AI-Augmented Capital-Skill Complementarity*

André Victor D. Luduvice†
Federal Reserve Bank of Cleveland

Roberto Pinheiro‡
Federal Reserve Bank of Cleveland

- PRELIMINARY -

Abstract

We propose a theory and measurement of AI as capital, and use it to study how rapid improvements in AI may reshape factor demands, production, and the income distribution. Building on Krusell et al. (2000), we consider a production function with two types of labor (skilled and unskilled) and three types of capital (structures, equipment, and AI), allowing each capital type to complement or substitute for different labor inputs. We also allow for uneven capital ownership, which provides a simple bridge from aggregate adjustments to distributional outcomes. On the measurement side, we construct a quality-adjusted AI price index from hedonic regressions on accelerator characteristics and build a corresponding AI capital stock via a perpetual inventory method. We estimate the production function parameters via GMM and find that AI and the high-skill–equipment composite are complements. Using a real business cycle model, we study perfect-foresight transition dynamics under permanent AI shocks. Doubling the AI usage share is contractionary, while halving the AI price is expansionary; crucially, the combined effect is super-additive, with output gains far exceeding the sum of the individual effects.

*We would like to thank seminar participants at the the EESP/FGV-SP, the York University, and the University of Toronto. First preliminary version: January 2026. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Cleveland or the Federal Reserve System. All errors are our own.

†E-mail: andrevictor.luduvice@clev.frb.org.

‡E-mail: roberto.pinheiro@clev.frb.org (corresponding author).

1 Introduction

A recurring theme in discussions surrounding the adoption of AI in production is that it may reduce firms' labor demand by allowing them to automate tasks currently performed by workers. This view is often expressed explicitly by major firms and is consistent with emerging evidence that labor-market outcomes are already shifting in more AI-exposed parts of the economy. In the words of Andy Jassy in a recent Amazon's internal memo:

> *We will need fewer people doing some of the jobs that are being done today, and more people doing other types of jobs. It's hard to know exactly where this nets out over time, but in the next few years, we expect that this will reduce our total corporate workforce as we get efficiency gains from using AI extensively across the company.*
>
> —*Andy Jassy*[1]

Along similar lines, Brynjolfsson et al. (2025) find that early-career workers (ages 22–25) in the most AI-exposed occupations have experienced a 13 percent relative decline in employment, even after controlling for firm-level shocks.

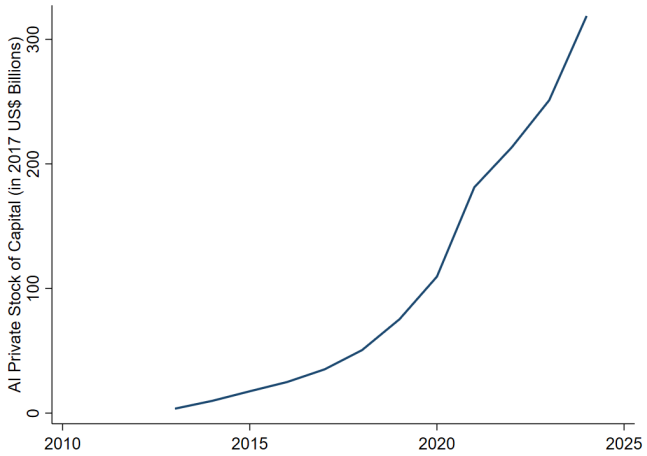
These labor-market concerns reflect just one of several channels through which AI may affect the economy. If AI is interpreted as a new form of capital that firms accumulate and deploy, then AI adoption influences aggregate production and measured productivity, as well as reshape the composition of investment and labor demand. Moreover, AI may affect groups differently, either because it interacts with skill in production or because it changes the task content of jobs in ways that are not uniform across workers. Building a macroeconomic model allows us to account for these interactions of AI with the different economic inputs in the production process, as well as AI's impact on economic aggregates. By estimating or calibrating the relevant parameters, we obtain estimates of the elasticity of substitution across the different production factors at the economy-wide or sectoral levels instead of particular markets. The rapid rise in the measured stock of AI capital, shown in Figure 1, is consistent with the idea that AI is becoming a quantitatively relevant input, making our approach empirically relevant.

The pace of recent technological change further motivates a macroeconomic approach. Performance benchmarks for AI models have improved quickly, and innovation activity has grown, as indicated by the level and share of AI-related patenting. In parallel, large investments in AI infrastructure suggest continued growth in installed capacity. Figures 2 and 3 summarize these patterns.

This paper studies these issues through the lens of a classic macroeconomic production function framework. A central ingredient is measurement: we propose to measure AI as capital and construct

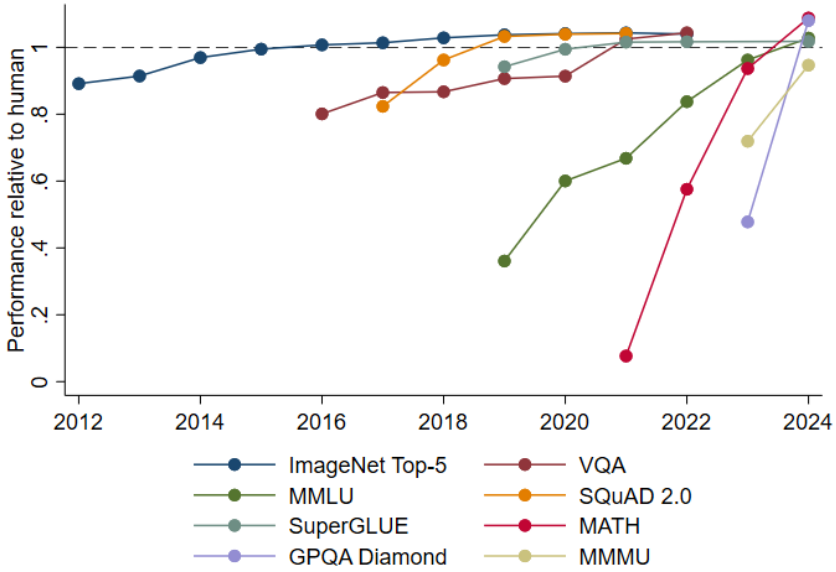
[1]See "Message from CEO Andy Jassy: Some thoughts on Generative AI": https://www.aboutamazon.com/news/company-news/amazon-ceo-andy-jassy-on-generative-ai

Figure 1: Evolution of AI Capital Stock



Note: This figure presents the stock of AI capital constructed through perpetual inventory method (PIM). Since AI is a brand new technology, we consider $K_0 = I_0$ as our initial benchmark capital. Data from Stanford HAI's 2025 AI Index Report. Data accessible here. Values deflated using BEA's ICT Price Index (series B935RG3A086NBEA).

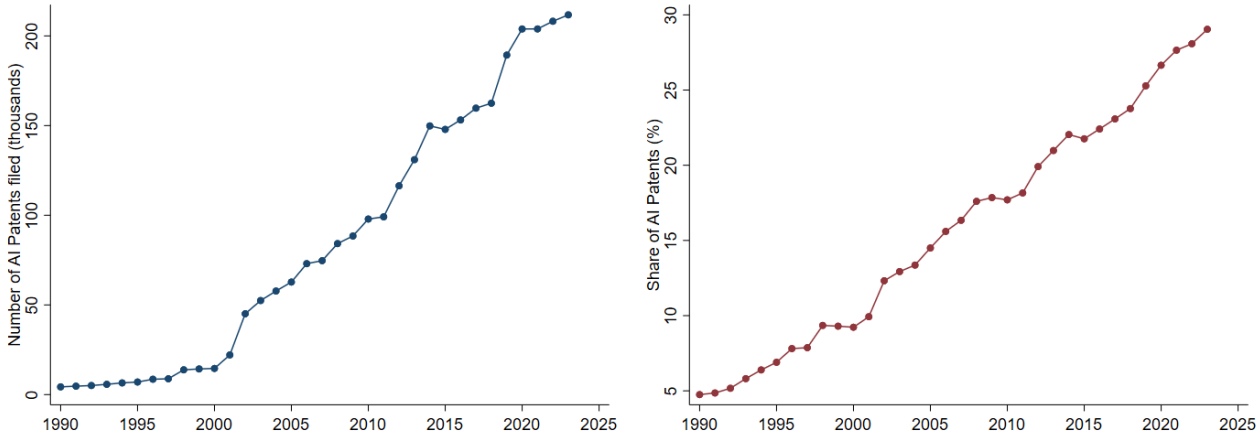
Figure 2: AI Models' Performance



Note: This figure presents some benchmark tests for model performance. In particular, we have: Image classification (ImageNet Top-5), Visual reasoning (VQA), Medium-level reading comprehension (SQuAD 2.0), English language understanding (SuperGLUE), Multitask language understanding (MMLU), Competition-level mathematics (MATH), PhD-level science questions (GPQA Diamond), and Multimodal understanding and reasoning (MMMU).

a quality-adjusted AI price index from hedonic regressions on accelerator characteristics, together with a corresponding AI capital stock via a perpetual inventory method. These series allow us to treat AI on the same footing as equipment or structures in a standard macro model. On the theory side, we build on the capital-skill complementarity literature (Krusell et al., 2000), augmenting the production function to include AI as a distinct capital input alongside structures and equipment.

Figure 3: AI Patents over Time



(a) Number of filed AI Patents

(b) Share of AI Patents

The model distinguishes between skilled and unskilled workers and allows each capital type to complement or substitute for different labor inputs. We embed this production structure into a classical real business cycles environment, allowing for the characterization and study of shocks on AI on model outcomes as well as transitional dynamics.

On the empirical side, the hedonic regressions estimate that quality-adjusted AI hardware prices have fallen by 97% between 2013 and 2024, roughly 27% per year, far outpacing the historical decline in ICT equipment prices. This rapid quality improvement translates, via the perpetual inventory method, into a quality-adjusted AI capital stock that has grown from US$ 15 billion in 2018 to US$ 625 billion in 2024 (in 2019 constant dollars), a 4,000% increase in six years. Despite this growth, AI capital remains small relative to the total capital stock, currently less than 5% of structures.

We then write and solve a general version of the model, characterizing its equilibrium as well as its steady-state. For the quantitative analysis, we start by simplifying the production function following the literature and proceed to the estimation and calibration of parameters to the US economy. We estimate the production function parameters via GMM, targeting moments that include the wage-bill ratio, the labor share, no-arbitrage conditions across capital types, and the AI expenditure share. A key finding is that AI and the high-skill–equipment composite are complements. Hence, fears of broad high-skill jobs eliminated by AI may be overblown given the currently available data. The estimate for the elasticity between high-skill labor and equipment confirms the original capital-skill complementarity result of Krusell et al. (2000), while the high elasticity between skill bundles implies that low- and high-skill composites are near-perfect substitutes in the aggregate.

With the estimated model in hand, we study perfect-foresight transition dynamics under three permanent AI shocks: doubling the AI usage share (a demand-side scenario), halving the AI price (a supply-side scenario), and both simultaneously. The two individual shocks produce different

outcomes while expanding the aggregate AI capital by roughly the same amount. Doubling the AI share is contractionary: output falls 0.95% in the new steady state because complementarity with a still-scarce input reduces the effective high-skill composite. Halving the AI price is expansionary: output rises 1.58% as cheaper AI capital accumulation raises the marginal product of the entire high-skill bundle without crowding out other investment.

When we combine both shocks, we uncover a novel feature of super-additivity in the outcomes where output rises more than three times the linear sum of the individual effects. This super-additivity arises because each channel removes the constraint binding on the other: the price drop makes it affordable to accumulate the AI capital that the higher weight demands, while the higher weight amplifies the productivity gain from each unit of cheaper AI. Because capital ownership is unevenly distributed in the economy, these exercises also generate distributional consequences: entrepreneurs capture the bulk of the gains from AI capital accumulation, the labor share declines, and the skill premium remains nearly neutral under the combined shock as the opposing effects on high-skill wages offset each other.

Road Map. The remainder of this section discusses the related literature. Section 2 describes the data sources and our measurement of AI investment, including the hedonic price index and the quality-adjusted capital stock. Section 3 presents the model, which extends a standard RBC framework by allowing AI capital to interact differently with skilled and unskilled labor while incorporating uneven capital ownership. Section 4 discusses the calibration and estimation strategy. Section 5 reports quantitative exercises on how AI shocks affect macroeconomic aggregates, transition dynamics, and the income distribution, including a robustness analysis across substitution elasticities. The final section concludes and outlines directions for future work.

Related Literature. Our paper builds on different strands of related extant literature. First, it builds on the capital-skill complementarity tradition. The canonical framework in Krusell et al. (2000) shows how changes in the relative price of equipment, combined with differential substitution patterns across skill groups, can generate shifts in labor demand and relative wages. Our approach to the technological change driven by AI follows this organizing principle by allowing multiple capital types, AI included, to interact differently with skilled and unskilled labor. Within that, we use the extended and improved estimates from Ohanian et al. (2023) and update them to recent years with the inclusion of our measurements for AI. A complementary literature emphasizes that "skill" differences often reflect deeper heterogeneity in occupations and tasks. The occupational and task-based approaches in Caunedo and Keller (2022), Caunedo et al. (2023), and the broader tasks framework in Acemoglu and Autor (2011) and Acemoglu and Restrepo (2022) motivate our focus on flexible substitution patterns as a key determinant of distributional outcomes.

Second, our paper is part of the burgeoning literature on the macroeconomics of AI. Recent work emphasizes that AI can simultaneously substitute for some labor inputs while complementing

others, with equilibrium consequences for output, wages, and inequality (Acemoglu et al., 2022; Acemoglu, 2025). We also draw from this literature to motivate and make key quantitative choices in our calibration, including benchmarks for AI/robot-capital income shares (Berg et al., 2018) and the mapping from investment to productive AI capital services and depreciation (Kudoh and Miyamoto, 2025). Our work is closest to Berg et al. (2018) in terms of the framework of analysis. However, different from Berg et al. (2018), our paper places greater emphasis on measurement and on disciplined calibration of substitution patterns across production inputs. Recent work by Guerrón-Quintana et al. (2024) supports the value of embedding AI mechanisms in general equilibrium to study aggregate and distributional responses jointly. We capture this central insight by introducing AI as an explicit capital category that can have different elasticities of substitution with each labor type.

Third, our model is anchored in the RBC tradition. We embed AI and capital-skill complementarity in a standard RBC environment to interpret AI as an investment-specific force operating through relative prices, input complementarity, and capital accumulation. Recent RBC-style quantitative work such as Kopytov et al. (2018) and Albanesi (2025) motivates treating capital composition and investment-specific forces as central for aggregate dynamics, which aligns with our focus on the quality-adjusted price of AI capital.

Fourth, the paper relates to the measurement of capital prices and quantities. We follow the hedonic measurement tradition in Gordon (1990) and the practical construction of quality-adjusted investment deflators in Cummins and Violante (2002). Related work on investment-specific technology and macro dynamics (DiCecio, 2009) and recent discussions of generative-AI measurement challenges (Baily et al., 2025) further motivate our approach to constructing AI price indexes and capital stocks. We contribute to this literature by proposing a new angle through which to measure AI, i.e., as capital, and use the lessons to compute a novel time series of quality-adjusted price of AI.

Overall, we integrate these strands by extending capital-skill complementarity to an explicit AI capital input, embedding it in an RBC framework, and disciplining the key AI-relative-price channel using the quality-adjustment measurement literature.

2 Empirical Analysis

2.1 Measuring AI as Capital

Measuring AI as a productive input raises three practical issues that are closely related, and that shape our empirical strategy. First, there is the distinction between usage-based measures and installed-capacity measures. Much of the macro literature that treats technology-specific capital as an input, including the capital-skill complementarity framework in Krusell et al. (2000), is

naturally aligned with installed capacity, typically constructed via perpetual-inventory methods, and adjusted for quality. Installed-capacity measures, however, may understate the importance of complementary infrastructure investment and organizational expenditures that are required for effective deployment. While our focus remains on an installed-capacity concept to match a standard macroeconomic model, we keep this limitation in mind when interpreting the results.

Second, data availability differs sharply across capital categories. For conventional aggregates such as structures, equipment, and ICT capital, we rely on standard BEA measures. For some ICT components, alternative sources based on micro data (such as the Aberdeen Ci Technology database) can provide additional granularity and cross-checks. For AI capital, by contrast, national accounts do not yet offer a dedicated series that separates AI-specific hardware and its quality evolution. We therefore combine external sources that track AI-related hardware characteristics and prices with investment data that can be mapped into an AI-capital concept consistent with the model.

Third, quality adjustment for AI is crucial. Because hardware performance is constantly and rapidly improving, nominal prices alone do not capture the relevant evolution in the effective price of AI capital. We therefore construct a quality-adjusted AI price index using a hedonic pricing approach in the spirit of Gordon (1990). For equipment, we follow the same principle and extend Gordon's quality-adjusted series using the approach in Cummins and Violante (2002) and DiCecio (2009). These choices ensure that the data inputs correspond to the variables in the model: quality-adjusted prices and quantities of capital that can be accumulated over time.

2.2 AI Epoch Data

Our main AI-hardware database comes from Epoch AI and includes information on over 170 AI accelerators, including GPUs and TPUs, that have been used to develop and deploy machine learning models. A key constraint is that only a subset of these accelerators are commercial products for which release prices can be observed, so the usable sample for hedonic price estimation is smaller than the universe of listed chips. In practice, Epoch AI provides price data for 61 processing units starting in 2008, with 56 of these starting in 2013.

Beyond prices, the database contains technical characteristics that can be used to control for changes in quality. Following the measurement emphasis highlighted by Baily et al. (2025), we focus on three attributes that are both economically relevant and broadly available across chips: *ML OP/s* (maximum performance in any format 8 bits or wider, measured in FLOP/s or OP/s), *memory size per board* (in bytes), and thermal design power (*TDP*), which captures the maximum power that can be dissipated as heat. These characteristics capture performance, capacity constraints, and energy requirements, which are all central drivers of effective prices in AI hardware markets.

2.3 Quality-Adjusted Price Index

To construct a quality-adjusted AI price index, we follow the hedonic methodology in Chapter 6 of Gordon (1990). For processor i released at time t, we estimate

$$\log p_{i,t} = \beta_0 + \sum_{s=1}^{N} \delta_s D_s + \sum_{j=1}^{m} \log y_{j,i,t} + u_{i,t}, \tag{1}$$

where $p_{i,t}$ denotes the release price, $y_{j,i,t}$ denotes characteristic j for processor i at time t, and D_s are year dummies. We focus on newly released processors, so that the regression traces the technology frontier, as recommended by Fisher et al. (1985). This mirrors the logic in Gordon's construction of quality-adjusted ICT price indexes, where frontier products are used to characterize the evolving cost of a unit of quality.

We recover the hedonic price index using the "dummy variable method": the quality-adjusted price index is given by the antilog of the year-dummy coefficients. Figure 4 reports the resulting series, which summarizes the evolution of the effective (quality-adjusted) price of AI hardware over time. The index equals 1 during the base year of 2019. Appendix Figure A.1 includes the confidence intervals. As we can see in that figure, confidence intervals were quite wide early on, while narrowing significantly in more recent years.

Figure 4: Evolution of AI Quality-Adjusted Price Index

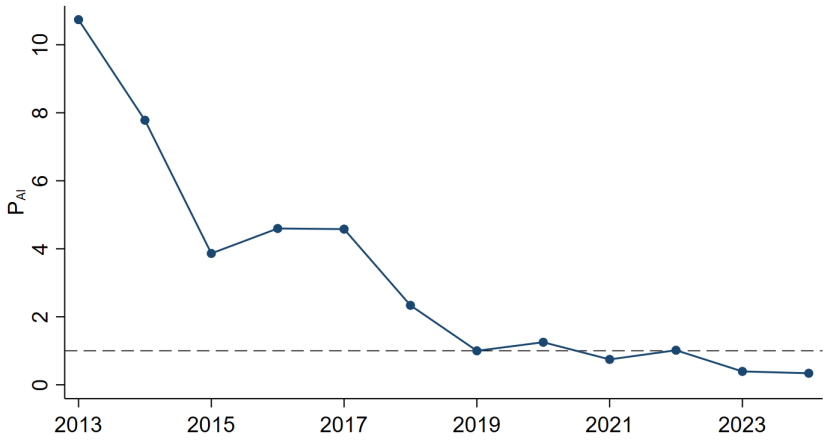


Figure 4 shows a sharp decline in quality-adjusted prices.[2] Considering the normalization of 1 in 2019, the index declines from 10.7 in 2013 to 0.34 in 2024, a decline of 97% in 11 years, or 27% per year. If we focus on the period with the largest investment in AI (2018–2024), the decrease in quality adjusted prices was 85% in 6 years, or 27.5% per year.

[2]Periods around the pandemic showed quite noisy and statistically insignificant coefficients. This is expected given the supply chain disruptions that pushed microchip and goods prices up during that period.

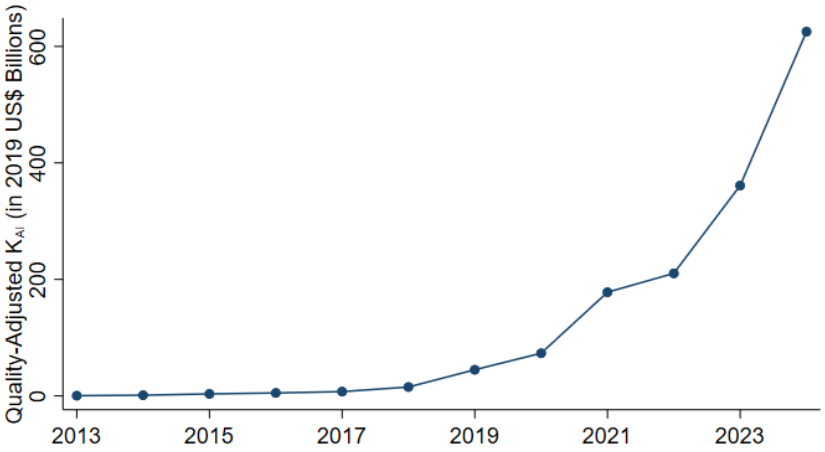
2.4 Quality-Adjusted Capital Stock

Given the price index, we construct a quality-adjusted stock of AI capital using a perpetual inventory method (PIM), following Krusell et al. (2000). Because AI is a relatively new technology in the aggregate data, we normalize the initial stock to $K_{AI} = 0$ in 2012 and then accumulate AI capital according to

$$K_{AI,t} = (1 - \delta_{AI})K_{AI,t-1} + I_{AI,t},$$

where $I_{AI,t}$ is investment in AI capital measured using data from Stanford HAI's 2025 AI Index Report, deflated by our quality-adjusted AI price index. We set $\delta_{AI} = 0.156$, following Kudoh and Miyamoto (2025). Figure 5 presents the resulting time series for the quality-adjusted AI capital stock. Appendix Figure A.2 includes the confidence intervals.

Figure 5: Evolution of AI Quality-Adjusted Capital Stock

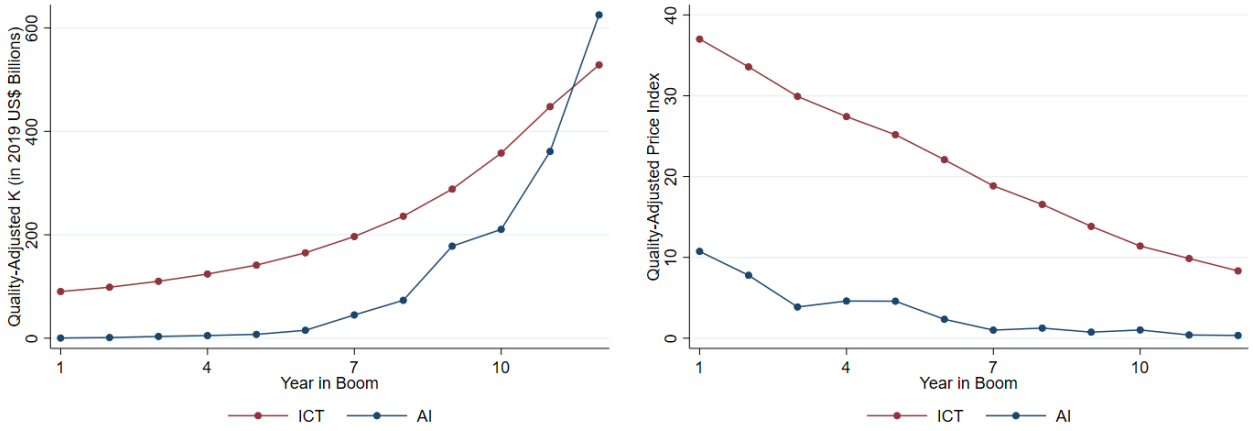


As we can see in Figure 5, there was a fast accumulation of capital, in particular after 2018. In fact, quality-adjusted K_{AI} increased from US$ 15 billions in 2018 to 625 billions in 2024 in 2019 constant dollars (a 4,000% increase in 6 years, representing an 85% annual increase). Furthermore, given that the gap between the point estimate and the upper limit of the 95% confidence interval has widened significantly, as shown in Appendix Figure A.2, the growth in the quality-adjusted capital stock may be even stronger.

However, this capital stock remains small in absolute terms. Even in 2024, quality-adjusted AI capital was less than 5% of NIPA's capital stock of structures. Despite rapid growth, AI capital is still a modest share of the total. In this sense, the boom in AI investment shares some similarities with the ICT boom. While ICT has seen a significant growth in the 1990s and early 2000s, the total stock of ICT capital has always remained a somewhat small share of the total capital stock. We focus on this comparison next.

Comparison: ICT revolution To put these findings in perspective, it is useful to benchmark AI capital against the ICT capital stock that sparked the literature following Krusell et al. (2000). Figure 6 below presents the evolution of the ICT quality-adjusted capital stock (panel 6a) and price index (panel 6b) from 1990 to 2001 and compare it against the growth in AI counterparts from 2013 to 2024.[3] Our ICT data comes from Ohanian et al. (2023).[4]

Figure 6: The ICT vs. AI Revolutions: Capital Stock and Prices



(a) Quality-Adjusted Capital Stock

(b) Quality-Adjusted Price Index

Note: This figure presents changes in the AI and ICT quality-adjusted capital stocks and price indexes. We consider years in boom 1990-2001 for ICT and 2013-2024 for AI.

As we can see in panel 6a, while ICT growth was significantly slower than the recent investment in AI. Between 1990 to 2001, the ICT capital has grown from US$ 90 billion to US$ 529 billion in 2019 constant dollars, representing a growth of 488% in 11 years or 17% per year. Notice that this is significantly slower growth than what we have seen for AI in the last 6 years.

In terms of decline in quality-adjusted ICT prices, as shown in panel 6b, we have also seen a faster AI price decline than during the ICT revolution. Normalizing the initial ICT Price index to 1 in 1990, by 2001 the ICT price index had declined to .23. In contrast, after normalizing the initial AI price index to 1 in 2013, the AI price index had declined to 0.032 in 2024.

2.5 Computers & Peripherals as AI Proxy

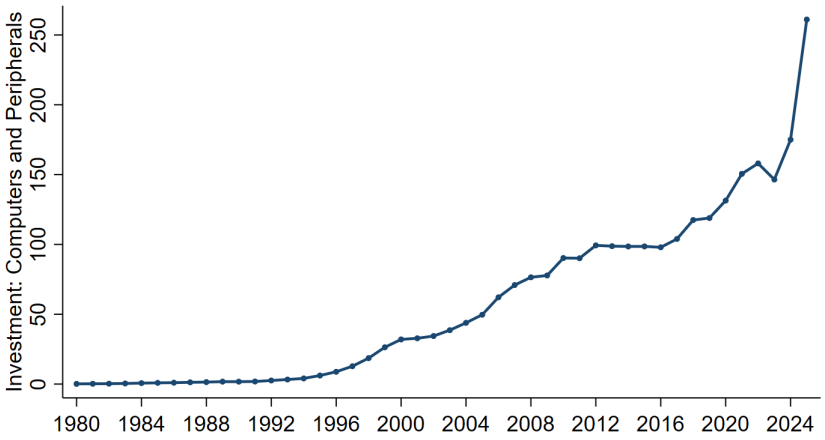
Because our AI capital series based on accelerator hardware begins only recently, we initially proxy AI capital using BEA real investment in Computers & Peripherals starting in 2000. This category encompasses the computing hardware that underpins AI workloads—including GPUs, TPUs, and

[3]In Appendix Figure A.3, we extend the ICT capital stock time series to 2019 to see how it evolved beyond the window presented here.

[4]We simply readjusted the base year to 2019 – Ohanian et al. (2023)'s base year was 2012.

the servers that house them— and is available over a longer sample. As we see in Figure 7, real investment in Computers & Peripherals captured the run-up in AI investment. Using Computers & Peripherals from 2000 onward allows us to estimate the production function parameters over a period that captures the early diffusion of machine-learning hardware into the capital stock, while maintaining consistency with the BEA capital accounts used for structures and equipment. In Section 5, we use the estimated parameters together with our AI-specific series for the counterfactual experiments.

Figure 7: Real Investment in Computers & Peripherals (US$ Billions)



2.6 Possible Alternative Approaches

Our baseline measurement strategy mirrors the approach used for equipment and ICT in the macro literature: we emphasize installed capacity, use quality-adjusted prices, and construct capital stocks with PIM. For AI, however, additional data sources may allow complementary approaches. In particular, some datasets provide measures of price per chip-hour, benchmarking information on performance in deployed tasks, and preliminary usage measures drawn from the Annual Business Survey and Regional Federal Reserve surveys. These alternative sources could be informative about utilization or effective services of AI capital, as opposed to installed capacity alone. In future work, these data may also support a multi-sector extension that distinguishes an AI-producing sector from AI-using sectors, allowing the model to capture both the supply of AI hardware and its diffusion into production.

3 Model

We build a quantitative environment that extends the capital-skill complementarity framework of Krusell et al. (2000) to include AI capital as a separate, embodied input. Time is discrete and infinitely-lived agents discount future at rate $\beta \in (0, 1)$. We consider a 5-factor production function

that allows for different elasticities of substitution between different types of capital and labor. Production combines structures, equipment, and AI with two labor types (low-skill and high-skill), allowing substitution patterns to differ across inputs.

Technological progress enters through a neutral productivity term A_t and through investment-specific components q_t (equipment) and ζ_t (AI), which map into their relative prices. We introduce a simple ownership structure in which entrepreneurs own the capital stock and receive profits, while workers earn labor income. This partition is meant to capture the idea that AI-related returns may accrue disproportionately through capital income.

3.1 Production

There are four final goods: consumption c_t, structures investment x_{st}, equipment investment x_{et}, and AI investment x_{AIt}.[5] Consumption and structures are produced with a constant returns to scale technology, while equipment and AI are produced with the same technology scaled by the equipment-specific and AI-specific technological progress q_t and ζ_t, respectively. Under these assumptions, the relative price of equipment is $\frac{1}{q_t}$, while the relative price of AI is $\frac{1}{\zeta_t}$. Hence, the aggregate production function is given by:

$$y_t = c_t + x_{st} + \frac{x_{et}}{q_t} + \frac{x_{AIt}}{\zeta_t} = A_t G(k_{st}, k_{et}, k_{AIt}, m_{Lt}, m_{Ht}) \tag{2}$$

where $G(k_{st}, k_{et}, k_{AIt}, m_{Lt}, m_{Ht})$ is given by:

$$k_{st}^{\alpha_1} \left\{ \begin{array}{l} \left[A_L \left(m_{Lt}^{\gamma_L} + k_{et}^{\gamma_L} \right)^{\frac{\lambda_L}{\gamma_L}} + \Psi_{AI} k_{AIt}^{\lambda_L} \right]^{\frac{\phi}{\lambda_L}} \\ + \left[A_H \left(m_{Ht}^{\gamma_H} + k_{et}^{\gamma_H} \right)^{\frac{\lambda_H}{\gamma_H}} + \Psi_{AI} k_{AIt}^{\lambda_H} \right]^{\frac{\phi}{\lambda_H}} \end{array} \right\}^{\frac{\alpha_2}{\phi}} \tag{3}$$

where k_s, k_e, and k_{AI} are capital structures, equipment, and AI, respectively. Similarly, m_L and m_H are low-skill and high-skill labor.

The production aggregator nests inputs in a way that mirrors the logic of capital-skill complementarity while allowing AI to enter symmetrically across skill groups. Structures enter in a Cobb-Douglas fashion, while the remaining inputs form a composite of two skill-specific bundles. Each bundle combines skill-j labor with equipment through a CES structure governed by γ_j, and then incorporates AI capital with strength governed by (λ_j, Ψ_{AI}). The outer parameter ϕ governs the aggregation across the two skill bundles. This structure allows AI to be complementary to, or a substitute for, each labor type depending on parameter values.

[5] Notice that investments are expressed in units of the relevant capital.

We consider that worker types are fixed, i.e., a worker is born with a type $j \in \{L, H\}$. As a result, a worker decides how much labor to supply, but cannot supply labor of a different type. We assume that the measure of workers of type j is M_j. We also consider the presence of entrepreneurs, who own the firm profits and other means of production. Entrepreneurs do not participate in production, apart from providing capital. We assume that the measure of entrepreneurs is M_E. We assume that the total population is $\overline{M}$, with $M_L + M_H + M_E = \overline{M}$.

3.2 Workers' Problems

3.2.1 Low-skill Worker

We assume that the only source of income for a low-skill worker is labor income or governmental benefits. In other words, we assume that low-skill workers do not have a savings technology. As a result, the low-skill worker problem is given by:

$$\max_{\{C_{L,t+s}, L_{L,t+s}\}_{s=0}^{\infty}} E_t \sum_s \beta^s b_{t+s} \left[\ln(C_{L,t+s} - \eta C_{L,t+s-1}) - \frac{L_{L,t+s}^{1+\frac{1}{\tau}}}{1+\frac{1}{\tau}} \right] \tag{4}$$

subject to:

$$C_{L,t+s} = w_{L,t+s} L_{L,t+s} - T_{L,t+s}$$

where β is the discount rate, $C_{L,t}$ is consumption, η is the degree of habit formation and b_t is a shock to the discount factor. $L_{L,t}$ is the individual's labor supply in period t. We assume that an individual has a unit of labor to supply per period. Hence $L_{L,t} \leq 1, \forall t$. Then, the Lagrangean becomes:

$$\mathcal{L} = E_t \sum_s \beta^s b_{t+s} \left\{ \begin{array}{c} \left[\ln(C_{L,t+s} - \eta C_{L,t+s-1}) - \frac{L_{L,t+s}^{1+\frac{1}{\tau}}}{1+\frac{1}{\tau}} \right] - \nu_{L,t}(L_{L,t} - 1) \\ + \mu_{L,t+s} \left(-C_{L,t+s} + w_{L,t+s} L_{L,t+s} - T_{L,t+s} \right) \end{array} \right\}$$

then, F.O.C.s are:

$$\left\{ \begin{array}{l} \frac{\partial \mathcal{L}}{\partial C_{L,t}} = \frac{b_t}{C_{L,t} - \eta C_{L,t-1}} - \mu_{L,t} - \beta \eta E_t \frac{b_{t+1}}{C_{L,t+1} - \eta C_{L,t}} = 0 \\ \frac{\partial \mathcal{L}}{\partial L_{L,t}} = -b_t L_{L,t}^{\frac{1}{\tau}} + \mu_{L,t} w_{L,t} - \nu_{L,t} = 0 \\ \frac{\partial \mathcal{L}}{\partial \mu_{L,t}} = -C_{L,t} + w_{L,t} L_{L,t} - T_{L,t} = 0 \end{array} \right.$$

and complementary slackness conditions:

$$\nu_{L,t}(L_{L,t} - 1) = 0 \quad \text{and} \quad \nu_{L,t} \geq 0$$

3.2.2 High-skill Worker

The problem for high-skill worker is similar. We again assume that high-skill workers have no access to a savings technology. This hand-to-mouth assumption is a simplification that isolates how factor prices and profits move with AI.[6] Hence, the problem for the high-skill worker is:

$$\max_{\{C_{H,t+s},L_{H,t+s}\}_{s=0}^{\infty}} E_t \sum_s \beta^s b_{t+s} \left[\ln(C_{H,t+s} - \eta C_{H,t+s-1}) - \frac{L_{H,t+s}^{1+\frac{1}{\tau}}}{1+\frac{1}{\tau}} \right] \tag{5}$$

subject to:

$$C_{H,t+s} = w_{H,t+s} L_{H,t+s} - T_{H,t+s}$$

Then, the Lagrangean becomes:

$$\mathcal{L} = E_t \sum_s \beta^s b_{t+s} \left\{ \begin{array}{c} \left[\ln(C_{H,t+s} - \eta C_{H,t+s-1}) - \frac{L_{H,t+s}^{1+\frac{1}{\tau}}}{1+\frac{1}{\tau}} \right] - \nu_{H,t}(L_{H,t} - 1) \\ +\mu_{H,t+s} \left(-C_{H,t+s} + w_{H,t+s} L_{H,t+s} - T_{H,t+s} \right) \end{array} \right\}$$

then, F.O.C.s are:

$$\left\{ \begin{array}{l} \frac{\partial \mathcal{L}}{\partial C_{H,t}} = \frac{b_t}{C_{H,t} - \eta C_{H,t-1}} - \mu_{H,t} - \beta \eta E_t \frac{b_{t+1}}{C_{H,t+1} - \eta C_{H,t}} = 0 \\ \frac{\partial \mathcal{L}}{\partial L_{H,t}} = -b_t L_{H,t}^{\frac{1}{\tau}} + \mu_{H,t} w_{H,t} - \nu_{H,t} = 0 \\ \frac{\partial \mathcal{L}}{\partial \mu_{H,t}} = -C_{H,t} + w_{H,t} L_{H,t} - T_{H,t} = 0 \end{array} \right.$$

and complementary slackness conditions:

$$\nu_{H,t}(L_{H,t} - 1) = 0 \quad \text{and} \quad \nu_{H,t} \geq 0$$

3.3 Entrepreneur's Problem

The entrepreneurs own the firm profits and other means of production and receive the firm's profits as compensation. Consequently, the entrepreneur must decide not only her own consumption, but also the investment in each type of capital. Therefore, the problem of the entrepreneur is:

$$\max_{\{C_{E,t+s},k_{st+s+1},k_{e,t+s+1},k_{AI,t+s+1}\}_{s=0}^{\infty}} E_t \sum_s \beta^s b_{t+s} \left[\ln(C_{E,t+s} - \eta C_{E,t+s-1}) \right] \tag{6}$$

[6] We plan to relax it in future versions by allowing at least high-skill workers to save.

subject to:

$$C_{E,t+s} + k_{st+s+1} + P_{e,t+s}k_{e,t+s+1} + P_{AI,t+s}k_{AI,t+s+1} = \left\{ \begin{array}{c} P_{e,t+s}(1-\delta_e)k_{e,t+s} \\ +P_{AI,t+s}(1-\delta_{AI})k_{AI,t+s} + (1-\delta_s)k_{st+s} \\ +r_{s,t+s}k_{st+s} + r_{AI,t+s}k_{AI,t+s} + r_{e,t+s}k_{e,t+s} \\ +\pi_{t+s-1} - T_{E,t+s} \end{array} \right\}$$

where π_{t+s-1} is the share of firm's profit, $P_{e,t+s} = \frac{1}{q_{t+s}}$ and $P_{AI,t+s} = \frac{1}{\zeta_{t+s}}$ are the prices for equipment and AI capital, respectively. Moreover, the stocks of capital evolve as:

$$\left\{ \begin{array}{ll} k_{s,t+s+1} = (1-\delta_s)k_{s,t+s} + x_{s,t+s} & (LOM_s) \\ k_{e,t+s+1} = (1-\delta_e)k_{e,t+s} + x_{e,t+s} & (LOM_e) \\ k_{AI,t+s+1} = (1-\delta_{AI})k_{AI,t+s} + x_{AI,t+s} & (LOM_{AI}) \end{array} \right.$$

Substituting the laws of motion and prices in the budget constraint, we have:

$$C_{E,t+s} + x_{st+s+1} + \frac{x_{e,t+s+1}}{q_{t+s}} + \frac{x_{AI,t+s+1}}{\zeta_{t+s}} = \left\{ \begin{array}{c} r_{e,t+s}k_{e,t+s} + r_{s,t+s}k_{st+s} \\ +r_{AI,t+s}k_{AI,t+s} + \pi_{t+s-1} - T_{E,t+s} \end{array} \right\}$$

where π_{t+s-1} is total firm profit divided by M_E. Hence, the Lagrangean becomes:

$$\mathcal{L} = \left\{ \begin{array}{c} E_t \sum_s \beta^s b_{t+s} \left[\ln(C_{E,t+s} - \eta C_{E,t+s-1})\right] \\ -\mu_{E,t+s} \left[\begin{array}{c} C_{E,t+s} + k_{st+s+1} + \frac{k_{e,t+s+1}}{q_{t+s}} + \frac{k_{AI,t+s+1}}{\zeta_{t+s}} \\ -\left(r_{e,t+s} + \frac{(1-\delta_e)}{q_{t+s}}\right) k_{e,t+s} \\ -(r_{s,t+s} + 1 - \delta_s)k_{st+s} - \pi_{t+s-1} + T_{E,t+s} \\ -\left(r_{AI,t+s} + \frac{(1-\delta_{AI})}{\zeta_{t+s}}\right) k_{AI,t+s} \end{array} \right] \end{array} \right\}$$

so the first-order conditions are:

$$\left\{ \begin{array}{ll} -\mu_{E,t} + E_t\mu_{E,t+1}(r_{s,t+1} + 1 - \delta_s) = 0 & (FOC_{k_s}) \\ -\frac{\mu_{E,t}}{q_t} + E_t\mu_{E,t+1}\left(r_{e,t+1} + \frac{(1-\delta_e)}{q_{t+1}}\right) = 0 & (FOC_{k_e}) \\ -\frac{\mu_{E,t}}{\zeta_t} + E_t\mu_{E,t+1}\left(r_{AI,t+1} + \frac{(1-\delta_{AI})}{\zeta_{t+1}}\right) = 0 & (FOC_{k_{AI}}) \\ \frac{b_t}{C_{E,t} - \eta C_{E,t-1}} - \mu_{E,t} - \beta E_t b_{t+1} \frac{\eta}{C_{E,t+1} - \eta C_{E,t}} = 0 & (FOC_C) \end{array} \right.$$

3.4 Firm's Problem

We consider a competitive firm that rents capital and labor in a competitive market. As a result, firms can optimize period by period. Hence, firm's problem is given by:

$$\max_{\substack{k_{st},k_{e,t},k_{AI,t},\\ m_{L,t},m_{H,t}}} \left\{ \begin{array}{l} A_t k_{st}^{\alpha_1} \left\{ \begin{array}{c} \left[A_L \left(m_{Lt}^{\gamma_L} + k_{et}^{\gamma_L} \right)^{\frac{\lambda_L}{\gamma_L}} + \Psi_{AI} k_{AIt}^{\lambda_L} \right]^{\frac{\phi}{\lambda_L}} \\ + \left[A_H \left(m_{Ht}^{\gamma_H} + k_{et}^{\gamma_H} \right)^{\frac{\lambda_H}{\gamma_H}} + \Psi_{AI} k_{AIt}^{\lambda_H} \right]^{\frac{\phi}{\lambda_H}} \end{array} \right\}^{\frac{\alpha_2}{\phi}} \\ -r_{s,t} k_{st} - r_{e,t} k_{e,t} - r_{AI,t} k_{AI,t} - w_{L,t} m_{L,t} - w_{H,t} m_{H,t} \end{array} \right\}$$

where k_s, k_e, and k_{AI} are capital structures, equipment, and AI, and r_{it} is the rental rate for capital of type $i \in \{s, e, AI\}$ at time t, m_L, m_H are the demand for low- and high-skill labor, $w_{i,t}$ the wage rate for skill $i \in \{H, L\}$.

Then, the first-order conditions are:

$$\left\{ \begin{array}{ll} \alpha_2 A_t k_{st}^{\alpha_1} \tilde{G}_t^{\frac{\alpha_2}{\phi}-1} X_{L,t}^{\frac{\phi}{\lambda_L}-1} A_L \left(m_{Lt}^{\gamma_L} + k_{et}^{\gamma_L} \right)^{\frac{\lambda_L}{\gamma_L}-1} m_{Lt}^{\gamma_L - 1} - w_{L,t} = 0 & (FOC_L) \\ \alpha_2 A_t k_{st}^{\alpha_1} \tilde{G}_t^{\frac{\alpha_2}{\phi}-1} X_{H,t}^{\frac{\phi}{\lambda_H}-1} A_H \left(m_{Ht}^{\gamma_H} + k_{et}^{\gamma_H} \right)^{\frac{\lambda_H}{\gamma_H}-1} m_{Ht}^{\gamma_H - 1} - w_{H,t} = 0 & (FOC_H) \\ \alpha_1 A_t k_{st}^{\alpha_1 - 1} \tilde{G}_t^{\frac{\alpha_2}{\phi}} - r_{s,t} = 0 & (FOC_{k_s}) \\ \alpha_2 A_t k_{st}^{\alpha_1} \tilde{G}_t^{\frac{\alpha_2}{\phi}-1} \left\{ \begin{array}{c} X_{L,t}^{\frac{\phi}{\lambda_L}-1} A_L \left(m_{Lt}^{\gamma_L} + k_{et}^{\gamma_L} \right)^{\frac{\lambda_L}{\gamma_L}-1} k_{et}^{\gamma_L - 1} \\ + X_{H,t}^{\frac{\phi}{\lambda_H}-1} A_H \left(m_{Ht}^{\gamma_H} + k_{et}^{\gamma_H} \right)^{\frac{\lambda_H}{\gamma_H}-1} k_{et}^{\gamma_H - 1} \end{array} \right\} - r_{e,t} = 0 & (FOC_{k_e}) \\ \alpha_2 A_t k_{st}^{\alpha_1} \tilde{G}_t^{\frac{\alpha_2}{\phi}-1} \Psi_{AI} \left\{ X_{L,t}^{\frac{\phi}{\lambda_L}-1} k_{AIt}^{\lambda_L - 1} + X_{H,t}^{\frac{\phi}{\lambda_H}-1} k_{AIt}^{\lambda_H - 1} \right\} - r_{AI,t} = 0 & (FOC_{k_{AI}}) \end{array} \right.$$

where:

$$\tilde{G}_t = \left\{ X_{L,t}^{\frac{\phi}{\lambda_L}} + X_{H,t}^{\frac{\phi}{\lambda_H}} \right\}$$

and

$$X_{L,t} \equiv \left[A_L \left(m_{Lt}^{\gamma_L} + k_{et}^{\gamma_L} \right)^{\frac{\lambda_L}{\gamma_L}} + \Psi_{AI} k_{AIt}^{\lambda_L} \right]$$

$$X_{H,t} \equiv \left[A_H \left(m_{Ht}^{\gamma_H} + k_{et}^{\gamma_H} \right)^{\frac{\lambda_H}{\gamma_H}} + \Psi_{AI} k_{AIt}^{\lambda_H} \right]$$

3.5 Market Clearing Conditions

Product market clearing condition:

$$c_t + x_{st} + \frac{x_{et}}{q_t} + \frac{x_{AIt}}{\zeta_t} = A_t G(k_{st}, k_{et}, k_{AIt}, m_{Lt}, m_{Ht})$$

using the capital laws of motion, we have:

$$\left\{ \begin{array}{c} M_L C_{L,t} + M_H C_{H,t} + M_E C_{E,t} \\ +k_{s,t+1} - (1-\delta_s) k_{s,t} \\ + \frac{k_{e,t+1} - (1-\delta_e) k_{e,t}}{q_t} \\ + \frac{k_{AI,t+1} - (1-\delta_{AI}) k_{AI,t}}{\zeta_t} \end{array} \right\} = A_t k_{st}^{\alpha_1} \left\{ \begin{array}{c} \left[A_L \left(m_{Lt}^{\gamma_L} + k_{et}^{\gamma_L} \right)^{\frac{\lambda_L}{\gamma_L}} + \Psi_{AI} k_{AIt}^{\lambda_L} \right]^{\frac{\phi}{\lambda_L}} \\ + \left[A_H \left(m_{Ht}^{\gamma_H} + k_{et}^{\gamma_H} \right)^{\frac{\lambda_H}{\gamma_H}} + \Psi_{AI} k_{AIt}^{\lambda_H} \right]^{\frac{\phi}{\lambda_H}} \end{array} \right\}^{\frac{\alpha_2}{\phi}}$$

Labor market clearing condition:

$$M_i L_{it} = m_{it}, \ \forall i \in \{L, H\} \text{ and } t \in \mathbb{N}$$

where M_i is the measure of agents in the economy of type $i \in \{L, H\}$.

3.6 Closing the Model

The equilibrium in this economy is pinned down by a standard set of optimality and feasibility restrictions: workers choose consumption and labor supply taking wages as given, entrepreneurs choose consumption and investment taking prices and returns as given, and competitive firms rent capital and hire labor until marginal products equal factor prices. Together with complementary slackness (for the labor upper bounds) and the goods- and labor-market clearing conditions, these restrictions form a system of dynamic equations that maps states and shocks into allocations and prices.

In the baseline version in which workers do not save and entrepreneurs own the productive capital, the equilibrium system consists of 22 equations in 22 unknowns. The unknowns are:

$$\{Y, C_L, C_H, C_E, L_L, L_H, k_s, k_e, k_{AI}, m_L, m_H, r_s, r_e, r_{AI}, w_L, w_H, \mu_L, \mu_H, \mu_E, \nu_L, \nu_H, \pi\}. \tag{7}$$

The 22 equilibrium conditions are: the 3 Euler equations for entrepreneurs' intertemporal choices (structures, equipment, and AI); the 3 consumption F.O.C.'s (one for each agent type); the 2 labor-supply F.O.C.'s (low- and high-skill); the 2 budget constraints for workers; the 1 budget constraint for entrepreneurs; the 5 firm first-order conditions (two for labor demands and three for capital demands); the 2 labor-market clearing conditions; the 1 goods-market clearing condition; and the 3 complementary slackness conditions associated with the labor upper bounds.[7]

[7] In the current baseline, entrepreneurs do not supply labor, so there is no labor-supply choice (and no associated slackness) for type E. If we later allow entrepreneurs to supply labor, we add the corresponding intratemporal F.O.C. and slackness condition and adjust the accounting accordingly.

3.7 Steady-State

We initially assume $b_t = 1$, $T_{i,t} = 0$, and constant technology parameters $q_t = q$ and $\zeta_t = \zeta$, for all $t \in \mathbb{N}$ and $i \in \{L, H\}$.

Then, from consumer L problem, we have:

$$\begin{cases} C_L = \frac{1-\beta\eta}{(1-\eta)\mu_L} & (FOC_{C_L}) \\ L_L = (\mu_L w_L - \nu_L)^\tau & (FOC_{L_L}) \\ C_L = w_L L_L & (FOC_{\mu_L}) \\ \nu_L(L_L - 1) = 0 & (CS_L) \end{cases}$$

similarly, from consumer H problem, we have:

$$\begin{cases} C_H = \frac{1-\beta\eta}{(1-\eta)\mu_H} & (FOC_{C_H}) \\ L_H = (\mu_H w_H - \nu_H)^\tau & (FOC_{L_H}) \\ C_H = w_H L_H & (FOC_{\mu_H}) \\ \nu_H(L_H - 1) = 0 & (CS_H) \end{cases}$$

Meanwhile, from the problem of the entrepreneur, the Euler equations collapse to user-cost conditions and the consumption FOC pins down C_E:

$$\begin{cases} r_s = \frac{1}{\beta} - 1 + \delta_s & (FOC_{k_s}) \\ r_e = \frac{\frac{1}{\beta} - 1 + \delta_e}{q} & (FOC_{k_e}) \\ r_{AI} = \frac{\frac{1}{\beta} - 1 + \delta_{AI}}{\zeta} & (FOC_{k_{AI}}) \\ C_E = \frac{1-\beta\eta}{(1-\eta)\mu_E} & (FOC_C) \end{cases}$$

and market clearing conditions give us:

$$\left\{ \begin{array}{c} M_L C_L + M_H C_H + M_E C_E \\ +\delta_s k_s + \frac{\delta_e}{q} k_e + \frac{\delta_{AI}}{\zeta} k_{AI} \end{array} \right\} = A k_s^{\alpha_1} \left\{ \begin{array}{c} \left[A_L \left(m_L^{\gamma_L} + k_e^{\gamma_L} \right)^{\frac{\lambda_L}{\gamma_L}} + \Psi_{AI} k_{AI}^{\lambda_L} \right]^{\frac{\phi}{\lambda_L}} \\ + \left[A_H \left(m_H^{\gamma_H} + k_e^{\gamma_H} \right)^{\frac{\lambda_H}{\gamma_H}} + \Psi_{AI} k_{AI}^{\lambda_H} \right]^{\frac{\phi}{\lambda_H}} \end{array} \right\}^{\frac{\alpha_2}{\phi}}$$

and

$$\begin{cases} m_L = M_L L_L & (MC_L) \\ m_H = M_H L_H & (MC_H) \end{cases}$$

while the problem of the firm, since it is static, remains the same in steady state.

Before we continue, let's consider the relaxed problem for low-skill worker and assume that

$\nu_L = 0$. From (FOC_{C_L}) and (FOC_{μ_L}), we have:

$$\frac{1-\beta\eta}{(1-\eta)w_L L_L} = \mu_L \tag{8}$$

while from (FOC_{L_L}), we have:

$$L_L^{\frac{1}{\tau}} = \mu_L w_L \tag{9}$$

then, from (9)/(8), we have:

$$\frac{L_L^{\frac{1}{\tau}}}{\frac{1-\beta\eta}{(1-\eta)w_L L_L}} = \frac{\mu_L w_L}{\mu_L}$$

simplifying it:

$$L_L = \left(\frac{1-\beta\eta}{1-\eta}\right)^{\frac{\tau}{1+\tau}}$$

since η, $\beta \in (0,1)$, $\frac{1-\beta\eta}{1-\eta} > 1$. Hence, in equilibrium, the restriction $L_L \leq 1$ is binding and $L_L = 1$. Hence $C_L = w_L$. Furthermore, since the problem for high-skill workers is symmetric to the low-skill worker one, we can easily show that $L_H = 1$ and $C_H = w_H$.

From the entrepreneur's problem, the Euler equations yield:

$$\begin{cases} r_s = \frac{1}{\beta} - 1 + \delta_s & (FOC_{k_s}) \\ r_e = \frac{\frac{1}{\beta}-1+\delta_e}{q} & (FOC_{k_e}) \\ r_{AI} = \frac{\frac{1}{\beta}-1+\delta_{AI}}{\zeta} & (FOC_{k_{AI}}) \\ C_E = \frac{1-\beta\eta}{(1-\eta)\mu_E} & (FOC_C) \end{cases}$$

while, from market clearing conditions, we have $M_L = m_L$, $M_H = m_H$ and:

$$\left\{ \begin{array}{c} M_L w_L + M_H w_H + \frac{M_E(1-\beta\eta)}{(1-\eta)\mu_E} \\ + \left(\frac{1}{\beta} - 1 + \delta_s\right) k_s + \frac{\frac{1}{\beta}-1+\delta_e}{q} k_e + \frac{\frac{1}{\beta}-1+\delta_{AI}}{\zeta} k_{AI} \end{array} \right\} = A k_s^{\alpha_1} \tilde{G}^{\frac{\alpha_2}{\phi}} \tag{10}$$

In steady state, we shut down time variation in shocks and taxes by setting $b_t \equiv 1$ and $T_{i,t} \equiv 0$. Given the no-saving assumption for workers, their budget constraints imply $C_i = w_i L_i$ for $i \in \{L, H\}$. Moreover, the intratemporal optimality conditions combined with habit formation imply that the unconstrained optimum would feature $L_i > 1$. Since labor supply is bounded above by one, the upper bound binds in steady state, so that $L_L = L_H = 1$ and therefore $C_L = w_L$ and $C_H = w_H$. On the entrepreneurs' side, the Euler equations collapse to static user-cost relationships, so that rental rates equal the standard discount-adjusted user costs: $r_j = (1/\beta - 1 + \delta_j)/p_j$, where p_j is the relative price of capital type j.

With these simplifications, the steady-state allocation and prices are pinned down by a system consisting of the firm's first-order conditions for factor demands together with the aggregate resource

constraint in (10). In this steady-state system, the endogenous unknowns are

$$\{\mu_E, w_L, w_H, k_s, k_e, k_{AI}\}, \tag{11}$$

where μ_E is the entrepreneur's marginal utility multiplier and $\{w_L, w_H, r_s, r_e, r_{AI}\}$ and $\{k_s, k_e, k_{AI}\}$ are factor prices and capital stocks consistent with firms' optimality and market clearing.

The resulting steady-state implications can be summarized as follows:

- Labor supplies are at the upper bound: $L_L = L_H = 1$.
- Workers consume their labor income: $C_i = w_i$ for $i \in \{L, H\}$.
- Rental rates satisfy user-cost conditions:

$$r_s = \frac{1}{\beta} - 1 + \delta_s, \qquad r_e = \frac{\frac{1}{\beta} - 1 + \delta_e}{q}, \qquad r_{AI} = \frac{\frac{1}{\beta} - 1 + \delta_{AI}}{\zeta}.$$

- Entrepreneur consumption is pinned down by the marginal-utility condition:

$$C_E = \frac{1 - \beta\eta}{(1-\eta)\mu_E}.$$

3.8 Simplified Production Function

For estimation and quantitative analysis, we work with a restricted version of the production function that shuts down the direct interaction between low-skill labor and both equipment and AI capital. This follows the approach of Krusell et al. (2000), who find that the data do not support an independent role for equipment in the low-skill bundle.

Specifically, we impose the following restrictions on the general production function in Section 3: (i) low-skill labor enters the top-level CES aggregator directly, without interacting with equipment or AI; and (ii) AI capital enters only through the high-skill bundle, where it combines with the high-skill–equipment composite. Under these restrictions, the production function becomes:

$$Y_t = A_t k_{st}^{\alpha_1} \left\{ \ell_{Lt}^{\phi} + \left[\mu_H \left(\xi_H \ell_{Ht}^{\gamma_H} + (1-\xi_H) k_{et}^{\gamma_H} \right)^{\frac{\lambda_H}{\gamma_H}} + (1-\mu_H) k_{AIt}^{\lambda_H} \right]^{\frac{\phi}{\lambda_H}} \right\}^{\frac{\alpha_2}{\phi}} \tag{12}$$

where $\ell_{Lt} = e^{\varphi_L} M_L L_{Lt}$ and $\ell_{Ht} = e^{\varphi_H} M_H L_{Ht}$ are efficiency-weighted aggregate labor inputs for low- and high-skill workers, respectively, with φ_j denoting the labor efficiency weight for skill type j. The parameter μ_H governs the weight of the high-skill–equipment composite relative to AI capital within the high-skill bundle, ξ_H governs the weight of high-skill labor relative to equipment within that composite, and $(1 - \mu_H)$ captures the AI share.

The nesting structure implies three distinct elasticities of substitution:

- $\sigma_\phi = \frac{1}{1-\phi}$: elasticity between the high-skill composite and low-skill labor at the top level;
- $\sigma_\lambda = \frac{1}{1-\lambda_H}$: elasticity between AI capital and the high-skill–equipment composite;
- $\sigma_\gamma = \frac{1}{1-\gamma_H}$: elasticity between high-skill labor and equipment.

When $\sigma_\lambda < 1$ (i.e., $\lambda_H < 0$), AI capital and the high-skill–equipment composite are *complements*: expanding AI is most productive when the skill–equipment bundle also expands. When $\sigma_\gamma < 1$ (i.e., $\gamma_H < 0$), high-skill labor and equipment are complements, consistent with the capital-skill complementarity finding of Krusell et al. (2000). Finally, when $\sigma_\phi \gg 1$ (i.e., ϕ close to 1), the two skill types are highly substitutable in production.

The calibration and all quantitative results in Sections 4–5 are based on this simplified production function.

4 Calibration and Estimation

4.1 Externally Calibrated Parameters

We split the parameterization into two groups. The first group consists of parameters that are set externally, either from standard calibration targets or from data averages. Table 1 reports these values.

4.2 Estimation of Production Function Parameters

The remaining parameters of the simplified production function (12)—α_1, α_2, ϕ, μ_H, ξ_H, γ_H, λ_H, and φ_L—are estimated by the Generalized Method of Moments (GMM). The estimation targets the following moments:

1. The wage-bill ratio: $\frac{w_{H,t}m_{H,t}}{w_{L,t}m_{L,t}}$;
2. The labor share: $\frac{w_{L,t}m_{L,t}+w_{H,t}m_{H,t}}{Y_t}$;
3. The equipment vs. structures no-arbitrage condition;
4. The AI vs. structures no-arbitrage condition;
5. The AI expenditure share.

Table 1: Externally calibrated parameters

Parameter	Value	Source
β	0.96	4% annual interest rate
M_E	0.1	2013 Survey of Consumer Finance (SCF)[8]
M_H	0.34	Adjusted share of high-skill in Ohanian et al. (2023)
M_L	0.56	Adjusted share of low-skill in Ohanian et al. (2023)
δ_s	0.03	2000–2024 Avg. Structures Depreciation rate from BEA
δ_e	0.17	2000–2019 Avg. Equipment Depreciation rate from Ohanian et al. (2023)
δ_{AI}	0.18	2000–2024 Avg. Computers Depreciation rate from BEA
η	0.37	Median value from Albanesi (2025)
τ	0.75	Aggregate Frisch Elasticity in Chetty et al. (2011)
q	0.78	2000–2024 Avg. Price of Equipment in Ohanian et al. (2023) (2019 US$=1)
ζ	0.91	2000–2024 Avg. Price of Computers (2019 US$=1)
A	1	Normalization
φ_H	2	High-skill labor efficiency weight, same as Krusell et al. (2000)

The instrument set consists of $\{k_{et}, k_{st}, k_{e,t-1}, k_{s,t-1}, P_{e,t-1}, \text{Year}\}$. Table 2 reports the point estimates and standard errors.

Table 2: Estimated parameters (GMM)

Parameter	Estimate
ϕ	0.94
	(0.023)
α_1	0.38
	(0.013)
α_2	0.57
	(0.012)
μ_H	0.97
	(0.005)
ξ_H	0.83
	(0.022)
γ_H	−0.19
	(0.097)
λ_H	−0.14
	(0.028)
φ_L^*	0.80

Standard errors in parentheses.

4.3 Key Takeaways from Estimation

The estimates yield the following implied elasticities of substitution:

- $\sigma_\lambda = \frac{1}{1-\lambda_H} \approx 0.88$: AI and the high-skill–equipment composite are *complements* ($\sigma_\lambda < 1$).
- $\sigma_\gamma = \frac{1}{1-\gamma_H} \approx 0.84$: High-skill labor and equipment are complements ($\sigma_\gamma < 1$), consistent with the capital-skill complementarity finding of Krusell et al. (2000).
- $\sigma_\phi = \frac{1}{1-\phi} \approx 16.7$: High- and low-skill composites are highly substitutable ($\sigma_\phi \gg 1$).

Several features of the estimates are noteworthy. First, the AI weight is still quite small: $(1-\mu_H) \approx 3\%$, reflecting the fact that AI capital remains a small fraction of the total capital stock. Second, the human capital weight within the high-skill bundle is high ($\xi_H \approx 0.83$), indicating that high-skill labor still dominates the skill–equipment composite. Third, and most importantly for the quantitative exercises that follow, the complementarity between AI and the high-skill–equipment bundle ($\sigma_\lambda < 1$) implies that expanding AI is most productive when high-skill labor and equipment also expand.

5 Quantitative Results

This section reports the results of counterfactual experiments designed to assess how advances in AI affect macroeconomic aggregates, factor prices, and the income distribution. We compute perfect-foresight transition dynamics starting from the calibrated steady state, in which agents learn at time $t = 0$ that a permanent shock will take effect and the economy transitions to a new steady state.

5.1 Counterfactual Experiments

We consider three permanent-shock experiments:

1. **Double AI usage share.** Increase the AI weight in the CES aggregator so that $(1-\mu_H)$ rises from 3% to 6%, holding prices fixed. This captures a scenario in which AI services become more easily deployed or integrated into production.
2. **Halve AI price.** Reduce the relative price of AI capital by 50%—equivalently, double ζ—holding the production function parameters fixed. This represents continued rapid investment-specific technological change in AI hardware.
3. **Both shocks simultaneously.** Apply both changes at once, to study the interaction between demand-side and supply-side AI improvements.

Table 3 summarizes the long-run (new steady-state) effects of the first two experiments.

Table 3: Counterfactual experiments: new steady-state effects

	Initial SS	**Double Share**	**Halve Price**
Output Y (% Δ)	—	-0.95	$+1.58$
Consumption C (% Δ)	—	-1.91	$+1.63$
AI Capital K_{AI} (% Δ)	—	$+81$	$+88$
Skill Premium w_H/w_L	2.34	2.23	2.41

5.2 Mechanisms and Transition Dynamics

The two individual experiments produce strikingly different macroeconomic outcomes, despite both generating large increases in the AI capital stock. The contrast is best understood as a demand-side shock (Experiment 1) versus a supply-side shock (Experiment 2). Accordingly, the two experiments generate strikingly different transition paths. The share increase (Experiment 1) produces a V-shaped adjustment: the production function changes immediately but capital stocks are pre-determined, creating a sharp capital–technology mismatch. With $\sigma_\lambda < 1$, this mismatch is costly—inputs must be used in near-fixed proportions, and tilting weights toward the scarce input creates a bottleneck. The price reduction (Experiment 2) produces a smooth, monotone expansion: the production function is unchanged at $t = 1$, and the shock only affects the resource cost of AI investment, an intertemporal effect that shifts the investment plan without disrupting current production.

We trace these mechanisms in turn.

5.2.1 Output and Aggregate Consumption

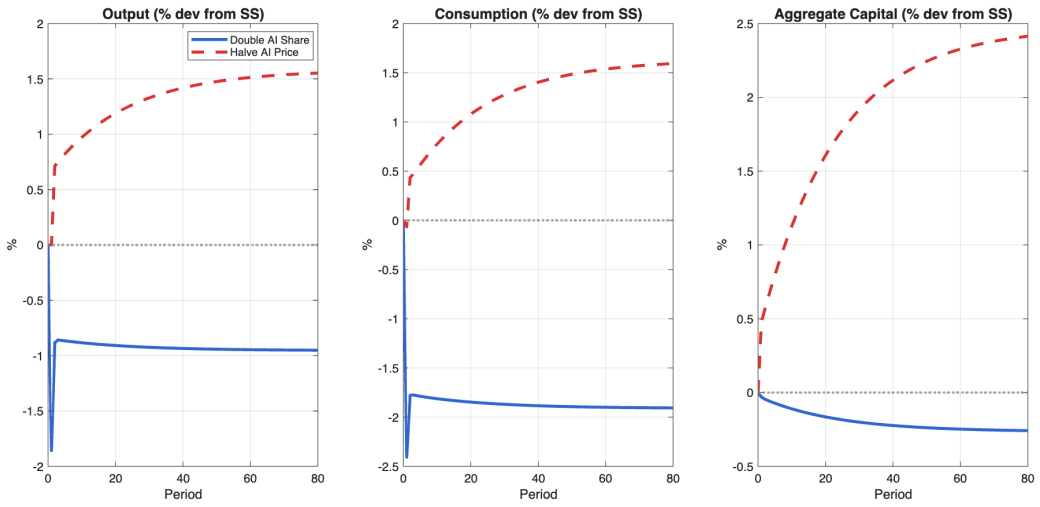
Under the share increase (Experiment 1), the production function changes on impact even though capital stocks have not yet adjusted. With $\sigma_\lambda \approx 0.88 < 1$ (complements), increasing the CES weight on the scarce factor—AI capital is only about 1% of total capital value—*reduces* the high-skill composite X_H. Output falls 1.86% on impact and, even after full capital adjustment, remains 0.95% below the initial steady state. AI capital is simply too small a share of the total for doubling its weight to compensate for the loss in the dominant equipment–skill channel. Aggregate consumption follows a similar path, falling 1.91% in the new steady state.

Under the price reduction (Experiment 2), output is unchanged on impact ($t = 1$): the same weights, capital stocks, and labor produce the same output. But the resource constraint has loosened—the same dollar of investment now buys twice as much AI capital. The entrepreneur immediately redirects spending toward AI, and more AI capital raises the high-skill composite X_H,

which raises the aggregate composite $\tilde{G}$ and hence Y. Output rises 1.58% and consumption rises 1.63% in the new steady state.

Figure 8 displays the perfect-foresight transition paths for output, consumption, and aggregate capital, illustrating the sharp V-shaped adjustment of Experiment 1 against the smooth expansion of Experiment 2.

Figure 8: Transition dynamics: Output, Consumption, and Aggregate Capital



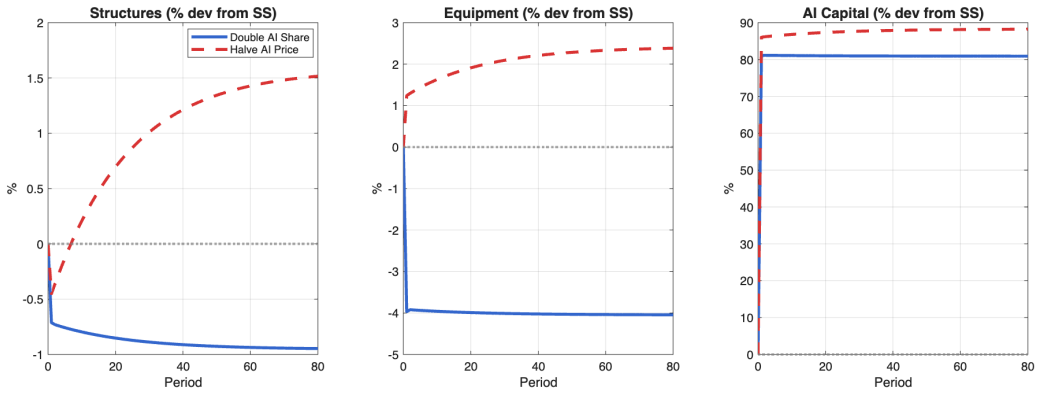
5.2.2 Capital Composition

The two experiments reshape the capital stock in fundamentally different ways. Under the share increase (Experiment 1), the entrepreneur reallocates investment: the marginal product of AI capital has jumped (the $(1-\mu_H)$ term in $\partial Y/\partial k_{AI}$ doubled), so AI investment surges and k_{AI} rises 81%. Meanwhile, the marginal product of equipment has fallen (μ_H dropped from 0.97 to 0.94), so equipment investment drops and k_e falls 4%. The Euler equations pin rental rates to user costs (MPK = user cost), which are independent of μ_H, but the composition of capital has shifted: more resources go to AI (small, expensive, high-depreciation) and fewer to equipment (large, productive).

Under the price reduction (Experiment 2), the parameter ζ governing the efficiency of AI capital production doubles permanently, so the relative investment price $P_{AI} = 1/\zeta$ halves. This is an investment-specific technological change (ISTC) in the spirit of Greenwood et al. (1997) and Krusell et al. (2000): a technology parameter that enters only through the resource constraint. k_{AI} surges 88%, and crucially, this surge is half as costly in goods terms (x_{AI}/ζ is the resource cost). The economy can accumulate massive AI capital without crowding out other investment: k_e actually *rises* 2.4% and k_s rises 1.6%. The Euler equation now requires a *lower* rental rate on AI capital ($r_{AI} = (1/\beta - 1 + \delta_{AI})/\zeta^{\text{new}}$), so k_{AI} rises until its marginal product falls to the new, lower required return.

Figure 9 shows these contrasting capital dynamics: crowding out under the share increase versus broad-based capital deepening under the price reduction.

Figure 9: Transition dynamics: Capital Stocks by Type



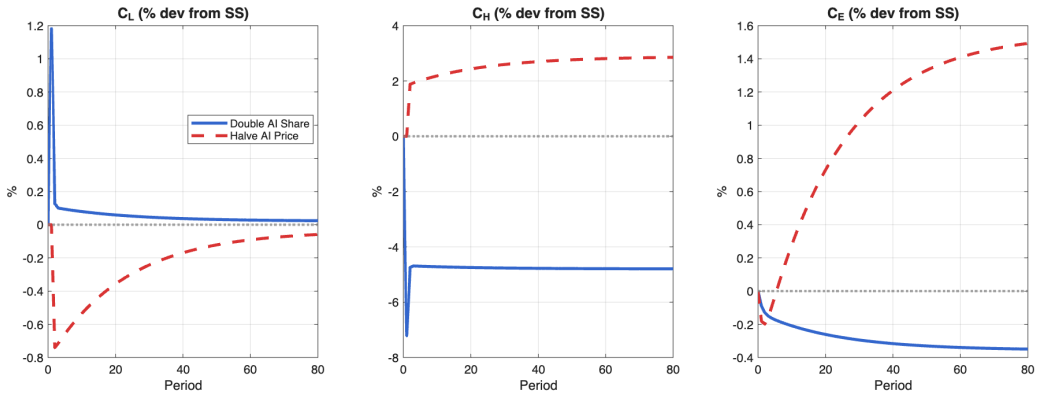
5.2.3 Consumption by Agent Type

Under the share increase (Experiment 1), entrepreneur consumption C_E adjusts gradually due to habit formation. On impact, C_E dips only modestly (-0.09% at $t = 1$) and then declines monotonically to -0.35% in the new steady state, reflecting the permanently lower output and worse capital composition. The smooth adjustment contrasts with output, which drops sharply at impact: the entrepreneur absorbs the shock by reallocating investment rather than cutting consumption abruptly.

Under the price reduction (Experiment 2), entrepreneur consumption traces a classic "invest-then-enjoy" path. C_E initially dips below the steady state (-0.18% at $t = 1$) as the entrepreneur redirects resources toward the now-cheaper AI capital. This investment boom absorbs resources for about 10 periods before the new capital stock generates enough additional output to push C_E above its initial level. In the terminal steady state, C_E rises $+1.55\%$: the larger, cheaper AI capital stock generates higher returns, and the positive spillovers to equipment and structures further increase total capital income.

Figure 10 displays the transition paths for consumption by agent type, highlighting how the timing and direction of entrepreneur consumption diverge across the two experiments.

Figure 10: Transition dynamics: Consumption by Agent Type



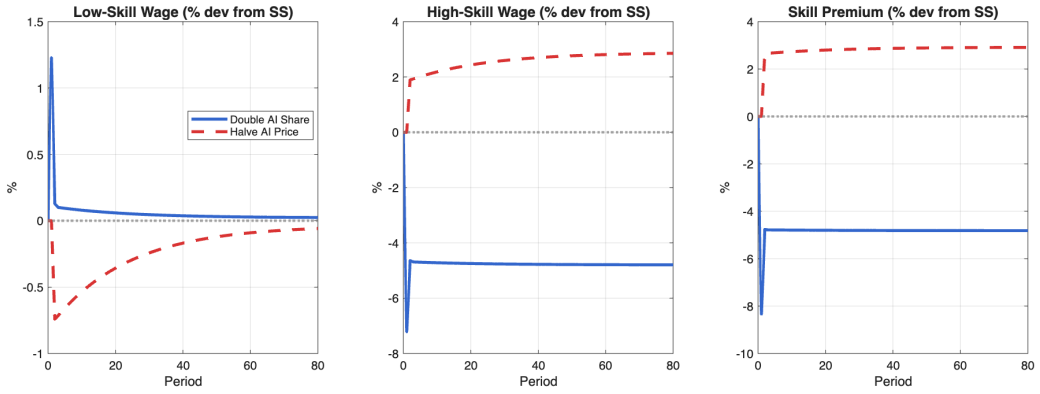
5.2.4 Wages and the Skill Premium

Under the share increase (Experiment 1), the high-skill wage falls (-4.8%) because the lower weight μ_H directly reduces the marginal product of the high-skill–equipment composite (i.e., the inner CES bundle $\xi_H \ell_H^{\gamma_H} + (1-\xi_H) k_e^{\gamma_H}$), compressing the skill premium from 2.34 to 2.23. Low-skill wages barely change ($+0.02\%$): with $\sigma_\phi \approx 16.7$, the low- and high-skill bundles are near-perfect substitutes in the aggregate composite $\tilde{G}$, so changes in the high-skill bundle composition have limited spillover to low-skill wages.

Under the price reduction (Experiment 2), the complementarity channel ($\sigma_\lambda < 1$) works in the opposite direction: more AI capital raises the marginal product of the high-skill–equipment composite, benefiting high-skill workers. The high-skill wage rises ($+2.9\%$) and the skill premium increases from 2.34 to 2.41—the classic capital–skill complementarity mechanism now operating through AI. Low-skill workers are again nearly unaffected (w_L changes by -0.04%), insulated by the high elasticity $\sigma_\phi \approx 16.7$ between skill bundles and by the general-equilibrium offset through structures adjustment.

Figure 11 confirms these patterns, showing the symmetric divergence in the skill premium across the two experiments.

Figure 11: Transition dynamics: Wages and Skill Premium



5.3 Combined Experiment and Super-Additivity

Experiment 3 applies both shocks simultaneously. The key finding is that the combined effect is *super-additive*: the output gain from both shocks together (+2.02%) far exceeds the sum of the individual effects ($-0.95\% + 1.58\% = +0.63\%$). The interaction term is +1.40 percentage points—more than twice the linear sum.

5.3.1 Intuition: Constraint Removal

Each individual experiment is constrained by what the other does not provide:

- **Experiment 1 alone:** The production function now "wants" more AI, but AI is expensive. The higher weight makes AI investment more attractive, but the high price limits how much k_{AI} the economy can actually accumulate. The economy is stuck in a bad spot: it "needs" AI but cannot afford enough.
- **Experiment 2 alone:** AI is cheap, but its weight is only 3%. Even with massive AI accumulation, the contribution is capped by the small $(1 - \mu_H) = 0.03$ in the CES aggregator. Cheap AI helps, but only modestly because the production function does not put much weight on it.

When both shocks hit simultaneously, each channel removes the binding constraint on the other. The price drop makes it cheap to accumulate the AI capital that the higher weight demands. The higher weight makes the price drop far more impactful because each unit of cheap AI now contributes more to production. The result is a virtuous cycle that neither channel can achieve alone.

5.3.2 Theoretical Foundation

The super-additivity in our setting is a form of *within-input supermodularity*: both the weight and the inverse price of the *same* input change simultaneously. This is distinct from the between-input complementarity ($\sigma < 1$) studied in the Milgrom–Roberts tradition. In a CES aggregator embedded in a general-equilibrium model with capital accumulation, the value function $V(\alpha, 1/p_A)$ satisfies increasing differences for all $\sigma > 0$: a higher weight raises the marginal product of AI, a lower price raises the equilibrium quantity of AI, and the output gain from cheaper AI is amplified when the weight is higher. This within-input mechanism is a general CES property, not specific to complementarity—our robustness analysis in Section 5.4 confirms that super-additivity holds (and indeed strengthens) even when $\sigma_\lambda > 1$.

The practical interpretation connects to the general-purpose technology (GPT) literature. In the framework of Bresnahan and Trajtenberg (1995), the price drop is the GPT improving, and the share increase is the development of complementary applications. Our results provide a quantitative CES mechanism for these informal GPT complementarities: AI policies targeting both adoption (demand: subsidies, infrastructure, regulation) and cost reduction (supply: R&D, semiconductor policy) simultaneously will be far more effective than either alone.

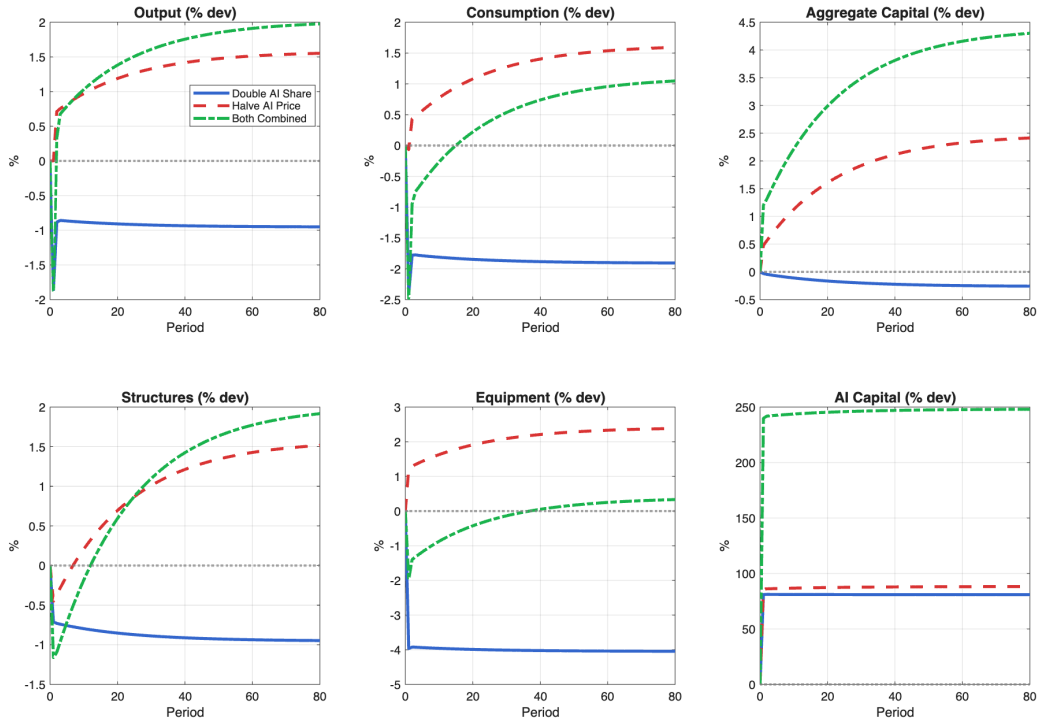
Table 4 reports the full set of results, including the hypothetical linear sum for comparison.

Table 4: Combined experiment: super-additivity

	Initial SS	**Dbl. Share**	**Halve Price**	**Both**	**Exp1+Exp2**
Aggregate (% Δ from initial SS)					
Output Y	—	−0.95	+1.58	+2.02	+0.63
Consumption C	—	−1.91	+1.63	+1.11	−0.28
AI Capital K_{AI}	—	+81	+88	+248	+169
Distributional					
Skill Premium w_H/w_L	2.34	2.23	2.41	2.35	—
Labor Share	0.495	0.486	0.495	0.486	—
C_E (Entrepreneur)	26.50	26.41	26.91	27.19	—

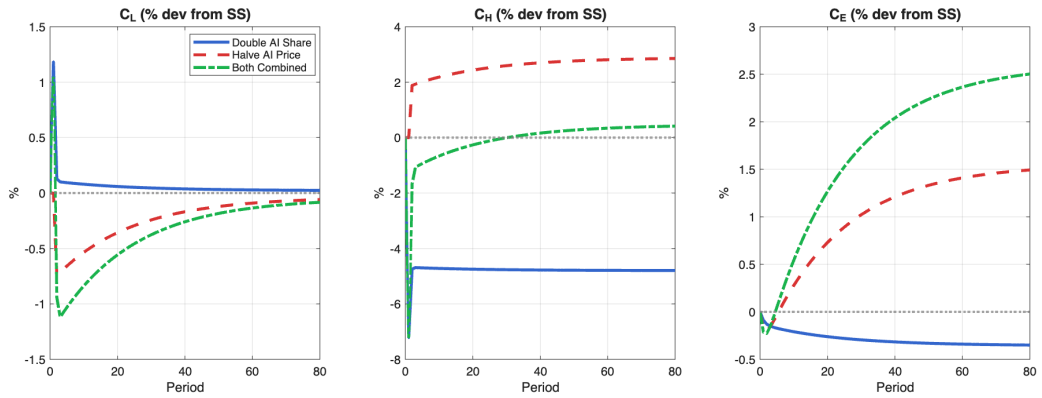
The output super-additivity is striking: the combined experiment produces more than 3× the output growth of the linear sum. The contractionary effect of Experiment 1 is entirely eliminated—what was −0.95% becomes part of a +2.02% expansion. Figure 12 displays the transition dynamics for output, aggregate consumption, and the capital stocks under the combined shock.

Figure 12: Combined experiment: Transition dynamics overview



The combined experiment is nearly skill-premium neutral (2.35 vs. 2.34), as the two effects on w_H systematically offset each other. Entrepreneurs are the biggest winners: their consumption rises +2.6% in the combined scenario because they capture the returns from the massive, cheap AI capital accumulation. The labor share falls (from 0.495 to 0.486) because the recomposition of capital favors returns to capital owners. Figure 13 traces these distributional dynamics over the transition.

Figure 13: Combined experiment: Consumption by Agent Type



5.4 Robustness

5.4.1 Varying σ_λ

To assess the sensitivity of our results to the estimated complementarity between AI and the high-skill–equipment bundle, we vary λ_H to produce three alternative values of σ_λ: the baseline estimate (0.88, complements), moderate substitutes (1.20), and strong substitutes (1.49). Table 5 reports the results.

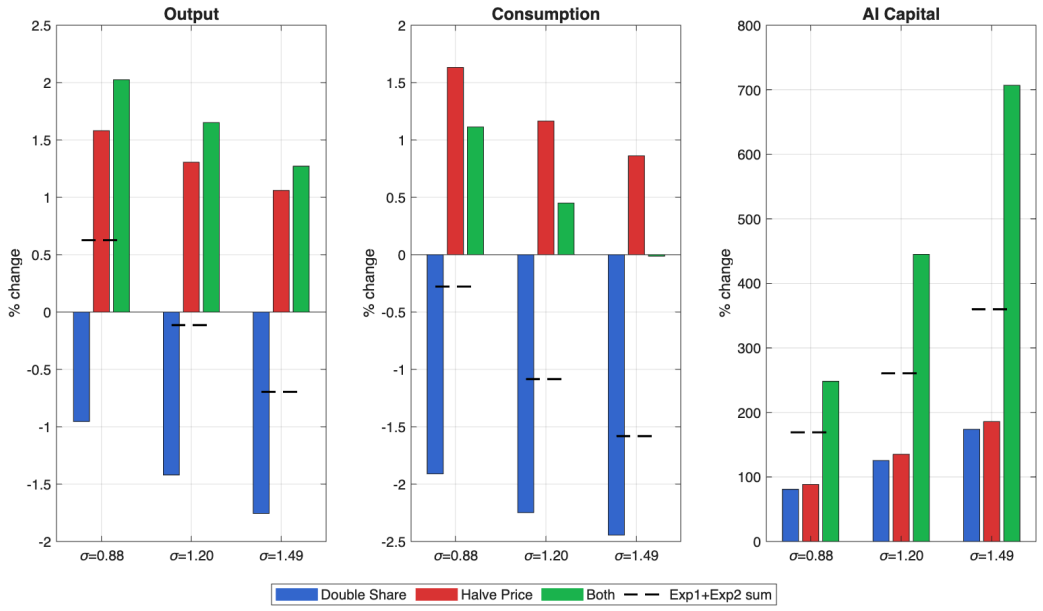
Table 5: Robustness: Results across σ_λ

	$\sigma_\lambda = 0.88$ (complements)	$\sigma_\lambda = 1.20$ (mod. subs.)	$\sigma_\lambda = 1.49$ (strong subs.)
Output Y (% Δ)			
Double Share	-0.95	-1.42	-1.76
Halve Price	$+1.58$	$+1.31$	$+1.06$
Both	$+2.02$	$+1.65$	$+1.27$
Skill Premium w_H/w_L (level)			
Initial SS	2.34	2.38	2.41
Both	2.35	2.36	2.36

The share increase is always contractionary, and the contraction deepens with higher substitutability ($-0.95\% \to -1.42\% \to -1.76\%$). With higher σ, AI substitutes more easily for the high-skill–equipment composite, so the economy loses more from reweighting away from the dominant equipment–skill channel. The price drop is always expansionary, but less so with substitutes: the output gain shrinks ($+1.58\% \to +1.31\% \to +1.06\%$) as σ_λ rises. With complements, cheaper AI strongly boosts the marginal product of the high-skill–equipment composite (and hence the whole production chain); with substitutes, AI simply replaces it rather than amplifying it.

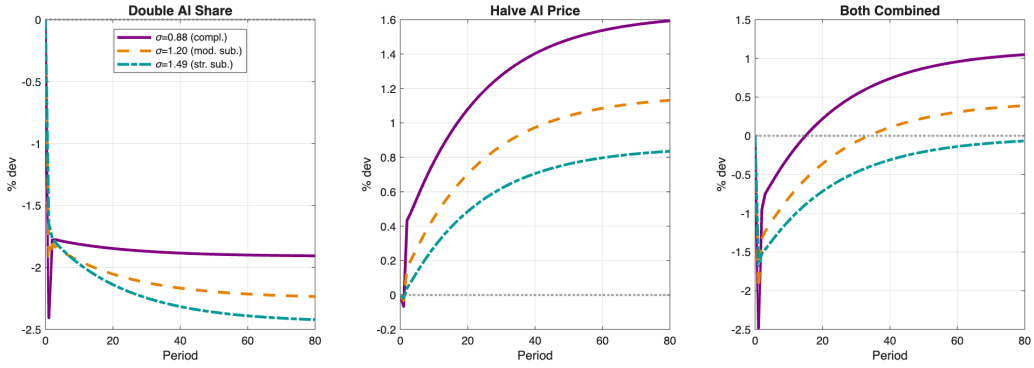
The surprising result is that super-additivity strengthens with substitutability. Even though individual experiments perform worse with higher σ_λ, the interaction gain grows: the super-additivity term rises from $+1.40$ pp at $\sigma_\lambda = 0.88$ to $+1.97$ pp at $\sigma_\lambda = 1.49$. The reason is that with substitutes, Experiment 1 alone is more contractionary and Experiment 2 alone is less expansionary, leaving "more room" for the interaction to matter. This confirms that the constraint-removal mechanism is a property of the CES functional form, not just of complementarity. Figure 14 displays this pattern.

Figure 14: Super-additivity across σ_λ



Finally, the skill premium is always near-neutral in the combined experiment: across all three σ_λ values, the combined experiment lands close to the initial skill premium (2.35, 2.36, 2.36), as the share-increase and price-drop effects on w_H systematically offset each other. Figure 15 shows the consumption dynamics across the three elasticity values, while Appendix Figures A.12–A.15 provide additional detail on output, entrepreneur consumption, the skill premium, and wages.

Figure 15: Consumption dynamics across σ_λ



5.4.2 Varying δ_{AI}

Because AI hardware depreciates rapidly and the precise rate is uncertain, we also assess robustness to the AI depreciation rate. We consider $\delta_{AI} \in \{0.10, 0.18, 0.25\}$, spanning from low to high obsolescence. Higher depreciation amplifies the contractionary effect of the share increase and slightly raises the expansionary effect of the price reduction, but super-additivity remains stable

across all values (+1.36 to +1.42 pp). The full results are reported in Appendix D.

6 Road Ahead

Government Sector and Policy Counterfactuals We plan to reintroduce a government sector and study policy counterfactuals that affect the user cost of capital and the speed of labor-market adjustment. Tax and transfer policies—such as investment tax credits for AI, payroll tax adjustments, or redistributive transfers—can alter the speed of AI adoption and shape its distributional consequences. We will also incorporate standard preference and technology shocks to examine how AI-specific improvements interact with business-cycle fluctuations, with particular attention to wages, rental rates, and factor shares.

Two-Sector Extension We intend to evaluate a simple split between an AI-producing and an AI-using sector, which may help connect the construction of AI prices and quantities more directly to equilibrium outcomes. This extension would allow us to study the feedback between the supply of AI hardware, its quality-adjusted price, and the diffusion of AI capital into downstream production.

7 Conclusion

This paper develops a quantitative framework to study how advances in artificial intelligence affect production, factor prices, and the income distribution. We propose a theory of AI-augmented complementarity that builds on the capital-skill complementarity approach of Krusell et al. (2000). AI is modeled as a distinct capital input, alongside structures and equipment, that can differentially substitute for or complement low- and high-skill labor. We also distinguish entrepreneurs from workers to capture concentrated capital ownership and its distributional implications.

We embed this mechanism in a classic real business cycle environment, extending recent work such as Albanesi (2025). A key component is measurement: because AI hardware improves rapidly, we construct a quality-adjusted series for AI prices and capital using a hedonic approach in the spirit of Gordon (1990). We estimate the production function parameters by GMM using data from 2000–2024, and find that AI capital and the high-skill–equipment composite are complements ($\sigma_\lambda \approx 0.88$), consistent with the broader capital-skill complementarity literature.

Our counterfactual experiments reveal that the macroeconomic effects of AI depend critically on the channel through which AI improves. Doubling the AI usage share is contractionary—because complementarity with a still-scarce input reduces the effective composite—while halving the AI price is expansionary. Crucially, the combined effect is super-additive: the two channels reinforce

each other, producing output gains that far exceed the sum of the individual effects. This finding is robust across a range of substitution elasticities and AI depreciation rates, and highlights the importance of jointly modeling demand-side adoption and supply-side cost reductions in AI. The framework provides a platform for studying how tax and transfer policies might shape AI adoption and its distributional consequences.

References

Acemoglu, Daron (2025). "The simple macroeconomics of AI*." *Economic Policy*, 40(121), pp. 13–58. doi:10.1093/epolic/eiae042. URL https://doi.org/10.1093/epolic/eiae042.

Acemoglu, Daron and David Autor (2011). "Skills, tasks and technologies: Implications for employment and earnings." In *Handbook of labor economics*, volume 4, pp. 1043–1171. Elsevier.

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo (2022). "Artificial Intelligence and Jobs: Evidence from Online Vacancies." *Journal of Labor Economics*, 40(S1), pp. S293–S340. doi:10.1086/718327. URL https://www.journals.uchicago.edu/doi/10.1086/718327.

Acemoglu, Daron and Pascual Restrepo (2022). "Tasks, automation, and the rise in us wage inequality." *Econometrica*, 90(5), pp. 1973–2016.

Ai, Epoch (2025). "Data on Machine Learning Hardware." URL https://epoch.ai/data/machine-learning-hardware.

Albanesi, Stefania (2025). "Changing business cycles: The role of women's employment." *American Economic Journal: Macroeconomics*, 17(4), pp. 345–390.

Baily, Martin, David Byrne, Aidan Kane, and Paul Soto (2025). "Generative ai at the crossroads: Light bulb, dynamo, or microscope?" *arXiv preprint arXiv:250514588*.

Berg, Andrew, Edward F. Buffie, and Luis-Felipe Zanna (2018). "Should we fear the robot revolution? (The correct answer is yes)." *Journal of Monetary Economics*, 97, pp. 117–148. doi:10.1016/j.jmoneco.2018.05.014. URL https://linkinghub.elsevier.com/retrieve/pii/S0304393218302204.

Bresnahan, Timothy F. and M. Trajtenberg (1995). "General purpose technologies 'Engines of growth'?" *Journal of Econometrics*, 65(1), pp. 83–108. doi:10.1016/0304-4076(94)01598-T. URL https://www.sciencedirect.com/science/article/pii/030440769401598T.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen (2025). "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence."

Caunedo, Julieta, David Jaume, and Elisa Keller (2023). "Occupational exposure to capital-embodied technical change." *American Economic Review*, 113(6), pp. 1642–1685.

Caunedo, Julieta and Elisa Keller (2022). "Technical change and the demand for talent." *Journal of Monetary Economics*, 129, pp. 65–88.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber (2011). "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins." *American Economic Review*, 101(3), pp. 471–475. doi:10.1257/aer.101.3.471. URL https://www.aeaweb.org/articles?id=10.1257/aer.101.3.471.

Cummins, Jason G. and Giovanni L. Violante (2002). "Investment-Specific Technical Change in the United States (1947–2000): Measurement and Macroeconomic Consequences." *Review of Economic Dynamics*, 5(2), pp. 243–284. doi:10.1006/redy.2002.0168. URL https://www.sciencedirect.com/science/article/pii/S1094202502901687.

DiCecio, Riccardo (2009). "Sticky wages and sectoral labor comovement." *Journal of Economic Dynamics and Control*, 33(3), pp. 538–553. doi:https://doi.org/10.1016/j.jedc.2008.08.003. URL https://www.sciencedirect.com/science/article/pii/S0165188908001449.

Fisher, Franklin M., Joen E. Greenwood, and John J. McGowan (1985). *Folded, Spindled, and Mutilated: Economic Analysis and U.S.v. IBM*. Regulation of Economic Activity. MIT Press, Cambridge, MA, USA.

Gordon, Robert J. (1990). *The Measurement of Durable Goods Prices*. University of Chicago Press. URL https://www.nber.org/books-and-chapters/measurement-durable-goods-prices, backup Publisher: National Bureau of Economic Research Type: Book.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell (1997). "Long-Run Implications of Investment-Specific Technological Change." *The American Economic Review*, 87(3), pp. 342–362. URL https://www.jstor.org/stable/2951349, publisher: American Economic Association.

Guerrón-Quintana, Pablo, Tomoaki Mikami, and Jaromir Nosal (2024). "The macroeconomic implications of the gen-ai economy." *Available at SSRN 4989979*.

Kopytov, Alexandr, Nikolai Roussanov, and Mathieu Taschereau-Dumouchel (2018). "Short-run pain, long-run gain? recessions and technological transformation." *Journal of Monetary Economics*, 97, pp. 29–44.

Krusell, Per, Lee E. Ohanian, Jose-Victor Rios-Rull, and Giovanni L. Violante (2000). "Capital-skill Complementarity and Inequality: A Macroeconomic Analysis." *Econometrica*, 68(5), pp. 1029–1053. doi:10.1111/1468-0262.00150. URL http://doi.wiley.com/10.1111/1468-0262.00150.

Kudoh, Noritaka and Hiroaki Miyamoto (2025). "Robots, AI, and unemployment." *Journal of Economic Dynamics and Control*, 174, p. 105,069. doi:10.1016/j.jedc.2025.105069. URL https://linkinghub.elsevier.com/retrieve/pii/S0165188925000351.

Ohanian, Lee E., Musa Orak, and Shihan Shen (2023). "Revisiting capital-skill complementarity, inequality, and labor share." *Review of Economic Dynamics*, 51, pp. 479–505. doi:10.1016/j.red.2023.05.002. URL https://linkinghub.elsevier.com/retrieve/pii/S1094202523000182.

Pilz, Konstantin F., James Sanders, Robi Rahman, and Lennart Heim (2025). "Trends in AI Supercomputers." doi:10.48550/arXiv.2504.16026. URL http://arxiv.org/abs/2504.16026, arXiv:2504.16026 [cs].

A Construction of Price Deflators

We follow the procedure described in Cummins and Violante (2002) to construct quality-adjusted price deflators. First of all, we gather the quality-adjusted price indexes constructed by Gordon (1990) for the period 1947–1983.[9] While there were changes on NIPA Table 5.6, we start discussing the approach to the NIPA Lines that stayed the same. We follow Cummins and Violante (2002) and run the following econometric specification in order to extend the data up to 2024:

$$\log(p_t^{i_j^*}) = c + \beta_1 t + \beta_2 \log(p_t^{i_j}) + \beta_3 \log(p_{t-1}^{i_j}) + \beta_4 \Delta y_{t-1} + \varepsilon_t^j \tag{13}$$

where $p_t^{i_j^*}$ is Gordon's quality-adjusted price index for asset category j, c is a constant, t is a linear time trend, $p_t^{i_j}$ and $p_{t-1}^{i_j}$ are, respectively, the current and lagged values of the NIPA price index, Δy_{t-1} is the growth rate of lagged GDP, and ε_t^j is a white noise. Then, using the coefficient estimates, we extrapolate for 1984-2024 the quality-adjusted price level for each asset from the original sample.

Moreover, while in equation (13) we assume only one lag, we follow a mixture of Akaike and Scharwz information criteria in order to select the optimal order lag in each equation.

Moreover, there are some asset categories in which NIPA's classification changed over time. An important case is the Information Processing Equipment and Software (IPES). The BEA now distinguishes explicitly among computers and peripherals and other office and accounting machinery. In contrast, office, computing, and accounting machinery (OCAM) was a single category in Gordon (1990).

[9]The majority of the durable aggregates' price indexes can be found in Table B.18 in Appendix B. Price indexes from Ships and Boats, Trucks, Special industry machinery, and Mining and Oil field machinery from Appendix C. Finally, we get information on quality-adjusted price index for computers from Gordon's Table 6.12, column 2

To circumvent this issue, to construct the quality-adjusted price index for computers, we again follow Cummins and Violante (2002) and combine two data sources. First, we use Gordon's quality-adjusted index for computers for 1947-1983 (Table 6.12, column 2) to pin down the values for 1947-1958. Second, we use NIPA index for the hedonic-based quality-adjusted price index for computers and peripherals for the period 1959-2024. In order to mend these two series (which have different basis), we start by rebasing Gordon's quality adjusted index such that the year base is 1958 (=100 in 1958). Second, we construct an adjustment factor $AdjFact = \frac{\text{Gordon}_{t=1958}}{\text{NIPA}_{t=1958}}$. Finally, we multiply all the observations for the NIPA index by the adjustment factor. Finally, the mended series has base year 1958, so we rebase to year 2009.

Other series also have shown changes over time. For example, Instruments and Photocopy were split into Medical equipment and instruments, Nonmedical instruments, and Photocopy and related equipment. Moreover, the current NIPA asset categories do not separate tractors from neither Agricultural Machinery nor construction machinery in most cases, while in Gordon's quality adjusted price index we do have that distinction. Finally, Other equipment asset category in Gordon's Appendix B is quite different from either the Other equipment (NIPA Line 29 in Table 5.5.4) or Other (NIPA Line 44 in Table 5.5.4).

In order to aggregate asset categories to match the different definitions, we use the Tornquist index. We first compute the nominal investment shares of each asset in the category for each year. The share of asset j is the ratio of the current dollar value of investment in asset j and the current dollar value of total investment in the asset category.

Let $s_t^{i_j}$ be the nominal share for investment good $j \in \{1, ..., \mathcal{J}\}$ and let $p_t^{i_j^*}$ be the corresponding quality-adjusted price index for investment of type j. Then, the change in the aggregate quality-adjusted price index for the asset category is:

$$\Delta p_t^{i_e^*} = \sum_{j=1}^{\mathcal{J}} \log\left(\frac{p_t^{i_j^*}}{p_{t-1}^{i_j^*}}\right)\left(\frac{s_t^{i_j} + s_{t-1}^{i_j}}{2}\right) \tag{14}$$

and the level of the price index is recovered recursively as:

$$p_t^{i_e^*} = p_{t-1}^{i_e^*} \times exp(\Delta p_t^{i_e^*})$$

While Gordon's study has no price index for software, NIPA provides price indexes beginning in 1959. Apart from the headline number, we have price indexes for own-account software,[10] pre-packaged, and custom software.[11] Following Cummins and Violante (2002), we use the NIPA's quality-adjusted software price index by itself.

[10] Software developed internally by firms themselves.

[11] Software tailored to the specifications of firms and purchased externally by these firms.

A.1 AI Price Index and Capital

We construct the AI price deflator following the methodology described by Gordon (1990)'s chapter 6. In particular, we run a hedonic price regression for AI:

$$\log p_{i,t} = \beta_0 + \sum_{s=1}^{N} \delta_s D_s + \sum_{j=1}^{m} \log y_{j,i,t} + u_{i,t} \tag{15}$$

where $p_{i,t}$ is the release price of processor i at time t, $y_{j,i,t}$ is the quality of characteristic j for processor i at time t, and D_t are year dummies. Since our database focuses on release dates and prices, we are focusing on new processors. As a result, our hedonic regression traces the technology frontier, the preferred methodology according to Fisher et al. (1985). This is also the case for Gordon's Computer and Peripheral's hedonic price index.

Our main database source comes from Epoch Ai (2025) that presents key data on over 170 AI accelerators, such as graphics processing units (GPUs) and tensor processing units (TPUs), used to develop and deploy machine learning models in the deep learning era. Unfortunately, only a sub-set of these processing units are commercialized and have price data. We combine the price data with Epoch AI's price data accessed through airtable and discussed by Pilz et al. (2025). At the end, we have price data on 61 processing units over the period starting in 2008, 56 starting in 2013.

In terms of the quality characteristics that we control for, we focus on *ML OP/s*, the maximum performance in any format 8 bits or wider, in units of FLOP/s or OP/s, *Memory size per board*, the hardware's amount of memory, in bytes, and Thermal design power (*TDP*), the theoretical maximum power that can be dissipated as heat.[12] We focus on these characteristics, not only because they are relevant to explain variation in prices, but there are few missing observations in these variables, boosting our sample size.

We present the regression results in Table A.1. In our analysis, we use specification (3). Then, following Gordon (1990), we calculate the log-ratio of actual prices and base-year imputed prices. We then drop the most overpriced processor, i.e., the ones above the 90th percentile.[13] We then re-run the model (3) from Table A.1. Following the "dummy variable method", the hedonic price index is the antilog of the coefficients for the year dummies. We assume the year dummy is equal to zero in years in which the year dummy is not statistically significant.

[12] In theory, this is the maximum sustainable power draw for a given chip.

[13] In Gordon (1990), he drops the processors in which the log-ratio is higher than 1.5 times the standard error. Given our small sample size, we decided to be more lenient.

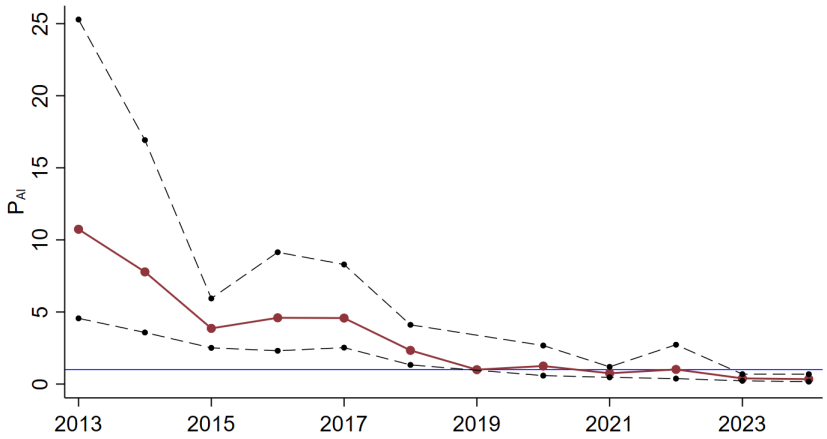
Table A.1: Hedonic Regressions: AI Hardware

	$\log(p_i)$		
	(1)	(2)	(3)
D_{2013}	2.48***	2.17***	2.73***
	(0.63)	(0.39)	(0.49)
D_{2014}	1.59**	1.20***	1.76***
	(0.65)	(0.28)	(0.47)
D_{2015}	1.98**	1.66***	1.91***
	(0.75)	(0.60)	(0.55)
D_{2016}	1.97***	1.01***	1.28***
	(0.48)	(0.36)	(0.40)
D_{2017}	1.20***	1.02***	1.35***
	(0.29)	(0.30)	(0.34)
D_{2018}	1.03**	0.51*	0.73**
	(0.46)	(0.26)	(0.28)
D_{2020}	0.39	-0.23	0.09
	(0.54)	(0.33)	(0.39)
D_{2021}	0.62**	-0.32*	-0.36*
	(0.25)	(0.18)	(0.20)
D_{2022}	0.02	-0.45	-0.05
	(0.52)	(0.40)	(0.47)
D_{2023}	-0.02	-1.16***	-0.95***
	(0.45)	(0.38)	(0.29)
D_{2024}	0.03	-0.73*	-1.07***
	(0.62)	(0.37)	(0.33)
$\log(ML\ OP/s)$	0.65***	0.27**	0.29**
	(0.12)	(0.11)	(0.12)
$\log(Memory\ Size)$		1.21***	1.55***
		(0.17)	(0.20)
$\log(TDP)$			-0.60*
			(0.32)
Adj. R^2	0.32	0.70	0.73
N	56	55	53

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

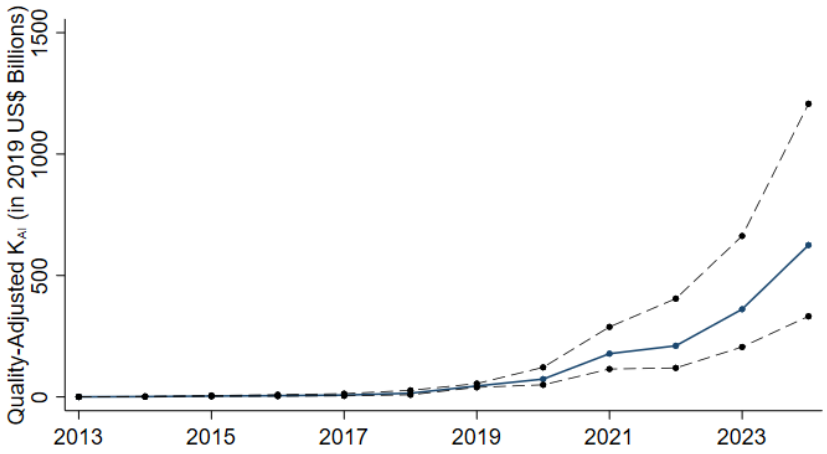
B Additional Figures

Figure A.1: Evolution of AI Quality-Adjusted Price Index



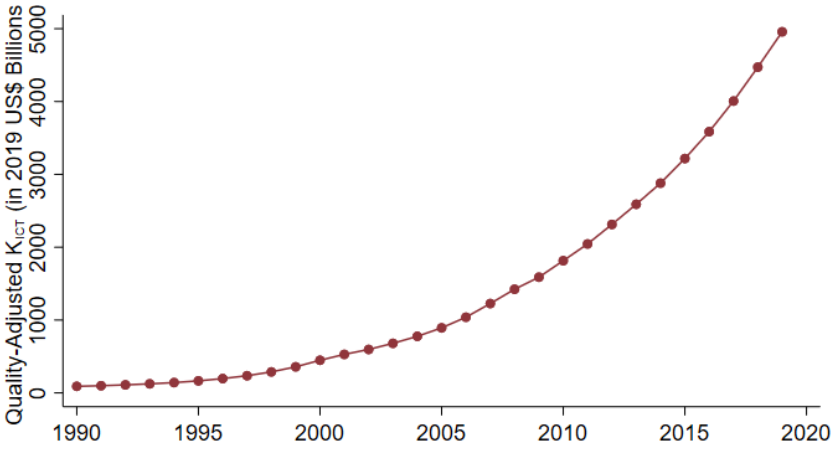
Note: This figure presents the quality-adjusted Price for AI using the hedonic regression presented in equation (1) and AI Epoch data for 2013-2024. Dashed lines present the 95% confidence interval. The blue line indicates $P_{AI} = 1$. The index is normalized to 1 in 2019.

Figure A.2: Evolution of AI Quality-Adjusted Capital Stock



Note: This figure presents the quality-adjusted Capital Stock for AI using the perpetual inventory method. Quality adjustment is based on the constructed hedonic AI Price index. Dashed lines present the 95% confidence interval, based on the bounds for the 95% confidence interval for the hedonic AI Price index. The Price index is normalized to 1 in 2019.

Figure A.3: Evolution of ICT Quality-Adjusted Capital Stock



Note: This figure presents the quality-adjusted ICT Capital For details in the data construction, please see Ohanian et al. (2023). The Price index is normalized to 1 in 2019.

C Additional Results Figures

C.1 Individual Experiments: Level Figures

Figure A.4: Transition dynamics (levels): Output, Consumption, and Aggregate Capital

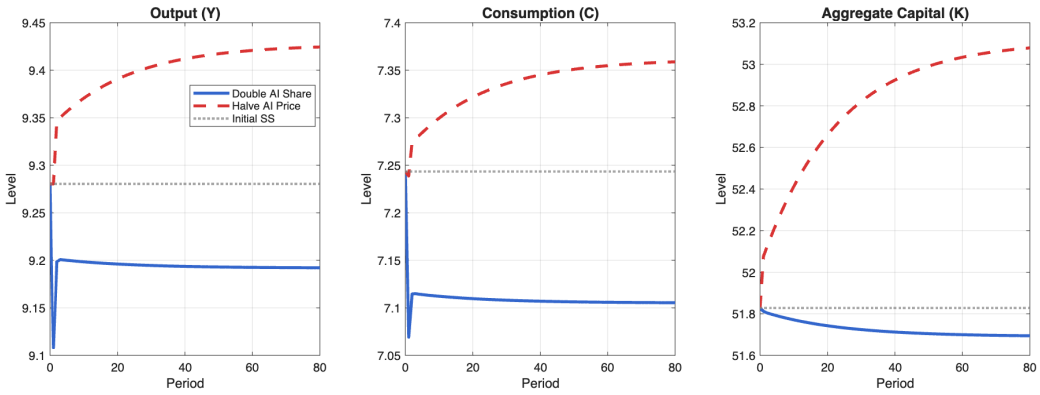


Figure A.5: Transition dynamics (levels): Capital Stocks by Type

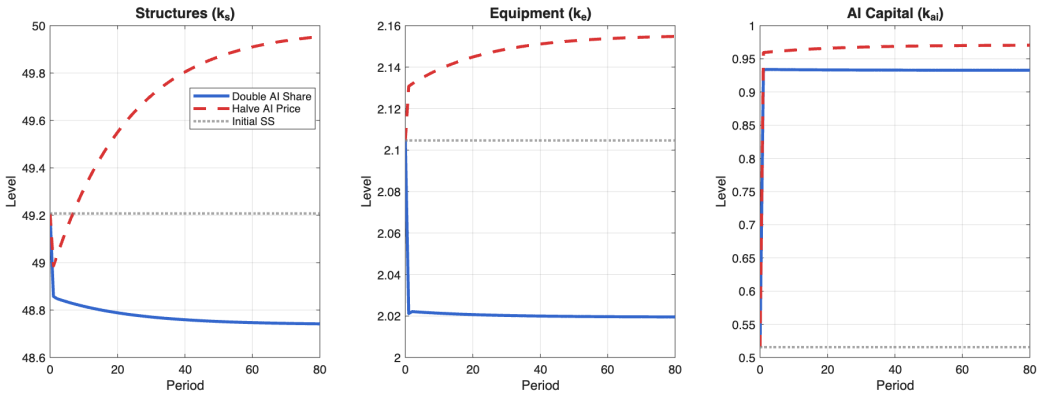


Figure A.6: Transition dynamics (levels): Consumption by Agent Type

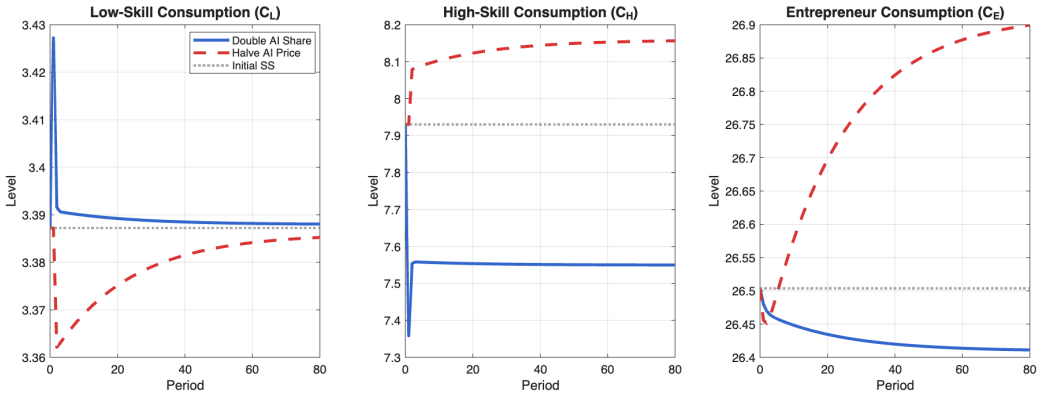
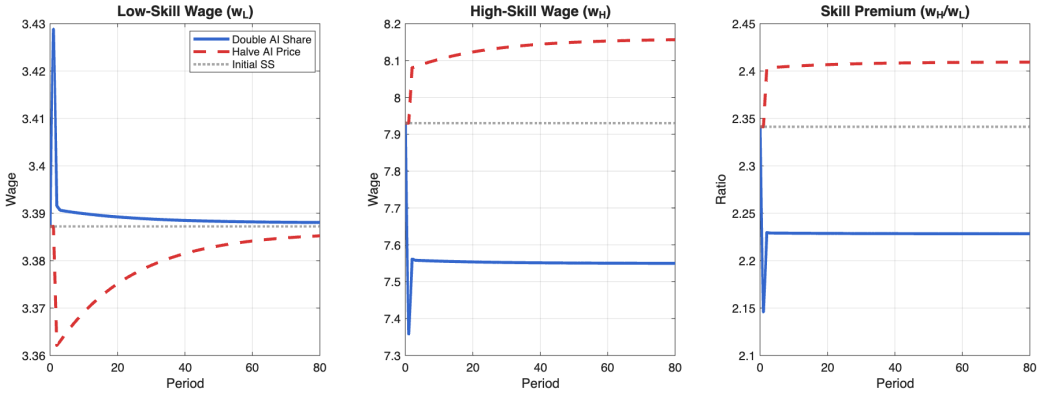


Figure A.7: Transition dynamics (levels): Wages



C.2 Combined Experiment: Additional Figures

Figure A.8: Combined experiment: Capital Stocks

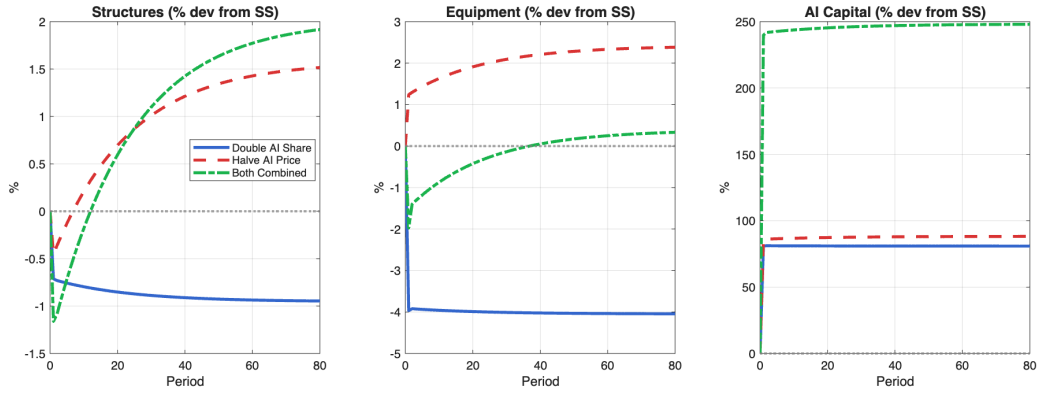


Figure A.9: Combined experiment: Wages

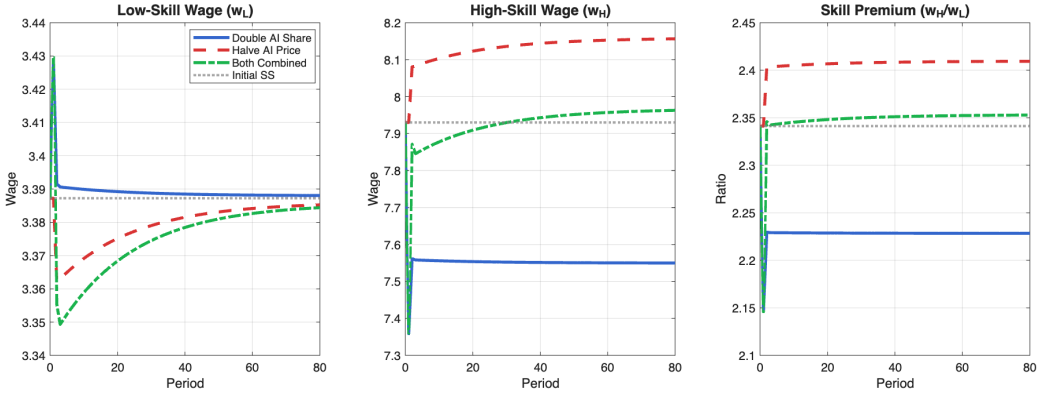


Figure A.10: Combined experiment: Consumption by Agent Type (levels)

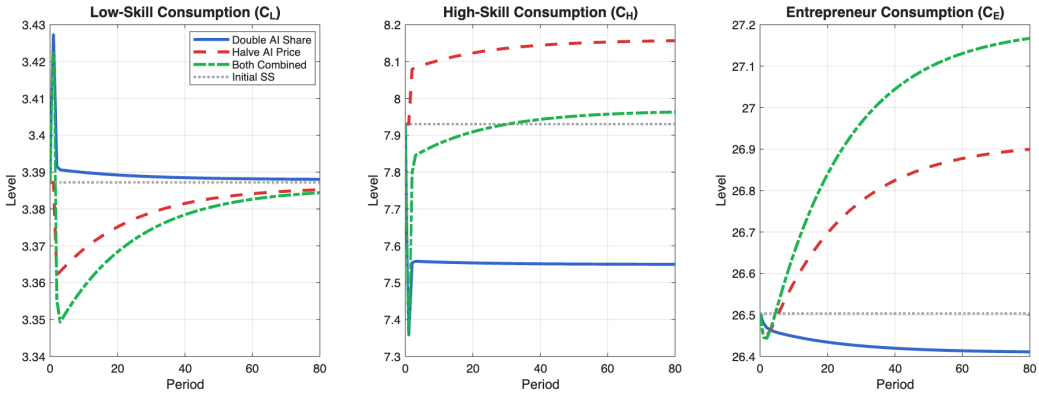
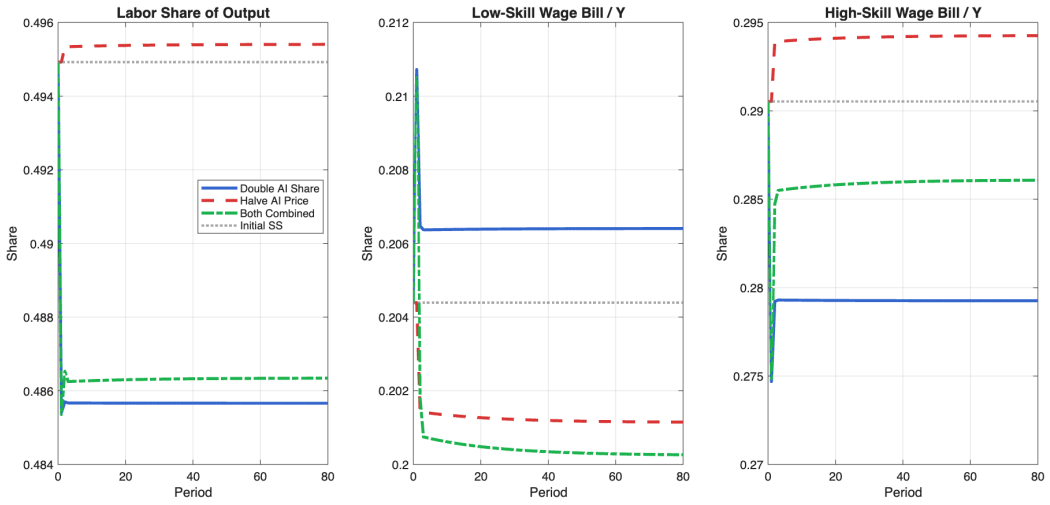


Figure A.11: Combined experiment: Income Distribution



C.3 Robustness: Additional Figures

Figure A.12: Robustness: Output across σ_λ

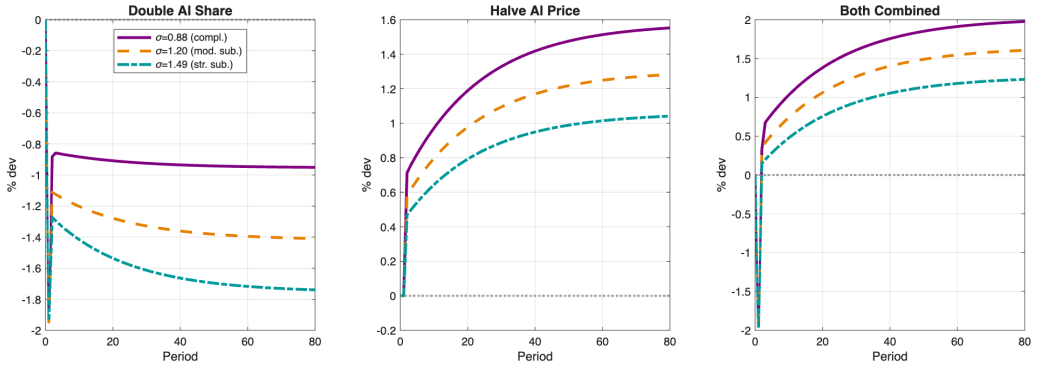


Figure A.13: Robustness: Entrepreneur Consumption across σ_λ

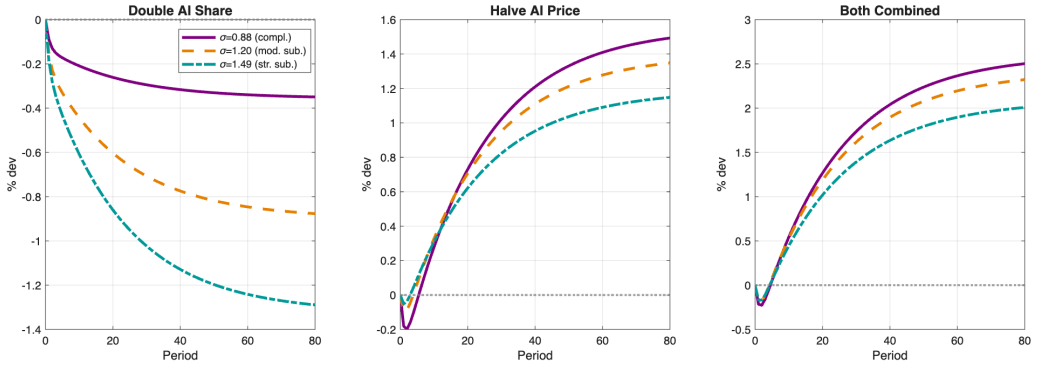


Figure A.14: Robustness: Skill Premium across σ_λ

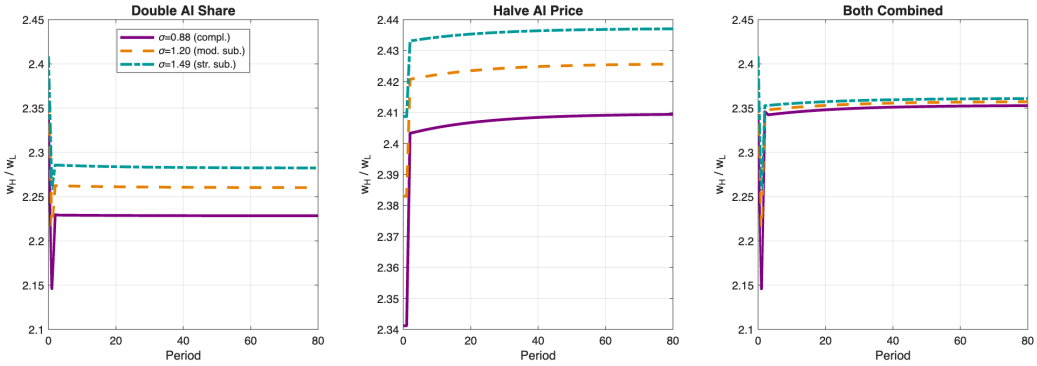
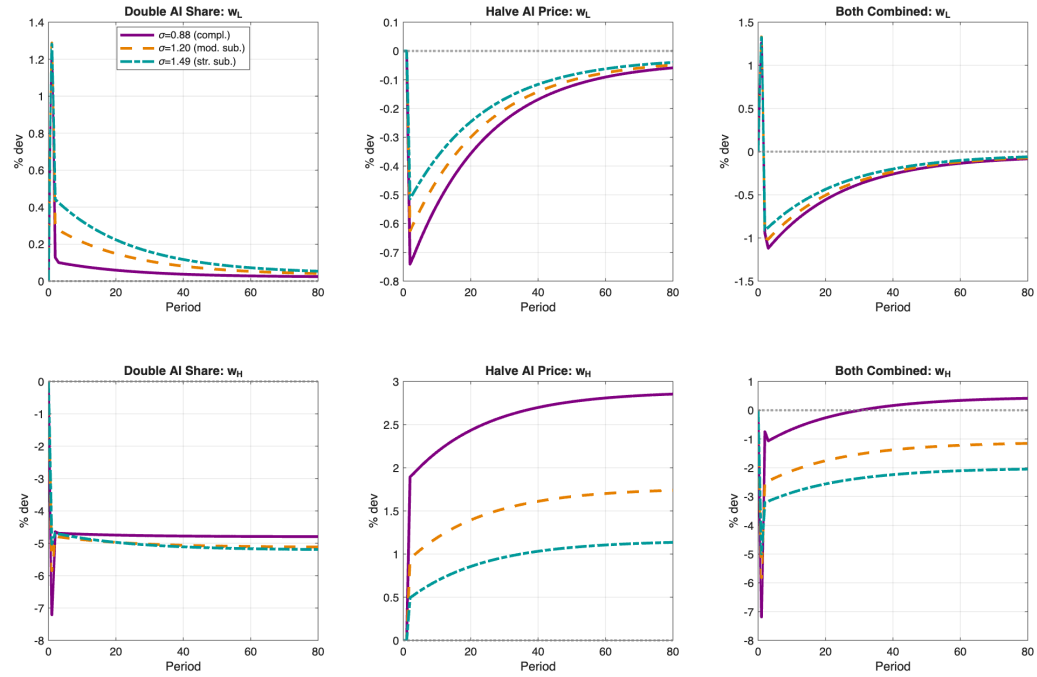


Figure A.15: Robustness: Wages across σ_λ



D Robustness: Varying δ_{AI}

Because the depreciation rate of AI capital is subject to considerable uncertainty, we assess the sensitivity of our results to δ_{AI}. Table A.2 reports the steady-state effects for three values: $\delta_{AI} = 0.10$ (low depreciation), $\delta_{AI} = 0.18$ (baseline, from the BEA average for Computers & Peripherals), and $\delta_{AI} = 0.25$ (high depreciation).

Table A.2: Robustness: Results across δ_{AI}

	$\delta_{AI} = 0.10$ (low)	$\delta_{AI} = 0.18$ (baseline)	$\delta_{AI} = 0.25$ (high)
Output Y (% Δ)			
Double Share	−0.05	−0.95	−1.53
Halve Price	+1.50	+1.58	+1.63
Both	+2.82	+2.02	+1.53
Skill Premium w_H/w_L (level)			
Initial SS	2.39	2.34	2.31
Both	2.43	2.35	2.30
Super-additivity (pp)			
Both − (Exp1+Exp2)	+1.36	+1.40	+1.42

Higher depreciation amplifies the contractionary effect of the share increase—when AI capital depreciates faster, the economy must devote more resources to replacement investment, making the weight mismatch costlier. Conversely, higher depreciation slightly increases the expansionary effect of the price reduction, since the user cost $r_{AI} = (1/\beta - 1 + \delta_{AI})/\zeta$ drops more when both δ_{AI} is high and ζ doubles. Super-additivity remains stable across all three values, confirming the robustness of the interaction mechanism.

Figures A.16 and A.17 display the super-additivity pattern and the output transition dynamics.

Figure A.16: Super-additivity across δ_{AI}

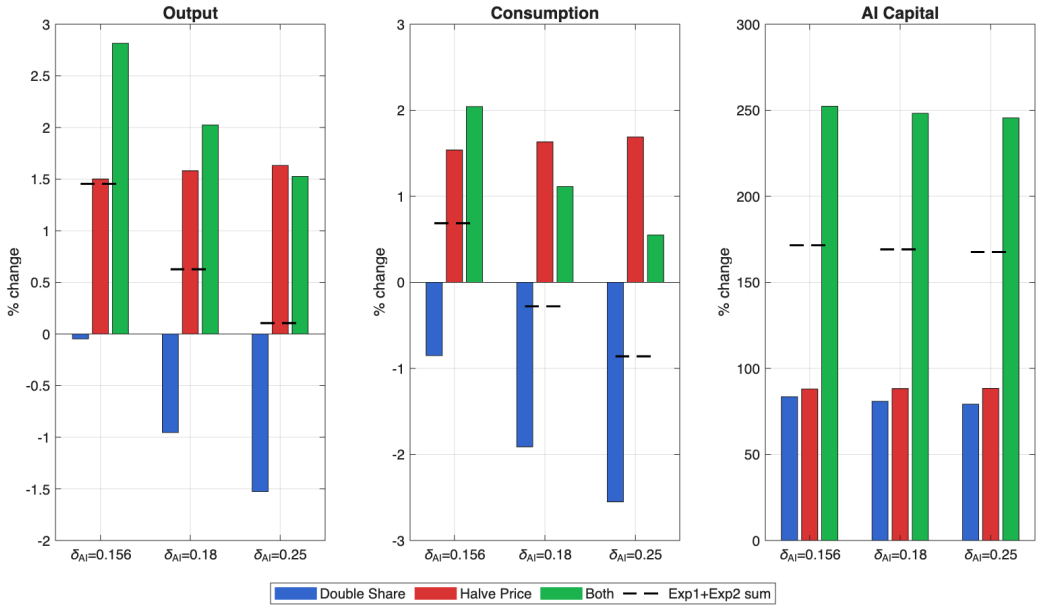


Figure A.17: Robustness: Output across δ_{AI}

