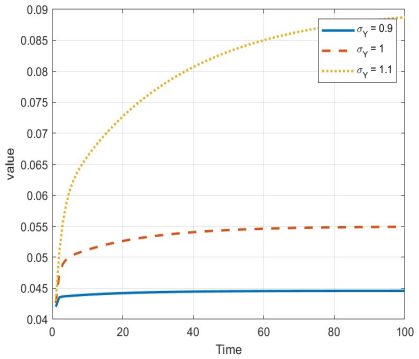
real wage: L^n



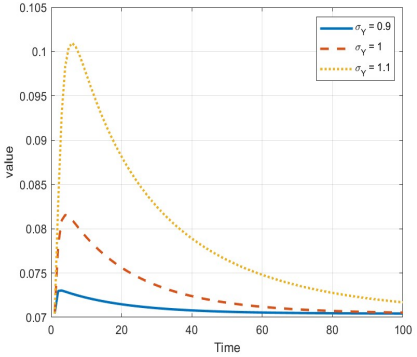
wage premium $(\frac{w}{w^u})$



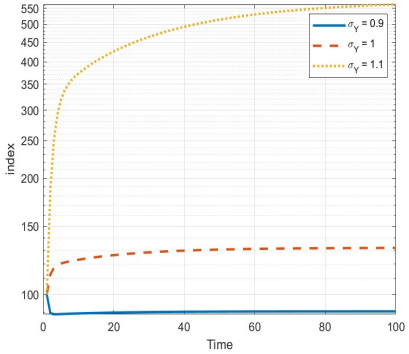
price for energy



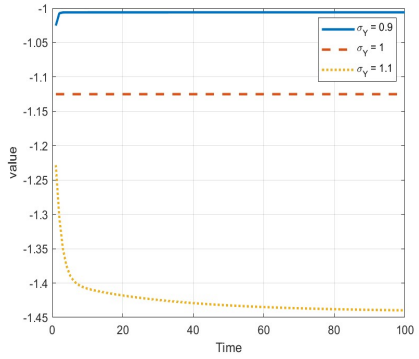
interest rate



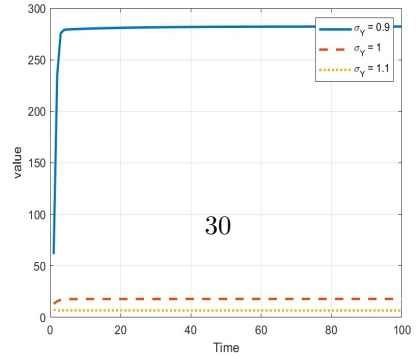
profits AI



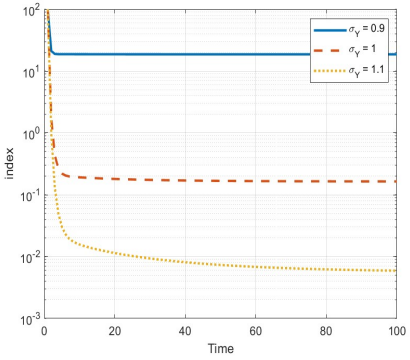
AI price elasticity $(\epsilon_{A,t})$



AI mark-up



Price of AI



Appendices

A Effective Elasticity of AI Demand

This appendix derives the elasticity of AI demand faced by the AI monopolist.

Markup in the Final-Goods Sector

Setting $Z_t = 1$, each final-goods variety i is produced with a CES technology:

$$Y_t(i) \;=\; \left[\sum_{j\in\{K,L^s,L^n,E,A\}} \omega_j^{1/\sigma_Y}\, x_{j,t}(i)^{\rho_Y}\right]^{1/\rho_Y}, \qquad \sigma_Y \equiv \frac{1}{1-\rho_Y}. \tag{A.1}$$

The unit cost index is:

$$\mathcal{U}_t \;\equiv\; \left(\sum_j \omega_j\, p_{j,t}^{1-\sigma_Y}\right)^{\frac{1}{1-\sigma_Y}}. \tag{A.2}$$

Under symmetry in the final-goods sector and CES demand over varieties with elasticity $\varepsilon > 1$, the final-goods markup is

$$\mu_Y \;\equiv\; \frac{\varepsilon}{\varepsilon-1}, \qquad \text{and} \quad p_t(i) \;=\; P_t \;=\; \mu_Y\, \mathcal{U}_t. \tag{A.3}$$

FOCs and Factor Demands

Each final-goods producer minimizes variable cost $\sum_j p_{j,t}x_{j,t}(i)$ subject to (A.1). The Lagrangian is

$$\mathcal{L} \;=\; \sum_j p_{j,t}\, x_{j,t}(i) \;+\; \lambda_t\left[Y_t(i) - \left(\sum_j \omega_j^{1/\sigma_Y} x_{j,t}(i)^{\rho_Y}\right)^{1/\rho_Y}\right].$$

FOC w.r.t. $x_{j,t}(i)$:

$$p_{j,t} \;-\; \lambda_t \left(\sum_k \omega_k^{1/\sigma_Y} x_{k,t}(i)^{\rho_Y}\right)^{\frac{1}{\rho_Y}-1} \omega_j^{1/\sigma_Y} x_{j,t}(i)^{\rho_Y-1} \;=\; 0. \tag{A.4}$$

Taking ratios across j, k and using the production function yields the factor demands:

$$x_{j,t}(i) \;=\; \omega_j \left(\frac{p_{j,t}}{\mathcal{U}_t}\right)^{-\sigma_Y} Y_t(i). \tag{A.5}$$

Let $s_{j,t}$ denote the cost share of input j in final-goods production:

$$s_{j,t} \;\equiv\; \frac{p_{j,t}x_{j,t}(i)}{\mathcal{U}_t Y_t(i)} \;=\; \frac{\omega_j\, p_{j,t}^{1-\sigma_Y}}{\sum_k \omega_k\, p_{k,t}^{1-\sigma_Y}} \;=\; \frac{\partial \ln \mathcal{U}_t}{\partial \ln p_{j,t}}. \tag{A.6}$$

Applying (A.5) to $j = A$ and aggregating with symmetry gives:

$$A_{Y,t} \;=\; \omega_A \left(\frac{p_{A,t}}{\mathcal{U}_t}\right)^{-\sigma_Y} Y_t \;=\; \omega_A\, \mathcal{U}_t^{\sigma_Y - 1}\, p_{A,t}^{-\sigma_Y}\, TC_t, \qquad TC_t \equiv \mathcal{U}_t Y_t. \tag{A.7}$$

How TC_t Responds to $p_{A,t}$

From $TC_t = \mathcal{U}_t Y_t$,

$$\frac{d\ln TC_t}{d\ln p_{A,t}} \;=\; \underbrace{\frac{d\ln \mathcal{U}_t}{d\ln p_{A,t}}}_{=\, s_{A,t}} + \frac{d\ln Y_t}{d\ln p_{A,t}}. \tag{A.8}$$

By (A.3), $P_t = \mu_Y \mathcal{U}_t$, so

$$\frac{d\ln P_t}{d\ln p_{A,t}} \;=\; \frac{d\ln \mathcal{U}_t}{d\ln p_{A,t}} \;=\; s_{A,t}.$$

Final demand satisfies $d\ln Y_t / d\ln P_t = -\varepsilon$; hence

$$\frac{d\ln Y_t}{d\ln p_{A,t}} \;=\; -\varepsilon\, s_{A,t}. \tag{A.9}$$

Substituting (A.9) into (A.8):

$$\frac{d\ln TC_t}{d\ln p_{A,t}} \;=\; (1-\varepsilon)\, s_{A,t}. \tag{A.10}$$

Effective Elasticity of AI Demand

Taking logs of (A.7) gives:

$$\ln A_{Y,t} = \ln \omega_A \;-\; \sigma_Y \ln p_{A,t} \;+\; (\sigma_Y - 1)\ln \mathcal{U}_t \;+\; \ln TC_t.$$

Then, differentiating $\ln A_{Y,t}$ w.r.t. $\ln p_{A,t}$, and using $d\ln \mathcal{U}_t / d\ln p_{A,t} = s_{A,t}$ and $\frac{d\ln TC_t}{d\ln p_{A,t}} = (1-\varepsilon)\, s_{A,t}$ gives:

$$\frac{d\ln A_{Y,t}}{d\ln p_{A,t}} \;=\; -\sigma_Y \;+\; (\sigma_Y - 1)\, s_{A,t} \;+\; (1-\varepsilon)\, s_{A,t} \;=\; -\sigma_Y \;+\; (\sigma_Y - \varepsilon)\, s_{A,t}.$$

Finally, defining the (positive) effective demand elasticity as

$$\epsilon \;\equiv\; -\frac{d\ln A_{Y,t}}{d\ln p_{A,t}}.$$

gives the effective elasticity facing the monopolist:

$$\epsilon \;=\; \sigma_Y(1 - s_{A,t}) \;+\; \varepsilon\, s_{A,t}. \tag{A.11}$$

Written this way, the expression for the effective elasticity separates the input substitution effect, $\sigma_Y(1 - s_{A,t})$, from the output scale effect, $\varepsilon\, s_{A,t}$. When the price of AI declines, AI demand rises both through input substitution (firms use more AI relative to other inputs) and through output expansion (lower costs enable firms to reduce prices and expand production). Notice that the input substitution effect dominates when AI has a small cost share, but the output scale effect dominates when AI's cost share is large.

B Parameter Values

Table 2: Parameter Values for Simulations

Symbol	Description	Value
Household / Intertemporal		
β	Discount factor	0.980
σ_c	Coefficient of relative risk aversion	2.000
ε	Elasticity of substitution across varieties	6.000
Capital		
δ	Depreciation rate	0.050
Technology: AI production (A_t)		
ϕ	Learning parameter	0.150
α_A	Recursive AI distribution parameter	0.425
α_K	Capital distribution parameter	0.425
α_{L^s}	Specialized labor distribution parameter	0.050
α_E	Energy distribution parameter	0.100
σ_A	Elasticity of substitution	0.900, 1.000, 1.100
Technology: final-goods production (Y_t)		
ω_K	Capital distribution parameter	0.350
ω_A	AI distribution parameter	0.025
ω_{L^s}	Specialized labor distribution parameter	0.125
ω_{L^n}	Nonspecialized labor distribution parameter	0.450
ω_E	Energy distribution parameter	0.050
σ_Y	Elasticity of substition	0.100-2.000
Endowments / Initial conditions		
$\bar{L}^s$	Specialized labor supply	0.050
$\bar{L}^n$	Nonspecialized labor supply	0.950
$\bar{E}$	Energy endowment	1.000
$A_{Y,0}$	Initial A_Y	1.187
$A_{A,0}$	Initial A_A	0.001

References

Daron Acemoglu and Pascual Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542, 2018.

Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Political Economy*, 128(6):2188–2244, 2020.

Philippe Aghion, Benjamin F. Jones, and Charles I. Jones. Artificial intelligence and economic growth. *Brookings Papers on Economic Activity*, 48(2 (Fall)):1–45, 2017.

Seth G Benzell and Erik Brynjolfsson. Digital abundance and scarce genius: Implications for wages, interest rates, and growth. Technical report, National Bureau of Economic Research, 2019.

Erik Brynjolfsson and Tom Mitchell. What can machine learning do? workforce implications. *Science*, 358(6370):1530–1534, 2017. doi: 10.1126/science.aap8062.

Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen. Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence. Technical report, Stanford Digital Economy Lab, Stanford University, August 2025.

Iain M Cockburn, Rebecca Henderson, and Scott Stern. 4. the impact of artificial intelligence on innovation: An exploratory analysis. In *The economics of artificial intelligence*, pages 115–148. University of Chicago Press, 2019.

Avinash K Dixit and Joseph E Stiglitz. Monopolistic competition and optimum product diversity. *The American economic review*, 67(3):297–308, 1977.

H1B Data. Anthropic pbc: Research engineer salaries in san francisco, 2025. URL https://h1bdata.info/index.php?city=SAN+FRANCISCO&em=ANTHROPIC+PBC&job=RESEARCH+ENGINEER. Accessed September 5, 2025.

Business Insider. Top ai startup companies' salaries: Openai, anthropic, perplexity, July 2025. URL https://www.businessinsider.com/top-ai-startup-companies-salaries-pay-data-openai-anthropic-perplexity-2025-7.

Anton Korinek and Joseph E. Stiglitz. Artificial intelligence and its implications for income distribution and unemployment. In Ajay Agrawal, Joshua Gans, and Avi Goldfarb, editors, *The Economics of Artificial Intelligence: An Agenda*, pages 349–390. University of Chicago Press, 2019.

William D Nordhaus. Are we approaching an economic singularity? information technology and the future of economic growth. *American Economic Journal: Macroeconomics*, 13(1):299–332, 2021.

OECD. The supply, demand and characteristics of the ai workforce across oecd countries. Technical report, OECD Publishing, 2023. URL https://www.oecd.org/content/dam/oecd/en/publications/reports/2023/02/the-supply-demand-and-characteristics-of-the-ai-workforce-across-oecd-countries_fcfe53fb/bb17314a-en.pdf.

PYMNTS. Ai startups: What openai, anthropic pay their technical staff, July 2025. URL https://www.pymnts.com/artificial-intelligence-2/2025/ai-startups-what-openai-anthropic-pay-their-technical-staff.

U.S. Bureau of Labor Statistics. Median weekly earnings of full-time workers with only a bachelor's degree: $1541 in q2 2024. The Economics Daily, July 2024. URL https://www.bls.gov/opub/ted/2024/median-weekly-earnings-of-full-time-workers-with-only-a-bachelors-degree-1541-in-q2-2024.htm.

Laura Veldkamp and Jinhee Chung. Data as a factor of production. *Journal of Political Economy*, 2024. Forthcoming.