AI-Augmented Capital-Skill Complementarity*

André Victor D. Luduvice†
Federal Reserve Bank of Cleveland

Roberto Pinheiro‡
Federal Reserve Bank of Cleveland

August 2026

Abstract

We model AI as a distinct capital input in a nested CES aggregate production function that extends Krusell et al. (2000) and embed it in a dynamic general equilibrium economy. By taking the model to the data, we estimate the degree to which artificial intelligence substitutes for or complements the other production inputs at the aggregate level. We use the model to study how this substitution pattern shapes the macroeconomic and distributional effects of AI adoption. We find that AI is complementary to the high-skill–equipment composite and that the AI weight in production remains small. The estimated model implies different effects of alternative AI shocks. A rise in the AI usage share is contractionary because it increases reliance on a scarce complementary input. A fall in AI prices is expansionary because it lowers the cost of accumulating AI capital. A tax on AI capital income raises limited revenue while the AI capital stock remains small, but can finance welfare-improving transfers as the AI price falls. A large-scale universal basic income (UBI) funded jointly by a consumption tax preserves AI investment with welfare gains for low-skill workers at the expense of losses for high-skill workers and entrepreneurs. On the measurement side, we construct a quality-adjusted AI price index from hedonic regressions and build a corresponding AI capital stock.

Keywords: AI, technological change, capital-skill complementarity, earnings inequality.
JEL Codes: E22, E25, J24, J31.

*We would like to thank Yvan Becard, Daniel Carroll, Sebastian Dyrda, Victor Hernandez Martinez, Barış Kaymak, Karen Kopecky, as well as seminar participants at the SAET 2026, the Cleveland Fed Brown Bag, the LAEF/UCSB conference "AI and the Macroeconomy," EESP/FGV-SP, York University, and the University of Toronto. We thank Riccardo DiCecio for generously sharing data on detailed equipment prices. We also thank Youha Kim and Rory Quinlan for outstanding research assistance. We acknowledge substantial use of generative AI and large language models in the development of this paper. All errors are our own. 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. This version: August 2026. First preliminary version: January 2026.

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

1 Introduction

The development and initial adoption of AI tools generated considerable uncertainty about the impact of AI on the productive process. Opinions are divided in terms of the degree to which AI may substitute for or complement skilled and unskilled labor. At one end of the spectrum, Anthropic CEO Dario Amodei believes in strong substitutability between AI and white-collar labor. According to Amodei, AI could wipe out half of all entry-level white-collar jobs – and spike unemployment to 10–20 percent in the next one to five years.[1] At the other end of the spectrum, Nvidia CEO Jensen Huang expresses a strong belief in the complementarity between labor and AI, stating that there will be more jobs, but every job will be augmented by AI.[2] The current empirical literature is mixed. 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. In contrast, Aghion et al. (2025) find that AI adoption is positively associated with an increase in total firm-level employment, even in occupations classified as likely to be displaced by AI. Overall, existing studies remain concentrated on particular industries, tasks, and occupations.

Our goal in this paper is to build and estimate a macroeconomic model that incorporates AI as a new form of capital that firms accumulate and deploy, embedded in a neoclassical economy with different agents. Building a macroeconomic model allows us to account for the interactions of AI with the different economic inputs in the production process, as well as AI's impact on economic aggregates. By taking the model to the data, we obtain estimates of the elasticity of substitution across the different production factors at the economy-wide level instead of particular markets. As a result, we can infer the degree of complementarity or substitutability between high-skill labor and AI capital, shedding light on the current debate. Moreover, the model allows us to evaluate the impact on macroeconomic aggregates of policies designed to partially reverse some of the negative effects that AI may bring. For example, we can study the impact of a digital dividend suggested by Marinescu (2025), which introduces a universal basic income mechanism funded by a tax on the digital sector.

Our modeling approach follows the classic macroeconomic production function framework, used in the past to understand the impact of information and communication technologies (ICT) on the macroeconomy. In particular, 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. 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 dynamic general equilibrium environment with a government sector featuring progressive labor taxation, a consumption tax, a capital income tax, and lump-sum transfers.

[1]See the Axios article on the interview.
[2]See the Axios article on Huang's interview.

High-skill workers hold a portfolio share of capital income, providing a bridge from aggregate adjustments to distributional outcomes and allowing for the study of fiscal policy counterfactuals.

Credible estimates, however, require a contribution in measurement: we propose to measure AI as capital and construct 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 empirical side, our hedonic regressions show that quality-adjusted AI hardware prices fell by 97 percent between 2013 and 2024, roughly 27 percent 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 grew from US$ 15 billion in 2018 to US$ 625 billion in 2024 (in 2019 constant dollars), a 4,000 percent increase in six years. Despite this growth, AI capital remains small relative to the total capital stock, currently less than 5 percent of structures.

We then write and solve a general version of the model, characterizing its equilibrium as well as its steady state. We estimate the production function parameters via GMM, targeting moments that include the wage-bill ratio, the labor share, no-arbitrage and user-cost conditions across capital types, and the capital-demand conditions. Our key finding is that AI and the high-skill–equipment composite are complements, with an estimated elasticity of substitution $\sigma_\lambda \approx 0.89$. Hence, the currently available data lend little support to a broad substitution of AI for high-skill labor, in line with the results of Aghion et al. (2025) and consistent with the view that AI adoption can complement white-collar workers. Our estimates also indicate that the AI weight in production is still small (approximately 2.4 percent), reflecting the fact that AI capital remains a small fraction of the total capital stock.

With the estimated model in hand, we study 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 doubling of the AI usage share can be interpreted as either an increase in the measure of tasks performed by AI (as presented by Korinek and Suh (2024) and Acemoglu et al. (2025)) or as a thought experiment in which CEOs promote broader adoption of AI. In contrast, the decline in the AI price is an investment-specific technological change in the spirit of Greenwood et al. (1997) and Krusell et al. (2000).

The two shocks increase AI capital in a similar magnitude, but pull output in opposite directions. Doubling the AI weight is contractionary: the economy is told to lean on an input it barely holds, and output settles about 1.9 percent below its initial level. Halving the AI price is expansionary, raising output by about 1 percent as cheaper AI capital accumulates. Together the two shocks leave output essentially unchanged, whereas adding the separate effects would predict a clear decline. The interaction accounts for the entire gap, because each shock relaxes the constraint that limits the other: the lower price supplies the AI capital that the higher weight calls for, and the higher weight raises what each unit of that capital contributes. Reinforcement is not dominance, however, and we find that the combined outcome still falls short of what the price decline delivers on its own.

Finally, we consider policy counterfactuals that introduce a separate tax on AI capital income and rebate the proceeds as per-capita transfers, tracing its Laffer curve following the standard in Trabandt and Uhlig (2011). This policy is similar to the digital dividend suggested by Marinescu (2025). Since AI is a small share of the capital stock, the AI tax raises only modest revenue on its own. However, when the AI tax is paired with a fall in the AI price – which enlarges the AI capital base – it becomes welfare-positive for low-skill workers and in aggregate, at a small cost to high-skill workers and entrepreneurs. We also conduct a policy counterfactual that pairs this logic with the possibility of financing a Universal Basic Income (UBI): a per-capita transfer of about \$12,000 per year – roughly a fifth of average worker income – funded by combining the AI tax with a broad consumption tax. This experiment connects our analysis to the quantitative literature on UBI (Luduvice, 2024; Daruich and Fernández, 2024; Conesa et al., 2023; Guner et al., 2023), with the distinctive feature that a growing AI capital base shoulders part of a financing burden that would otherwise fall entirely on broad-based taxation. We find that such a policy would halt the AI capital accumulation the price decline would otherwise deliver, contract output, and generate welfare increases for low-skill workers at the expense of costs for high-skill workers and entrepreneurs.

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 dynamic general equilibrium framework by allowing AI capital to interact differently with skilled and unskilled labor, incorporating a government sector with progressive taxation and redistribution. Section 4 discusses the calibration and estimation strategy. Section 5 reports the quantitative exercises – transitory impulse responses and the permanent AI shocks and their super-additivity – with robustness checks gathered in the appendix. Section 6 then studies policy counterfactuals that tax AI capital income to fund transfers, together with the accompanying welfare analysis. The final section concludes.

Related Literature. Our paper builds on different strands of related 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. Quasi-experimental evidence supports this production structure: Hernandez Martinez (2025) exploits the US shale boom to provide causal evidence of capital-skill complementarity in manufacturing, with firms substituting capital for low-skill labor, consistent with capital being more substitutable with low-skill than with high-skill labor. 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. Related work on CES estimation and the role of IT capital includes Beaudry et al. (2010), Eeckhout et al. (2026), Jerbashian (2026), and Graetz and Michaels (2018). We add to this literature by augmenting the production function to include AI capital and estimating the elasticity of substitution between AI capital and the equipment–high-skill composite. We also add by incorporating the estimated production function in a fully-specified general equilibrium model.

Second, our paper is part of the burgeoning literature on the macroeconomics of AI. The tradition of modeling automation as capital that can take over productive tasks from labor goes back to Zeira (1998). Aghion et al. (2018) examine the growth implications of AI within this tradition, and Korinek and Suh (2024) develop scenarios for economies undergoing an AI transition. 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; Acemoglu and Restrepo, 2018). Empirical studies document heterogeneous AI exposure across occupations and tasks (Webb, 2020; Brynjolfsson et al., 2018; Hampole et al., 2025). Autor (2024) argues that AI could extend expert tasks to a broader set of workers, while de Souza (2026) provides firm-level evidence on AI adoption patterns and its impact on administrative and production workers' employment and wages. We 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; Eden and Gaggl, 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. Overall, we add to this literature by merging the capital-skill complementarity time series analysis and production function estimation with a neoclassical macroeconomic model with heterogeneous agents and a detailed public finance structure that allows for mechanism and policy counterfactuals.

Third, our policy experiments are related to the literature on the optimal taxation of automation capital: Guerreiro et al. (2022) show that it can be optimal to tax robots while routine workers remain in the labor force, Thuemmel (2023) shows that the optimal robot tax can go either way once occupational adjustments are accounted for, and Costinot and Werning (2023) derive sufficient statistics for optimal technology regulation. Our analysis is positive rather than normative – we trace the revenue and welfare consequences of a given AI tax-and-transfer scheme rather than solving for the optimal instrument – and this literature provides a natural benchmark for interpreting the results. 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 and add to this literature by introducing AI as an explicit capital category – allowing, in the general framework, different elasticities of substitution with each labor type, with our quantitative specification restricting AI to the high-skill bundle – and an AI capital income tax that can be paired with other tax policy instruments.

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 using the lessons to compute a novel time series of the quality-adjusted price of AI.

Overall, we integrate these strands by extending capital-skill complementarity to an explicit AI capital input, embedding it in a dynamic general equilibrium 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 shape our empirical strategy. First, there is the distinction between usage-based measures and installed-capacity measures. Much of the macro literature treats technology-specific capital as an input. 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 Bureau of Economic Analysis

(BEA) measures.[3] 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).

2.2 Epoch AI Data

Our main AI-hardware database comes from Epoch AI (Ai, 2025), which tracks the accelerators – graphics and tensor processing units – used to train and deploy machine learning models. Only a subset of them are commercial products with observable release prices, so the sample available for hedonic estimation is much smaller than the universe of listed chips. Appendix A.2 reports the sample counts and the auxiliary price sources.

Beyond prices, the database records technical characteristics that let us control for quality. We use three: *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*), the maximum power that can be dissipated as heat. These capture performance, capacity constraints, and energy requirements, the dimensions Baily et al. (2025) emphasize in discussing the measurement of generative AI. They are also the characteristics with the fewest missing observations, which keeps the estimation sample as large as the data allow.

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

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

where $p_{i,t}$ denotes the release price, $y_{j,i,t}$ denotes characteristic j (of m in total) for processor i at time t with implicit price β_j, D_s are year dummies for the N release years with coefficients δ_s, and

[3] For some ICT components, alternative sources based on micro data (such as the Aberdeen Ci Technology database) can provide additional granularity and cross-checks.

$u_{i,t}$ is the regression residual. We focus on newly released processors, so that the regression traces the technology frontier, as recommended by Fisher et al. (1985). This approach 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, with statistically insignificant year dummies set to zero (Appendix A.2).[4] Figure 1 reports the resulting series, which summarizes the evolution of the quality-adjusted price of AI hardware over time. The index equals 1 in the base year, 2019. Appendix Figure A.1 includes the confidence intervals, which are wide early on but narrow in recent years.

Figure 1: Evolution of AI Quality-Adjusted Price Index

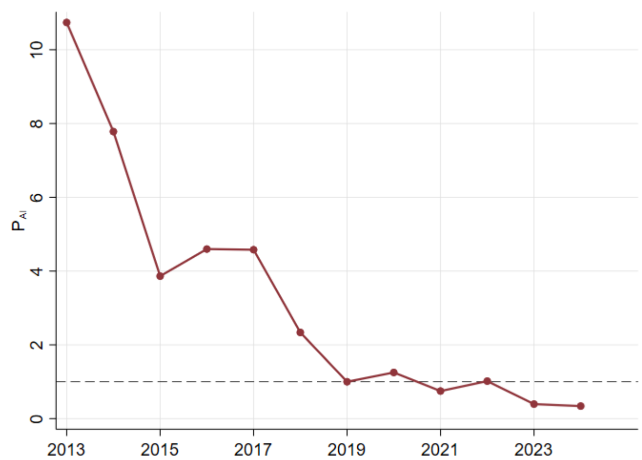

Notes: Quality-adjusted AI price index from the hedonic dummy-variable method, normalized to 1 in 2019. The index falls from 10.7 in 2013 to 0.34 in 2024, a 97% decline over the sample. Confidence intervals are reported in Appendix Figure A.1.

Figure 1 shows a sharp decline in quality-adjusted prices.[5] Considering the normalization of 1 in 2019, the index declines from 10.7 in 2013 to 0.34 in 2024, a decline of 97 percent in 11 years, or 27 percent per year. If we focus on the period with the largest investment in AI (2018–2024), the decrease in quality-adjusted prices was 85 percent in six years, or 27.5 percent per year.

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

[4]Setting insignificant year effects to zero is a pretest device that mainly smooths the intervening path of the index. The year effects at the endpoints of the sample are statistically significant (Appendix Table A.1), so the headline 2013–2024 decline does not hinge on this rule.

[5]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.

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.[6] For the PIM data construction, we set $\delta_{AI} = 0.156$, following Kudoh and Miyamoto (2025).[7] Figure 2 presents the resulting time series for the quality-adjusted AI capital stock. Appendix Figure A.2 includes the confidence intervals.

Figure 2: Evolution of AI Quality-Adjusted Capital Stock

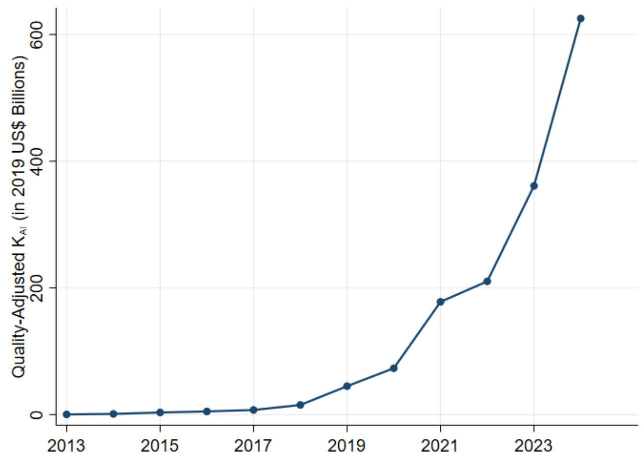

Notes: Quality-adjusted AI capital stock constructed by the perpetual inventory method, with $K_{AI} = 0$ in 2012, using Stanford HAI's 2025 AI Index investment data deflated by the hedonic AI price index. Confidence intervals are reported in Appendix Figure A.2.

Figure 2 shows a fast accumulation of capital, in particular after 2018. Quality-adjusted K_{AI} increased from US$ 15 billion in 2018 to US$ 625 billion in 2024 in 2019 constant dollars (a 4,000 percent increase in six years, representing an 85 percent annual increase). Furthermore, given that the gap between the point estimate and the upper limit of the 95 percent 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 percent of the National Income and Product Accounts' (NIPA) capital stock of structures. The AI boom resembles the ICT boom in that one respect – the stock stays small relative to total capital – but not in growth speed, as the comparison below makes clear.

[6]Stanford HAI's investment data is based on the Quid database. The Quid investment data captures four categories of capital flowing into AI companies, namely mergers and acquisitions, minority stake investments, private investment, and public offerings. Stanford HAI's private investment data focuses more narrowly on private financing events, such as venture capital or private equity funding, directed at AI companies that have received over $1.5 million in funding since 2013.

[7]The model calibration (Section 4) uses the same rate. In the estimation, where AI capital is proxied by the broader Computers & Peripherals category, the AI no-arbitrage condition instead carries that category's own BEA depreciation series.

Comparison: ICT revolution. To put these findings in perspective, we benchmark AI capital against the ICT capital stock that sparked the literature following Krusell et al. (2000). Figure 3 below presents the evolution of the ICT quality-adjusted capital stock (panel 3a) and price index (panel 3b) from 1990 to 2001 and compares it against the growth in AI counterparts from 2013 to 2024.[8] Our ICT data comes from Ohanian et al. (2023).[9]

Figure 3: The ICT vs. AI Revolutions: Capital Stock and Prices

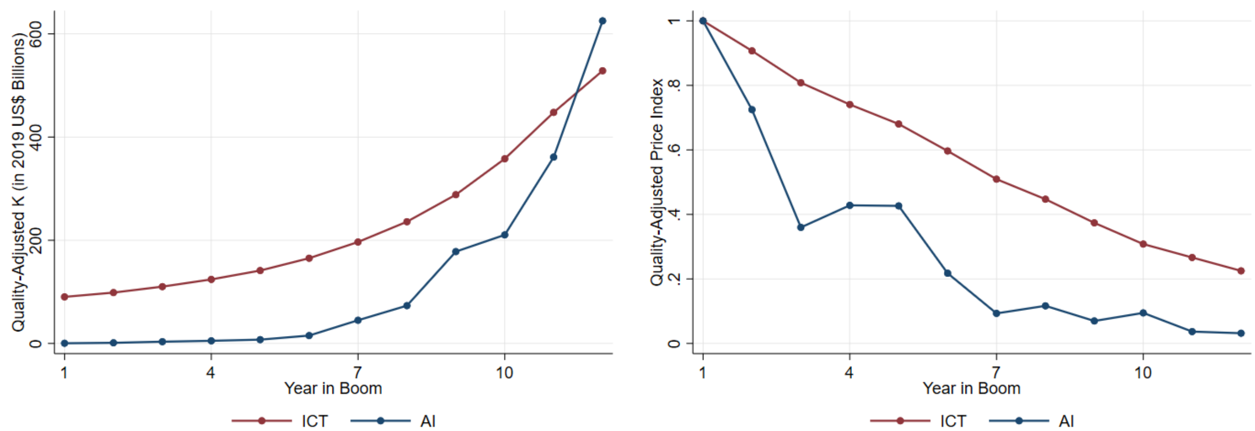

(a) Quality-Adjusted Capital Stock

(b) Quality-Adjusted Price Index

Notes: This figure presents changes in the AI and ICT quality-adjusted capital stocks and price indexes. We consider the boom years 1990–2001 for ICT and 2013–2024 for AI. We normalize the price indexes to equal 1 in the first year of each boom.

Panel 3a shows that ICT growth was significantly slower than the recent investment in AI. Between 1990 and 2001, ICT capital grew from US$ 90 billion to US$ 529 billion in 2019 constant dollars, representing growth of 488 percent in 11 years, or roughly 17 percent per year. The comparison is indicative rather than like-for-like: the AI stock is accumulated from a near-zero base immediately before its window, while the 1990 ICT stock was already mature, and percentage growth rates are sensitive to that difference in starting points.

In terms of decline in quality-adjusted ICT prices, as shown in panel 3b, 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 0.23. In contrast, after normalizing the initial AI price index to 1 in 2013, the AI price index had declined to 0.032 in 2024.

[8] 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.

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

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 the future at rate $\beta \in (0, 1)$. The economy is populated by three types of agents – low-skill workers, high-skill workers, and entrepreneurs – together with a government. 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 the two labor types, 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: entrepreneurs own the capital stock and make the investment decisions, while a share θ_H of after-tax ownership income is paid out to high-skill workers. Low-skill workers earn only labor income. Both types of workers are hand-to-mouth, i.e., they do not have access to a savings technology. The government levies progressive labor income taxes, a consumption tax, and a capital income tax, using the proceeds to finance a per-capita transfer and government purchases.

3.1 Production

There are four final goods: consumption c_t, structures investment x_{st}, equipment investment x_{et}, and AI investment x_{AIt}.[10] Consumption and structures are produced with a decreasing 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:

$$y_t = c_t + x_{st} + \frac{x_{et}}{q_t} + \frac{x_{AIt}}{\zeta_t} + G_t = A_t G(k_{s,t}, k_{e,t}, k_{AI,t}, m_{L,t}, m_{H,t}) \tag{2}$$

where $G(k_{s,t}, k_{e,t}, k_{AI,t}, m_{L,t}, m_{H,t})$ is given by:

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

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.

[10] Investments are expressed in units of the relative capital.

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}), $j \in \{L, H\}$. 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.

We consider that worker types are fixed, i.e., a worker is born with a type $j \in \{L, H\}$. Following Krusell et al. (2000), the skill classification corresponds to educational attainment: high-skill workers are those with a college degree or above, while low-skill workers are those without. 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 $M_L + M_H + M_E = \overline{M}$. We normalize $\overline{M} = 1$.

3.2 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, the firm's problem is given by:

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

where k_s, k_e, and k_{AI} are capital structures, equipment, and AI. 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, respectively. $w_{i,t}$ is the wage rate for skill $i \in \{H, L\}$.

Under decreasing returns to scale ($\alpha_1 + \alpha_2 < 1$), total factor payments do not exhaust output, leaving aggregate profits $\pi_t = y_t - r_{s,t}k_{s,t} - r_{e,t}k_{e,t} - r_{AI,t}k_{AI,t} - w_{L,t}m_{L,t} - w_{H,t}m_{H,t} > 0$ that accrue to entrepreneurs. Ownership income aggregates rental income and profits: $\Omega_t \equiv r_{s,t}k_{s,t} + r_{e,t}k_{e,t} + r_{AI,t}k_{AI,t} + \pi_t$. Throughout, the capital stocks $k_{i,t}$ and ownership income Ω_t are economy-wide (aggregate) quantities.

3.3 Entrepreneur's Problem

Entrepreneurs own all three types of capital (structures, equipment, and AI).[11] As capital owners, entrepreneurs make the investment decisions that determine the economy's capital composition, and therefore how the gains from AI accumulation are distributed. The representative entrepreneur maximizes:[12]

$$\max_{\{C_{E,t},k_{s,t+1},k_{e,t+1},k_{AI,t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \ln C_{E,t} \tag{4}$$

subject to:

$$\begin{aligned} M_E(1+\tau_c)\, C_{E,t} + k_{s,t+1} + P_{e,t}\, k_{e,t+1} + P_{AI,t}\, k_{AI,t+1} = k_{s,t} + P_{e,t}\, k_{e,t} \\ + P_{AI,t}\, k_{AI,t} + (1-\tau_k)\, \Omega_t^{\text{net}} - (1-\tau_k)\, \theta_H\, \overline{\Omega}_t^{\text{net}} + M_E\, T_{pc,t} \end{aligned}$$

where $P_{e,t} = 1/q_t$ and $P_{AI,t} = 1/\zeta_t$ are the equipment and AI prices, Ω_t is aggregate ownership income, Ω_t^{net} is the same measure net of economic depreciation (defined in equation (5) below), and $T_{pc,t}$ is a per-capita lump-sum transfer from the government.

The parameter $\theta_H \in [0,1]$ is the share of after-tax ownership income accruing to high-skill workers, which reflects the empirical observation that some capital income reaches workers through retirement accounts and equity compensation. That payment is indexed to economy-wide average net ownership income $\overline{\Omega}_t^{\text{net}}$, which the individual entrepreneur takes as given. In equilibrium, $\overline{\Omega}_t^{\text{net}} = \Omega_t^{\text{net}}$, so entrepreneurs retain a fraction $1-\theta_H$ of after-tax net ownership income.

Following Trabandt and Uhlig (2011) and Carroll et al. (2024), the capital income tax applies to ownership income *net of economic depreciation*:

$$\Omega_t^{\text{net}} = \Omega_t - D_t, \qquad D_t = \delta_s k_{s,t} + \frac{\delta_e}{q_t} k_{e,t} + \frac{\delta_{AI}}{\zeta_t} k_{AI,t}, \tag{5}$$

where D_t is depreciation measured in consumption-good units. Because the depreciation allowance is shielded from taxation, the entrepreneur retains it to fund replacement investment, so the budget above carries capital forward at full value and is equivalently stated in terms of *net* investment $k_{i,t+1} - k_{i,t}$.

[11] This ownership structure captures the empirical concentration of capital income: in the 2022 Survey of Consumer Finances, the top 10 percent of households by income receive roughly 89 percent of capital income from interest, dividends, and capital gains.

[12] Entrepreneurs form a continuum of mass M_E holding equal shares of the economy's capital. We state the problem in aggregate units for capital and ownership income, while $C_{E,t}$ denotes consumption *per entrepreneur*. The resource constraint (Appendix B.1) accordingly uses $M_E C_E$ for aggregate entrepreneur consumption. Each entrepreneur is atomistic and takes the economy-wide payment $\theta_H\, \overline{\Omega}_t^{\text{net}}$ as given.

The stocks of capital evolve as:

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

Substituting the laws of motion and prices, the budget constraint becomes:

$$M_E(1+\tau_c)\, C_{E,t} + (k_{s,t+1} - k_{s,t}) + \frac{k_{e,t+1} - k_{e,t}}{q_t} + \frac{k_{AI,t+1} - k_{AI,t}}{\zeta_t} = (1-\tau_k)\, \Omega_t^{\text{net}} - (1-\tau_k)\, \theta_H\, \overline{\Omega}_t^{\text{net}} + M_E\, T_{pc,t}$$

The first-order conditions, with $\Lambda_{E,t} = 1/\left((1+\tau_c)\, C_{E,t}\right)$, yield the Euler equations. The share θ_H does not appear in them because the payment to high-skill workers is indexed to aggregate net ownership income, which the individual entrepreneur takes as given: it shifts the level of resources without changing the marginal return to an additional unit of capital.[13] Note also that $(1+\tau_c)$ cancels between $\Lambda_{E,t}$ and $\Lambda_{E,t+1}$ under log utility, so the consumption tax does not distort the intertemporal margin when τ_c is constant. When τ_c varies over time, the ratio $(1+\tau_{c,t})/(1+\tau_{c,t+1})$ enters the Euler equations.

3.4 Workers' Problems

3.4.1 Low-skill Worker

We assume that the only source of income for a low-skill worker is after-tax labor income and a per-capita government transfer. We assume that low-skill workers do not have a savings technology. As a result, the low-skill worker problem is given by:

$$\max_{\{C_{L,t}, L_{L,t}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \left[\ln C_{L,t} - B \frac{L_{L,t}^{1+\frac{1}{\vartheta}}}{1+\frac{1}{\vartheta}} \right] \tag{6}$$

subject to:

$$(1+\tau_c)\, C_{L,t} = (1-\tau_L)\, w_{L,t} L_{L,t} + T_{pc,t}$$

where β is the discount rate, $C_{L,t}$ is consumption, $L_{L,t}$ is the individual's labor supply in period t, $\vartheta > 0$ is the Frisch elasticity of labor supply, $B > 0$ is the disutility-of-labor scaling parameter, τ_c is the consumption tax rate, τ_L is the labor income tax rate on low-skill workers, and $T_{pc,t}$ is a per-capita lump-sum transfer from the government. The parameter B is calibrated so that average

[13] The indexing convention is innocuous for the results. Under the alternative, equity-like convention in which each entrepreneur pays θ_H of *own* net ownership income, every Euler equation carries a constant wedge $(1-\theta_H)$ on the after-tax net return. Because β is calibrated to the capital–output ratio, that wedge is fully absorbed by the recalibrated discount factor, leaving the required pre-tax user costs, and hence the allocations and policy experiments, unchanged.

steady-state labor supply equals one-third of the time endowment. From the first order conditions, we obtain the intratemporal optimality condition:

$$BL_{L,t}^{\frac{1}{\vartheta}}(1+\tau_c)\,C_{L,t} = (1-\tau_L)\,w_{L,t}$$

3.4.2 High-skill Worker

The problem for the high-skill worker is similar, with two differences. First, high-skill workers face a higher marginal tax rate $\tau_H > \tau_L$ on their labor income, capturing the progressivity of the tax code. Second, high-skill workers receive the share θ_H of after-tax ownership income defined in the entrepreneur's problem. The problem is:

$$\max_{\{C_{H,t},L_{H,t}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \left[\ln C_{H,t} - B\frac{L_{H,t}^{1+\frac{1}{\vartheta}}}{1+\frac{1}{\vartheta}}\right] \tag{7}$$

subject to:

$$(1+\tau_c)\,C_{H,t} = (1-\tau_H)\,w_{H,t}L_{H,t} + (1-\tau_k)\,\theta_H\,\frac{\Omega_t^{\text{net}}}{M_H} + T_{pc,t}$$

where all variables are as previously defined. From the first order conditions, we obtain the intratemporal optimality condition:

$$BL_{H,t}^{\frac{1}{\vartheta}}(1+\tau_c)\,C_{H,t} = (1-\tau_H)\,w_{H,t}$$

3.5 Government

The government levies three types of taxes: a consumption tax at rate τ_c, skill-specific labor income taxes at rates τ_L and τ_H (with $\tau_H > \tau_L$ to capture the progressivity of the U.S. tax code), and a capital income tax at rate τ_k levied on ownership income (rental income and profits) net of economic depreciation. The two-bracket labor tax is a simplified representation of the progressive income tax schedule commonly modeled via the log-linear tax function of Benabou (2002) and Heathcote et al. (2017).[14] Tax revenues finance a per-capita lump-sum transfer $T_{pc,t}$ paid to all $\overline{M}$ agents and government purchases G_t.

The government budget constraint is:

$$\tau_c \sum_{j\in\{L,H,E\}} M_j C_{j,t} + M_L\tau_L w_{L,t}L_{L,t} + M_H\tau_H w_{H,t}L_{H,t} + \tau_k\,\Omega_t^{\text{net}} = T_{pc,t}\cdot\overline{M} + G_t \tag{8}$$

[14]This family of functions is commonly used in quantitative macroeconomic studies of fiscal policy. For recent examples, see Holter et al. (2019); Kindermann and Krueger (2022).

where $\Omega_t^{\text{net}} = \Omega_t - D_t$ is aggregate ownership income net of economic depreciation D_t, defined in equation (5). It is the economy-wide net capital tax base.

In our baseline calibration, the tax rates $(\tau_c, \tau_L, \tau_H, \tau_k)$ and the transfer T_{pc} are set to match U.S. fiscal aggregates, and G clears the government budget residually. In our policy counterfactuals shown later in the text, G is held fixed at its baseline level, a separate rate on AI capital income is introduced with the additional revenue rebated as per-capita transfers, and in the largest experiment the consumption tax also adjusts to fund a fixed transfer target.

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 and taxes as given, entrepreneurs choose consumption and investment taking prices, returns, and taxes as given, competitive firms rent capital and hire labor until marginal products equal factor prices, and the government satisfies its budget constraint. Together with the goods- and labor-market clearing conditions, these restrictions form a system of dynamic equations that determines allocations and prices.

In the baseline version in which workers do not save and entrepreneurs own the productive capital, the equilibrium system consists of 20 equations in 20 unknowns per period. With log utility, the marginal utility multipliers satisfy $\Lambda_{j,t} = 1/\left((1+\tau_c)\, C_{j,t}\right)$ and can be substituted out. The unknowns are:

$$\{Y, C_L, C_H, C_E, L_L, L_H, k_s, k_e, k_{AI}, m_L, m_H, r_s, r_e, r_{AI}, w_L, w_H, \pi, T_{pc}, \Omega, G\}. \tag{9}$$

The 20 conditions are the production function, the five firm first-order conditions, the three entrepreneur Euler equations, the two worker intratemporal conditions, the two worker budget constraints, the two labor-market clearing conditions, the goods-market clearing condition, the definitions of profits and of ownership income, the transfer rule (T_{pc} set as a share of output in the baseline – in the policy counterfactuals the transfer instead adjusts and G is fixed), and the government budget constraint, which determines G residually in the baseline. The entrepreneur's budget constraint is implied by the other conditions through Walras's law. Progressive taxation $(\tau_H > \tau_L)$ implies that low-skill and high-skill workers generally supply different amounts of labor in equilibrium.

In Appendix B.2, we present and fully characterize the steady-state equilibrium conditions.

3.7 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. We also reparameterize the CES weights in share-weight form (μ_H, ξ_H) and express labor in efficiency units $\ell_{jt} = e^{\varphi_j} m_{jt}$, where the efficiency weights φ_j capture skill-specific human capital differences following Krusell et al. (2000). Under these restrictions, the production function becomes:

$$Y_t = A_t k_{s,t}^{\alpha_1} \left\{ \ell_{Lt}^{\phi} + \left[\mu_H \left(\xi_H \ell_{Ht}^{\gamma_H} + (1-\xi_H) k_{et}^{\gamma_H} \right)^{\frac{\lambda_H}{\gamma_H}} + (1-\mu_H) k_{AIt}^{\lambda_H} \right]^{\frac{\phi}{\lambda_H}} \right\}^{\frac{\alpha_2}{\phi}} \tag{10}$$

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, $j \in \{H, L\}$. 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 $(1-\mu_H)$ is the AI weight.

The nesting structure implies three distinct elasticities of substitution that are the focus of our estimation:

- $\sigma_\phi = \frac{1}{1-\phi}$: elasticity between the high-skill bundle 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.

The calibration and all quantitative results in Sections 4 through 6 are based on this simplified production function.

4 Estimation and Calibration

In this section we discuss how we bring the model to the data and our approach to the calibration and estimation of the model parameters. We proceed in three steps. We first estimate the parameters of the simplified production function in the Krusell et al. (2000) tradition by GMM, extending the analysis presented by Ohanian et al. (2023). We then set the externally calibrated parameters from standard targets and data averages computed in our data analysis. Finally, given the estimated production function and external parameters, we internally calibrate the remaining parameters to match the relevant aggregate features of the U.S. economy.

4.1 Estimation of Production Function Parameters

AI proxy: Computers & Peripherals. Because our accelerator-based AI capital series begins only recently, in the estimation we proxy AI capital using a capital stock built from BEA real investment in Computers & Peripherals – the computing hardware that underpins AI workloads, including GPUs, TPUs, and the servers that house them – which is available over a longer sample. Figure 4 shows that this series captures the run-up in AI-related investment. The two AI-specific moments below (the AI no-arbitrage and AI user-cost conditions) enter the GMM objective only from the AI start year onward, which is 2013 in our baseline estimates.

Using the BEA's Computers & Peripherals series not only gives us a longer time series, but also keeps the AI capital series construction consistent with the BEA capital accounts used for structures and equipment.[15]

Figure 4: Real Investment in Computers & Peripherals (US$ Billions)

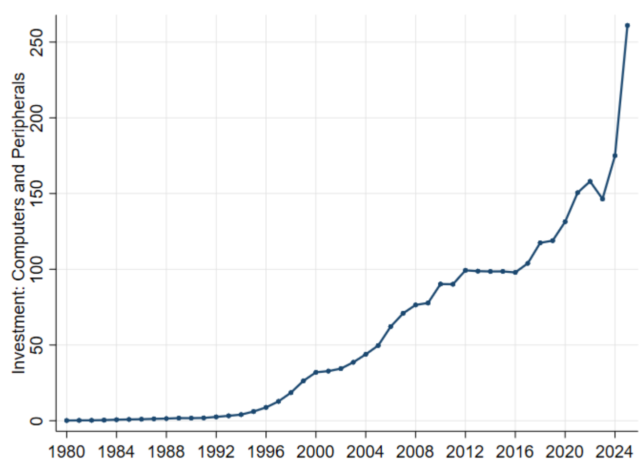

Notes: Real investment in Computers & Peripherals from the BEA NIPA accounts, in billions of dollars.

The parameters of the simplified production function (10) are determined in two tiers. The structures exponent α_1 is calibrated inside the estimation as the sample mean of the structures income share, $\alpha_1 = \text{mean}\,[(\bar{\rho} + \delta_{s,t})k_{s,t}/Y_t]$, evaluated at an externally declared required return $\bar{\rho}$

[15]Our proxy is deliberately narrower than the broad information-technology capital concept used by Eden and Gaggl (2018) to measure automation capital: software and the remaining IT categories exhibit price and depreciation dynamics different from those of AI hardware, so restricting to Computers & Peripherals keeps the proxy aligned with the quality-adjusted price behavior that identifies the AI-specific moments.

taken from the all-capital return estimates of **?**; α_2 then follows from the adding-up condition $\alpha_2 = 1-\alpha_1-\bar{\pi}$, with the profit share $\bar{\pi}$ declared at its NIPA value. The remaining seven parameters – ϕ, μ_H, ξ_H, γ_H, λ_H, the labor-efficiency term ϵ_u, and the scale $\bar{A}$ – are estimated by the Generalized Method of Moments (GMM). The CES aggregators are estimated in normalized form, with every input measured relative to its sample geometric mean, so the estimated weights μ_H and ξ_H are income shares at the sample reference point rather than unit-dependent constants.[16] 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. A **?** relative-demand condition for the two skill types;
6. The output level;
7. The user-cost (Euler) conditions for structures and AI at the required return $\bar{\rho}$;
8. The capital-demand conditions for structures, equipment, and AI.

The moments are interacted with an instrument set consisting of a constant, the structures stock $k_{s,t}$, the lagged equipment price $P_{e,t-1}$, relative skill supply for the Katz–Murphy condition, and the hedonic AI price of Section 2 as an instrument for the two gated AI conditions, yielding 22 orthogonality conditions. Appendix C.1 documents the implementation in detail: the sample and the gating of the AI-specific moments, the mapping from moments and instruments to the degrees of freedom of the overidentification test, the two-step weighting matrix and HAC inference, and the estimation constraints. Table 1 reports the point estimates and standard errors.

The estimates yield the following implied elasticities of substitution:

- $\sigma_\lambda = \frac{1}{1-\lambda_H} \approx 0.89$: AI and the high-skill–equipment composite are *complements* ($\sigma_\lambda < 1$).
- $\sigma_\gamma = \frac{1}{1-\gamma_H} \approx 0.24$: High-skill labor and equipment are strong complements ($\sigma_\gamma < 1$), consistent with the capital-skill complementarity finding of Krusell et al. (2000).
- $\sigma_\phi = \frac{1}{1-\phi} \approx 2.7$: High- and low-skill composites are substitutable ($\sigma_\phi > 1$).

[16] Labor enters in efficiency units following the specification of Krusell et al. (2000) as implemented by Ohanian et al. (2023), with the relative efficiency of the two skill types carried by the estimated term ϵ_u. Under the normalization, the units constants that the legacy parametrization absorbs into the CES weights are instead carried explicitly, which is what makes the weights interpretable as reference-point income shares.

Table 1: Estimated parameters (GMM)

Parameter	Estimate
ϕ	0.630 (0.005)
α_1^*	0.094
α_2^*	0.864
μ_H	0.976 (0.001)
ξ_H	0.676 (0.005)
γ_H	−3.174 (0.098)
λ_H	−0.118 (0.010)
ϵ_u	−0.729 (0.013)

Conditional standard errors in parentheses. *Calibrated in-run and held fixed during estimation; see text. The scale A is also estimated.

The AI weight is small, $(1 - \mu_H) \approx 2.4$ percent – under the normalization, the AI share of the high-skill bundle's income at the sample reference point – reflecting the fact that AI capital remains a small fraction of the total capital stock. The weight on human capital inside the high-skill–equipment composite is high, $\xi_H \approx 0.68$, so high-skill labor still dominates that composite. The new dimension is the complementarity between AI and the high-skill–equipment composite ($\sigma_\lambda < 1$), which implies that expanding AI is most productive when high-skill labor and equipment expand alongside it.

4.1.1 Identification and Robustness of Estimates

Identification of σ_λ comes from the AI-specific moments. The AI vs. structures no-arbitrage condition, instrumented by the hedonic AI price of Section 2, identifies λ_H and hence the AI–skill complementarity. The Katz–Murphy relative-demand condition pins down the top-level substitutability ϕ, which the wage-bill and labor-share moments alone leave loosely determined. The three Krusell et al. (2000) moments, namely the wage-bill ratio, labor share, and equipment vs. structures no-arbitrage condition, cannot alone identify the full parameter vector. The overidentifying restrictions test yields $J = 15.3$ with $p = 0.43$ (15 degrees of freedom, counting the two calibrated-and-pinned exponents among the parameters), so the model is not rejected at conventional significance levels. The standard errors in Table 1 are conditional on the calibrated (α_1, α_2); propagating the calibration uncertainty in $\bar{\rho}$ and $\bar{\pi}$ widens the standard error of λ_H by roughly a factor of four while leaving the complementarity finding intact, and the block bootstrap below measures the same total uncertainty directly.

The appendix documents the robustness of these estimates along several dimensions. The profile of the GMM objective in λ_H has a clear interior minimum, so the complementarity is estimated rather than imposed by the estimation constraints (Appendix C.1); a block bootstrap keeps both implied elasticities below one in every replication (Appendix C.2); and the estimates are insensitive both to how we split AI and equipment to avoid double counting of AI capital (Appendix C.3) and, finally, to alternative start years for the AI capital series (Appendix C.4).

4.2 Externally Calibrated Parameters

The population shares of low- and high-skill workers, $M_L = 0.575$ and $M_H = 0.350$, follow the college-attainment-based skill classification of the U.S. workforce used by Ohanian et al. (2023), with the worker shares adjusted for the entrepreneur mass M_E so that population shares sum to one (authors' calculations). The entrepreneur share $M_E = 0.075$ is set to the fraction of households in the 2022 Survey of Consumer Finances that own a privately held business but are not themselves self-employed, which we compute from the SCF microdata. The corresponding self-employment share of 11.2 percent matches Kuhn and Ríos-Rull (2025). The three depreciation rates are set to long-run averages of the corresponding BEA series: structures depreciate slowly ($\delta_s = 0.028$), while equipment and AI hardware depreciate rapidly ($\delta_e = 0.102$, the mean of the efficiency-unit equipment rate implied by the quality-adjusted stocks, and $\delta_{AI} = 0.156$, the rate we also use in the measurement of Section 2, following Kudoh and Miyamoto, 2025). The relative prices of the two reproducible capital goods enter through the investment-specific productivity terms q and ζ, with $P_e = 1/q$ and $P_{AI} = 1/\zeta$, calibrated to 2000–2024 averages of the corresponding quality-adjusted price series deflated by the GDP deflator. The Frisch elasticity of labor supply is set to $\vartheta = 0.50$ following Keane (2011). Finally, the consumption and capital income tax rates, $\tau_c = 0.064$ and $\tau_k = 0.273$, are the effective U.S. rates estimated by Carey and Rabesona (2003), and neutral productivity A is normalized to one. Table 2 reports these parameters and their sources.

4.3 Internally Calibrated Parameters

The remaining parameters are set jointly to match features of the U.S. fiscal system and aggregate labor supply. The two-bracket labor tax rates τ_L and τ_H target a labor-income-weighted average labor tax of approximately 21 percent, consistent with the estimates in McDaniel (2007) and the calibration in Carroll et al. (2024).[17] The discount factor β is set so that the steady-state capital–output ratio equals the U.S. value of 2.35 in Trabandt and Uhlig (2011), whose fiscal calibration we follow throughout.[18] The ownership share θ_H targets a capital income share of

[17] The average-tax target does not by itself pin down the spread between the two bracket rates. We impose $\tau_H = 3\tau_L$ and calibrate the level, and the distributional results should be read as conditional on that degree of progressivity.

[18] The three capital stocks are valued at investment-good prices, $K = k_s + k_e/q + k_{AI}/\zeta$, so the ratio is measured in consumption units on both sides.

Table 2: Externally calibrated parameters

Parameter	Value	Source
M_E	0.075	Business owners, not self-employed (2022 SCF; Kuhn and Ríos-Rull, 2025)
M_H	0.350	Adjusted share of high-skill in Ohanian et al. (2023)
M_L	0.575	Adjusted share of low-skill in Ohanian et al. (2023)
δ_s	0.028	Avg. Structures Depreciation rate from BEA
δ_e	0.102	Avg. efficiency-unit Equipment Depreciation rate (BEA; authors' calc.)
δ_{AI}	0.156	AI Depreciation rate, Kudoh and Miyamoto (2025)
ϑ	0.50	Frisch elasticity, Keane (2011)
q	0.74	Equipment technology parameter ($P_e = 1/q$); 2000–2024 avg.
ζ	0.77	AI technology parameter ($P_{AI} = 1/\zeta$); 2000–2024 avg.
A	1	Normalization
τ_c	0.064	Consumption tax rate; Carey and Rabesona (2003)
τ_k	0.273	Capital income tax rate; Carey and Rabesona (2003)

approximately 17 percent in high-skill workers' total income, consistent with the 2022 Survey of Consumer Finances for college-educated households (Kuhn and Ríos-Rull, 2025). The transfer share T/Y is set to 2.2 percent of GDP, within the range of 1.3–2.8 percent identified from CBO mandatory spending data (Congressional Budget Office, 2019) and the average mandatory outlays for the income security system tabulated by the Office of Management and Budget (Office Of Management And Budget, 2023), following the calibration in Luduvice (2024) and Carroll et al. (2024). Finally, the disutility scaling parameter B is calibrated so that the population-weighted average labor supply equals one-third of the time endowment:

$$\frac{M_L L_L + M_H L_H}{M_L + M_H} = \frac{1}{3}$$

Table 3 reports these parameters alongside their targets and model values.

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 first describe key properties of the calibrated steady state, then present transitory impulse responses as a validation exercise, and finally study permanent AI shocks and the interaction between them.

Table 3: Internally calibrated parameters

Parameter	Value	Target / Source	Model
β	0.964	Capital–output ratio K/Y = 2.35 (Trabandt and Uhlig, 2011)	2.35
τ_L τ_H	0.104 0.312	Wtd. avg. labor tax ≈ 21% (McDaniel, 2007), with $\tau_H = 3\tau_L$	21.0%
θ_H	0.427	H capital income ≈ 17% of H income (SCF; Kuhn and Ríos-Rull, 2025)	16.8%
T/Y	0.022	Transfers = 2.2% of GDP (Congressional Budget Office, 2019; Office Of Management And Budget, 2023)	2.2%
B	24.28	Avg. labor supply = 1/3	0.333

5.1 Baseline Calibration Results

Table 4 collects the main features of the calibrated steady state. Low-skill workers supply slightly more labor than high-skill workers, reflecting the interaction between progressive taxation and the wage differential: the higher marginal tax rate on high-skill earnings reduces the after-tax return to their labor supply. Government purchases absorb about a fifth of output broadly consistent with the values in Trabandt and Uhlig (2011) and Conesa et al. (2023). The capital–output ratio is 2.35 by construction, and pure firm profits are a small share of GDP.

The Euler equations equalize after-tax net returns across all capital types, as given by the user-cost conditions in Appendix B.2. Because the capital income tax applies to income net of depreciation, the net return per unit of capital value is the same for structures, equipment, and AI. Thus, what differs across types is the *gross* rental rate per unit of capital value, which must also cover depreciation and is therefore highest for AI and lowest for structures. The capital stock divides its value roughly evenly between structures and equipment, mirroring the income-share split the estimation fits in the data, while AI remains a small slice of both value and rental income.

5.2 Transitory Impulse Responses

Before turning to permanent shocks, we study two transitory impulse-response exercises to study and validate the model's dynamic properties.

Transitory Halving of AI Price. We subject the economy to a transitory reduction in the AI price: ζ_t doubles at $t = 1$ and then decays deterministically back to baseline along an AR(1) process

Table 4: Baseline steady state: key moments

Ratio / Moment	**Model**
Labor market	
L_L, L_H	0.341, 0.321
Skill premium w_H/w_L	1.819
Wage-bill ratio $w_H m_H/w_L m_L$	1.04
Government	
G/Y	20.7%
Capital	
k_s, k_e/q, k_{AI}/ζ (value shares)	50.6%, 46.8%, 2.6%
Rental income shares $r_i k_i/Y$	9.4%, 16.7%, 1.3%
Profit share π/Y	4.2%

Notes: Key ratios and moments of the calibrated baseline steady state. $r_i^{\text{net}} \equiv r_i/P_i - \delta_i$ is the return per unit of capital value net of depreciation. The entrepreneur's Euler equations equate it across the three capital types (Appendix B.2). Capital value shares and rental income shares are evaluated at investment-good prices ($P_s = 1$, $P_e = 1/q$, $P_{AI} = 1/\zeta$). The three rental income shares, the labor share (68.4%), and the profit share sum to one.

with $\rho = 0.90$ (a half-life of roughly seven years). This exercise isolates the investment-specific channel. Because AI is temporarily cheap, the entrepreneur front-loads AI investment, producing a sharp rise in K_{AI} that peaks early and then unwinds as the price reverts to baseline. Output barely moves on impact: at $t = 1$ all capital stocks are predetermined, so only the labor-supply response can shift it. Output rises from $t = 2$ onward, once the freshly accumulated AI capital enters production, and then decays as the price reverts. The skill premium rises temporarily, as AI–skill complementarity raises the marginal product of the high-skill bundle. Entrepreneur consumption traces a characteristic "invest-then-enjoy" path: C_E dips on impact as resources are redirected toward the now-cheap AI, then rises above its initial level as the returns to the additional AI capital materialize. The initial dip and the later gain are of comparable size. Figure 5 displays the impulse responses.

1 Percent TFP Shock. We also subject the economy to a 1 percent positive innovation in neutral productivity A_t, which follows an AR(1) process with persistence $\rho = 0.90$. Output, consumption, and investment all rise on impact and decay gradually as the shock dissipates. The responses are consistent with standard neoclassical dynamics: the entrepreneur increases investment in all three capital types, wages rise for both skill groups, and the skill premium is essentially unchanged. The impulse responses are reported in Appendix Figure A.5.

Figure 5: Transitory impulse response: halve AI price ($\rho = 0.90$)

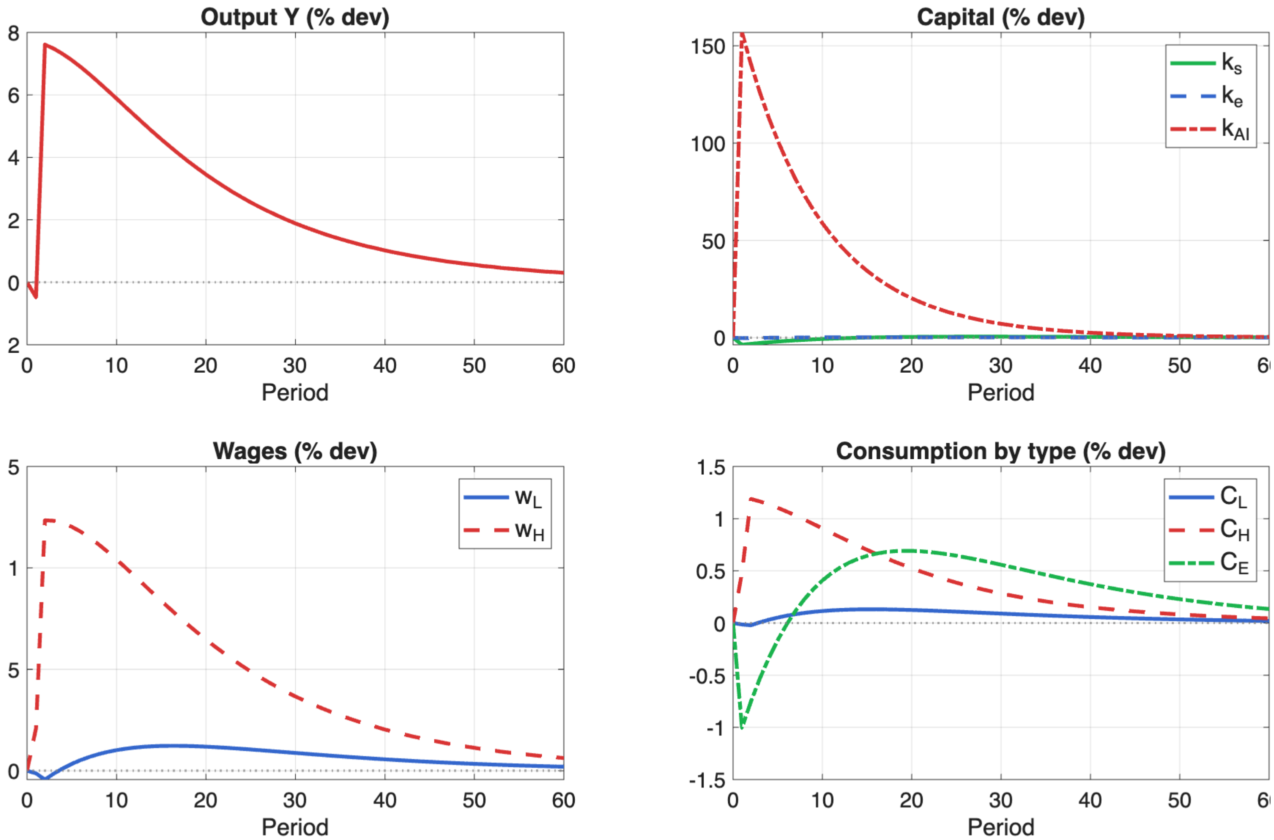

Notes: Perfect-foresight impulse responses (percent deviation from the initial steady state) to a transitory halving of the AI price: ζ_t doubles on impact and decays back to baseline following an AR(1) process with persistence $\rho = 0.90$. AI investment is front-loaded while the price is low, output rises with a one-period lag because capital is predetermined, and the skill premium rises temporarily through AI–skill complementarity.

5.3 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 about 2.4 percent to about 4.8 percent, 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 percent – 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 5 summarizes the long-run (new steady-state) effects of the first two experiments.[19]

[19] To keep the discussion focused on economic mechanisms, the full set of perfect-foresight transition-path figures

Table 5: Counterfactual experiments: new steady-state effects

	Initial SS	**Double Share**	**Halve Price**
Output Y (% Δ)	—	-1.87	$+1.00$
AI Capital K_{AI} (% Δ)	—	$+81$	$+89$
Skill Premium w_H/w_L	1.819	1.726	1.849
C_L (% Δ)	—	-0.68	$+0.36$
C_H (% Δ)	—	-5.01	$+1.83$
C_E (% Δ)	—	-0.28	$+0.72$

The two individual experiments induce distinct 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).

5.3.1 Experiment 1: Doubling the AI Usage Share

The first experiment is contractionary (see Appendix Figure A.8). Increasing the AI weight raises the desired role of AI in production, but AI is initially a small share of the capital stock. Since capital stocks are predetermined on impact, the production function immediately demands more of a scarce input. With $\sigma_\lambda < 1$, this reweighting creates a costly capital–technology mismatch: the composite the economy is well stocked with cannot easily substitute for the AI capital it lacks. Tilting weights toward the scarce input and away from the dominant high-skill–equipment composite creates a bottleneck. As a result, output falls sharply on impact and remains about 1.9 percent below the initial steady state in the long run.

The adjustment in capital composition reinforces this result. The higher AI weight raises the marginal product of AI, inducing a large increase in AI investment and an eventual 81 percent increase in the AI stock. At the same time, the marginal product of equipment falls, and equipment is crowded out. The Euler equations pin rental rates to user costs, which are independent of μ_H, but the composition of capital shifts toward AI – small, expensive, and fast-depreciating – and away from equipment. Appendix Figure A.6 shows this reallocation across capital types.

The distributional effects are also asymmetric. Low-skill consumption falls only modestly, while high-skill consumption falls substantially, due to a lower skill premium and reduced capital income share. Entrepreneur consumption barely moves: the permanently lower output is largely offset by the returns on the much larger AI stock. Appendix Figures A.7 and A.8 show these consumption, wage, and output dynamics.

is collected in Appendix D.2.

5.3.2 Experiment 2: Halving the AI Price

The price reduction produces a smooth, monotone output expansion. Unlike the usage-share shock, the AI price shock does not change the production function directly. Instead, it relaxes the resource constraint by making AI capital cheaper to accumulate. Since capital is predetermined, output does not jump on impact. The expansion occurs as entrepreneurs redirect investment toward AI and the larger AI stock enters production over time.

This is an investment-specific technological change in the spirit of Greenwood et al. (1997) and Krusell et al. (2000). When ζ doubles, the relative investment price $P_{AI} = 1/\zeta$ falls by half. Each unit of resources devoted to AI investment now purchases twice as much effective AI capital. The Euler equation requires a lower rental rate on AI, and AI accumulates until its marginal product falls to the new lower user cost. In the long run, the AI capital stock rises by about 90 percent. Unlike in the share-increase experiment, this capital deepening does not require crowding out the dominant high-skill–equipment composite. Appendix Figure A.6 shows that the price reduction generates broader capital deepening rather than a sharp reallocation away from equipment.

The output gain is positive but moderate, about 1 percent in the new steady state. The reason is that AI remains a small-weight input in the production function. Cheaper AI capital raises the high-skill bundle and therefore output, but the small value of $(1 - \mu_H)$ limits the aggregate effect.

The distributional effects move in the opposite direction from Experiment 1. The larger AI stock raises the marginal product of the high-skill–equipment composite, so the skill premium increases. High-skill workers gain through higher wages and capital-income exposure. Entrepreneurs also gain in the long run, although their consumption initially reflects the usual invest-then-enjoy pattern: resources are shifted toward AI accumulation early in the transition, and consumption rises once the larger AI stock generates returns. Low-skill workers remain nearly unaffected. These paths are reported in Appendix Figures A.7 and A.8.

5.4 Combined Experiment and Super-Additivity

The third experiment applies both shocks simultaneously. The two shocks are *super-additive*: applied together they leave output essentially unchanged, whereas the sum of their separate effects is a decline of about 0.9 percent. The interaction term is roughly 0.9 percentage points and accounts for the entire difference between the combined outcome and the naive sum. The combined shock still delivers less output than the price decline on its own.

Table 6 reports the full set of results, including the hypothetical linear sum for comparison.

Table 6: Combined experiment: super-additivity

	Initial SS	**Dbl. Share**	**Halve Price**	**Both**	**Exp1+Exp2**
Aggregate (% Δ from initial SS)					
Output Y	—	-1.87	$+1.00$	-0.01	-0.87
AI Capital K_{AI}	—	+81	+89	+245	+169
Distributional					
Skill Premium w_H/w_L	1.819	1.726	1.849	1.781	—
C_L (% Δ)	—	-0.68	$+0.36$	-0.00	—
C_H (% Δ)	—	-5.01	$+1.83$	-1.68	—
C_E (% Δ)	—	-0.28	$+0.72$	$+1.13$	—
Super-additivity (pp)					
Both $-$ (Exp1+Exp2)	—	—	—	$+0.86$	—

Notes: The super-additivity row is computed from unrounded model output. Individual entries are rounded independently.

We can observe that the contraction generated by the share increase in isolation is almost exactly undone once the price drop is layered on, even though the price drop alone would deliver a clear expansion. The combined-shock paths appear as the "Both" series in the overview of Appendix Figures A.6, A.7, and A.8.

The combined experiment compresses the skill premium only modestly, as the two effects on high-skill wages largely offset each other. Entrepreneurs are the biggest winners, capturing the returns from the expansive and cheap AI capital accumulation. Low-skill workers are essentially unaffected, while high-skill workers see a slight consumption decline. Appendix Figures A.13 and A.14 provide additional detail on consumption by type and income distribution dynamics.

Appendix D.5 tests the robustness of this result and repeats all three experiments across three values of σ_λ (0.89, 1.20, and 1.49) and three values of δ_{AI} (0.10, 0.156, and 0.25). The interaction gain stays between 0.6 and 0.9 percentage points throughout, including where AI and the high-skill–equipment composite are substitutes. The level of the combined output effect differs: it falls as either parameter rises, and is positive only at the lowest values we consider.

5.4.1 Mechanism Behind the Super-Additivity

Each experiment is limited by what the other does not supply. Experiment 1 raises the marginal product of AI without lowering its price, so the economy would like more AI capital than it can profitably accumulate. The higher weight is applied to a stock that stays small. Experiment 2 lowers the price without raising the weight, so AI capital nearly doubles, but a CES weight of about 2.4 percent caps what the extra capital can contribute.

When both shocks hit simultaneously, each channel relaxes the constraint that binds the other. The price drop makes it cheap to accumulate the AI capital that the higher weight calls for. The higher weight raises what each unit of that cheap capital contributes to production.

The mechanism operates within a single input: the weight and the inverse price of the same input move together, in contrast to the between-input complementarity ($\sigma_\lambda < 1$) that the estimation identifies. In our calibration, 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 larger when the weight is higher.

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 arrival of complementary applications. Our results are thus able to capture qualitatively how the two reinforce each other.

6 Policy Experiments

We now study policy counterfactuals that introduce a separate tax on AI capital income and use the proceeds to fund redistributive transfers. We conduct the following taxation counterfactual: decouple the capital income tax into two rates, a rate $\tau_{k,AI}$ on AI capital rental income and the baseline rate τ_k on structures, equipment, and profits and vary the AI capital income tax.[20] Government purchases G are held fixed at their baseline level, and all additional revenue raised by the AI capital tax is rebated as per-capita transfers T_{pc}. This experiment operationalizes, within our model, a policy in the spirit of the "digital dividend" proposed by Marinescu (2025) (and others) where the government can levy a tax on AI-related capital income and rebate it lump sum to the population.

This setup traces a "spending" Laffer curve for the AI capital tax in the spirit of Trabandt and Uhlig (2011), focusing on the positive mapping the counterfactual tax experiment. The intuition for the Laffer curve is standard. As $\tau_{k,AI}$ rises, revenue (and transfers) first increases, peaks, and then declines as the shrinking AI capital base dominates.

Figure 6 traces these transfer Laffer curves. The AI-capital Laffer curve is relatively flat and peaks at a high rate, around $\tau_{k,AI} \approx 92$ percent in the baseline economy, with revenue within a tenth of a percent of its maximum for any rate between 87 and 94 percent – the peak's location is loosely identified precisely because the curve is so flat. The peak rate is far above the baseline capital income tax because the tax falls on capital income net of depreciation, so it distorts only

[20] Formally, net ownership income splits as $\Omega^{\text{net}} = \Omega^{\text{net}}_{-AI} + \Omega^{\text{net}}_{AI}$, where $\Omega^{\text{net}}_{AI} = r_{AI}k_{AI} - (\delta_{AI}/\zeta)\, k_{AI}$ is AI rental income net of AI depreciation and $\Omega^{\text{net}}_{-AI}$ collects structures and equipment rental income and profits, net of their depreciation. Capital tax revenue is $\tau_k\, \Omega^{\text{net}}_{-AI} + \tau_{k,AI}\, \Omega^{\text{net}}_{AI}$, the payment to high-skill workers becomes $\theta_H\, [(1-\tau_k)\, \Omega^{\text{net}}_{-AI} + (1-\tau_{k,AI})\, \Omega^{\text{net}}_{AI}]$, and the AI Euler equation carries $(1-\tau_{k,AI,t+1})$ in place of $(1-\tau_{k,t+1})$. The other Euler equations are unchanged.

the net return; expressed as a rate on AI capital income gross of depreciation, the same peak corresponds to roughly 69 percent. Because AI capital is a small share of the total capital stock, the AI tax alone generates only modest additional revenue, leaving limited fiscal space for transfers. Even at the Laffer curve's peak, the per-capita transfer rises only to a few percent of average worker income. Halving the AI price enlarges the AI capital base, shifting the entire Laffer curve upward so that the same tax supports a somewhat larger transfer.[21]

Figure 6: AI capital income tax: transfer Laffer curves

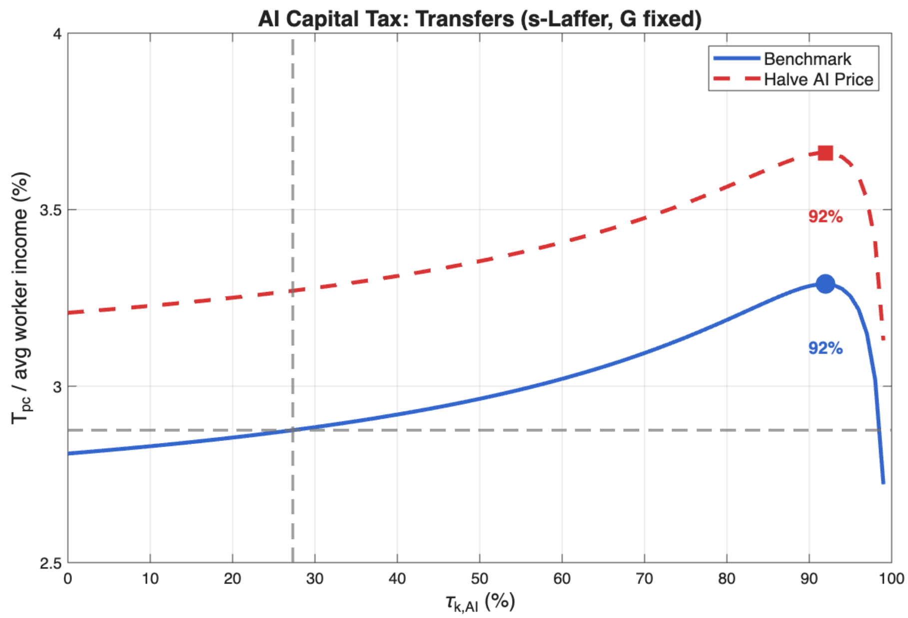

Notes: Per-capita transfer T_{pc} as a percent of average worker income financed by the AI capital income tax $\tau_{k,AI}$, with government purchases G held fixed at baseline (spending Laffer curve). Solid: baseline economy. Dashed: halved AI price. Markers indicate the revenue-maximizing rate, $\tau_{k,AI} \approx 92\%$ in both economies. The dashed gray lines mark the baseline $\tau_{k,AI} = 27.3\%$ and the corresponding transfer.

We consider three policy allocations.

Policy (i): AI tax at the Laffer peak (baseline economy). Holding the AI price at its baseline, we set $\tau_{k,AI}$ to its revenue-maximizing value of roughly 92 percent. The higher tax depresses AI accumulation, so output falls and the additional revenue funds only a marginally larger transfer – a direct consequence of AI's small tax base.

Policy (ii): Halve AI price + AI tax at the Laffer peak. When the AI price is permanently halved (as in Experiment 2), the AI capital base expands, so the same tax – which peaks at the same 92 percent rate – raises more revenue, and the larger base cushions the output cost of the tax. The transfer climbs to closer to four percent of average worker income.

[21]The corresponding revenue Laffer curve is shown in Appendix Figure A.20.

Policy (iii): Halve AI price + larger transfer. Finally, we ask how far transfers can be pushed by combining the AI capital tax with a higher consumption tax. Funding a per-capita transfer of about \$12,000 per year – roughly a fifth of average worker income – requires raising τ_c to about 28 percent alongside an AI tax of 84 percent.[22] The scale of this program places it in the territory of the quantitative literature on universal basic income (Luduvice, 2024; Daruich and Fernández, 2024; Conesa et al., 2023; Guner et al., 2023).

That literature consistently finds that a UBI of this size requires a large increase in broad-based taxation. The difference here is how much of the fiscal effort the AI base can carry: the allocation that funds the transfer pairs a consumption tax of about 28 percent with an AI capital tax of 84 percent, three times the baseline capital income rate and most of the way to the AI Laffer peak, so the AI base is taxed well above the baseline capital income rate. Appendix Figure A.21 displays the iso-revenue surface over $(\tau_c, \tau_{k,AI})$ and the Laffer ridge that locates this allocation. The downward-sloping ridge shows that the consumption tax and the AI capital tax act as substitutes in raising revenue.

Table 7 reports the three policy allocations. The first two keep the consumption tax at its baseline level, while Policy (iii) raises it along the Laffer ridge.

Table 7: Policy experiments: AI capital income tax at the Laffer peak

	Baseline	**Policy (i)** (base, peak)	**Policy (ii)** (halve, peak)	**Policy (iii)** (halve, \$12k)
τ_c	0.064	0.064	0.064	0.280
$\tau_{k,AI}$	0.273	0.92	0.92	0.84
Output Y (% Δ)	—	-1.83	-0.85	-5.17
K_{AI} (% Δ)	—	-63	-31	$+1$
Skill Premium w_H/w_L	1.819	1.766	1.799	1.804
T/avg. worker income	2.9%	3.4%	3.7%	20.0%

Notes: Policies (i) and (ii) set $\tau_{k,AI}$ at its revenue-maximizing (Laffer-peak) value in the baseline and halved-price economies, respectively, keeping τ_c at baseline. Policy (iii) funds the \$12,000 per-capita transfer under the halved price, with $(\tau_c, \tau_{k,AI})$ chosen on the Laffer ridge. The halved-price economy without any policy change is reported in Table 5.

6.1 Welfare Analysis

We evaluate welfare using a consumption-equivalent variation (CEV). Let V_j^{shock} denote agent j's discounted lifetime utility along the transition path, and V_j^0 the discounted utility of remaining at

[22] The transfer target is set at 18.9 percent of baseline average worker income, which corresponds to \$12,000 at the average worker income implied by the calibration, about \$63,500.

the initial steady state forever. The CEV λ_j solves:

$$\sum_{t=0}^{\infty} \beta^t \left[\ln\big((1+\lambda_j)\, C_{j,0}\big) - B \frac{L_{j,0}^{1+1/\vartheta}}{1+1/\vartheta} \right] = V_j^{\text{shock}}$$

yielding $\lambda_j = \exp\big[(1-\beta)(V_j^{\text{shock}} - V_j^0)\big] - 1$. The CEV captures changes in both consumption and labor disutility along the full transition path, expressed as a permanent consumption equivalent.[23] Table 8 reports the CEV for each experiment and agent type, along with a population-weighted average.

Table 8: Welfare: consumption-equivalent variation (%)

	AI shocks			**Policy experiments**		
	Dbl. Share	Halve Price	Both	(i)	(ii)	(iii)
CEV_L	−0.70	+0.22	−0.22	+0.38	+1.05	+3.70
CEV_H	−4.61	+1.46	−1.87	−2.24	−0.42	−4.62
CEV_E	−0.45	−0.20	−0.30	−0.74	−0.71	−8.17
CEV_{wtd}	−2.05	+0.62	−0.80	−0.63	+0.40	−0.10

Notes: Policy experiments: (i) AI tax at the Laffer peak in the baseline economy; (ii) halved AI price with AI tax at the Laffer peak; (iii) halved AI price with a larger $12,000 per-capita transfer. CEV_{wtd} is the population-weighted average.

Low-skill workers are close to insulated from the pure AI shocks: their consumption-equivalent gains and losses stay within seven tenths of a percentage point, reflecting the substitutability between the two skill bundles. The combined shock is welfare-negative in aggregate even though it leaves output essentially unchanged, because the gain to entrepreneurs does not offset the losses to both worker types. Among the policy experiments, taxing AI capital alone (Policy i) is modestly welfare-negative in aggregate, with the low-skill gain outweighed by the high-skill and entrepreneur losses

Pairing the AI tax with a fall in the AI price (Policy ii) turns the package welfare-positive in aggregate and for low-skill workers. Pushing redistribution further (Policy iii) delivers by far the largest low-skill welfare gain, but at a steep cost to high-skill workers and entrepreneurs, leaving the population-weighted effect roughly neutral: the large population share of low-skill workers almost exactly absorbs the concentrated losses.

[23] Given that labor supply adjusts along the transition, an experiment that raises steady-state consumption can be welfare-reducing if the transition requires increased labor supply or front-loaded investment. This explains, for example, why CEV_E is negative under the halve-price experiment despite inducing a higher terminal C_E as shown in the counterfactual analysis.

7 Conclusion

This paper proposes a theory and measurement of AI as capital and studies its macroeconomic and distributional consequences. We find that AI and the high-skill–equipment composite are complements. Moreover, our results show that the channel through which AI improves matters: a broader AI mandate and a cheaper AI price have opposite aggregate effects. These two channels are also super-additive. Moving both together produces more output than the sum of the separate effects, because each relaxes the constraint that limits the other, and the interaction is positive across every elasticity and depreciation rate we consider. In our calibration, the combined shock still yields less output than the price decline alone. Finally, a tax on AI capital income can fund redistributive transfers, though its revenue potential is limited while AI is a small share of the capital stock. Pairing the AI tax with a fall in the AI price makes such redistribution welfare-improving for low-skill workers and in the aggregate at modest cost to output.

Several extensions merit future investigation. The production function could be generalized to allow low-skill workers to interact with equipment or AI, which would alter the insulation result. A multi-sector extension would allow the AI-producing and AI-using sectors to be modeled separately, connecting AI price dynamics more directly to equilibrium. Finally, incorporating richer household heterogeneity – including savings decisions for workers and a distribution of capital holdings – would sharpen the welfare analysis and connect more tightly to the distributional evidence from household surveys.

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 David Card and Orley Ashenfelter, editors, *Handbook of Labor Economics*, volume 4, pp. 1043–1171. Elsevier. doi:10.1016/S0169-7218(11)02410-5. URL `https://www.sciencedirect.com/science/article/pii/S0169721811024105`.

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, Fredric Kong, and Pascual Restrepo (2025). "Tasks at work: comparative advantage, technology and labor demand." In Christian Dustmann and Thomas Lemieux, editors, *Handbook of Labor Economics*, volume 6 of *Handbook of Labor Economics*, pp. 1–114. Elsevier. doi:10.1016/bs.heslab.2025.08.003. URL `https://www.sciencedirect.com/science/article/pii/S1573446325000100`.

Acemoglu, Daron and Pascual Restrepo (2018). "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review*, 108(6), pp. 1488–1542. doi:10.1257/aer.20160696. URL `https://www.aeaweb.org/articles?id=10.1257/aer.20160696`.

Acemoglu, Daron and Pascual Restrepo (2022). "Tasks, Automation, and the Rise in U.S. Wage Inequality." *Econometrica*, 90(5). doi:https://doi.org/10.3982/ECTA19815. URL `https://onlinelibrary.wiley.com/doi/full/10.3982/ECTA19815`.

Aghion, Philippe, Simon Bunel, Xavier Jaravel, Thomas Mikaelsen, Alexandra Roulet, and Jakob Søgaard (2025). "How Different Uses of AI Shape Labor Demand: Evidence from France." *AEA Papers and Proceedings*, 115, pp. 62–67. doi:10.1257/pandp.20251047. URL `https://pubs.aeaweb.org/doi/10.1257/pandp.20251047`.

Aghion, Philippe, Benjamin Jones, and Charles Jones (2018). "Artificial Intelligence and Economic Growth." In *The Economics of Artificial Intelligence: An Agenda*, pp. 237–282. National Bureau of Economic Research, Inc. URL `https://EconPapers.repec.org/RePEc:nbr:nberch:14015`, series Title: NBER Chapters.

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

Autor, David (2024). "Applying AI to Rebuild Middle Class Jobs." doi:10.3386/w32140. URL https://www.nber.org/papers/w32140.

Baily, Martin, David Byrne, Aidan Kane, and Paul Soto (2025). "Generative AI at the Crossroads: Light Bulb, Dynamo, or Microscope?" Working paper, arXiv.org. URL https://EconPapers.repec.org/RePEc:arx:papers:2505.14588.

Beaudry, Paul, Mark Doms, and Ethan Lewis (2010). "Should the Personal Computer Be Considered a Technological Revolution? Evidence from U.S. Metropolitan Areas." *Journal of Political Economy*, 118(5), pp. 988–1036. doi:10.1086/658371. URL https://www.journals.uchicago.edu/doi/10.1086/658371.

Benabou, Roland (2002). "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?" *Econometrica*, 70(2), pp. 481–517. URL https://ideas.repec.org//a/ecm/emetrp/v70y2002i2p481-517.html.

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."

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock (2018). "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?" *AEA Papers and Proceedings*, 108, pp. 43–47. doi:10.1257/pandp.20181019. URL https://www.aeaweb.org/articles?id=10.1257/pandp.20181019.

Carey, David and Josette Rabesona (2003). "Tax Ratios on Labour and Capital Income and on Consumption." *OECD Economic Studies*, 2002(2), pp. 129–174. URL https://ideas.repec.org//a/oec/ecokaa/5lmqcr2jj2kh.html.

Carroll, Daniel, André Victor Doherty Luduvice, and Eric R. Young (2024). "Optimal Fiscal Reform with Many Taxes." doi:10.2139/ssrn.4352673. URL https://papers.ssrn.com/abstract=4352673.

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

1685. doi:10.1257/aer.20211478. URL https://www.aeaweb.org/articles?id=10.1257/aer.20211478.

Caunedo, Julieta and Elisa Keller (2022). "Technical change and the demand for talent." *Journal of Monetary Economics*, 129, pp. 65–88. doi:10.1016/j.jmoneco.2022.04.006. URL https://www.sciencedirect.com/science/article/pii/S0304393222000605.

Conesa, Juan Carlos, Bo Li, and Qian Li (2023). "A quantitative evaluation of universal basic income." *Journal of Public Economics*, 223, p. 104,881. doi:10.1016/j.jpubeco.2023.104881. URL https://www.sciencedirect.com/science/article/pii/S0047272723000634.

Congressional Budget Office (2019). "The Distribution of Household Income, 2016." URL https://www.cbo.gov/publication/55413.

Costinot, Arnaud and Iván Werning (2023). "Robots, Trade, and Luddism: A Sufficient Statistic Approach to Optimal Technology Regulation." *The Review of Economic Studies*, 90(5), pp. 2261–2291. doi:10.1093/restud/rdac076. URL https://EconPapers.repec.org/RePEc:oup:restud:v:90:y:2023:i:5:p:2261-2291.

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.

Daruich, Diego and Raquel Fernández (2024). "Universal Basic Income: A Dynamic Assessment." *American Economic Review*, 114(1), pp. 38–88. doi:10.1257/aer.20221099. URL https://www.aeaweb.org/articles?id=10.1257/aer.20221099.

de Souza, Gustavo (2026). "AI in the Office and the Factory: Evidence from Administrative Software Registry Data." doi:10.2139/ssrn.5375463. URL https://papers.ssrn.com/abstract=5375463.

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

Eden, Maya and Paul Gaggl (2018). "On the welfare implications of automation." *Review of Economic Dynamics*, 29, pp. 15–43. doi:10.1016/j.red.2017.12.003. URL https://www.sciencedirect.com/science/article/pii/S1094202517301205.

Eeckhout, Jan, Christoph Hedtrich, and Roberto Pinheiro (2026). "IT and Urban Polarization." *American Economic Journal: Macroeconomics*, 18(1), pp. 223–259. doi:10.1257/mac.20220306. URL https://www.aeaweb.org/articles?id=10.1257/mac.20220306.

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.

Graetz, Georg and Guy Michaels (2018). "Robots at Work." *Review of Economics and Statistics*, 100(5), pp. 753–768. doi:https://doi.org/10.1162/rest_a_00754. URL https://direct.mit.edu/rest/article/100/5/753/58489/Robots-at-Work.

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.

Guerreiro, Joao, Sergio Rebelo, and Pedro Teles (2022). "Should Robots Be Taxed?" *The Review of Economic Studies*, 89(1 (324)), pp. 279–311. URL https://www.jstor.org/stable/27164156.

Guerrón-Quintana, Pablo, Tomoaki Mikami, and Jaromir Nosal (2024). "The Macroeconomic Implications of the Gen-AI Economy." doi:10.2139/ssrn.4989979. URL https://papers.ssrn.com/abstract=4989979.

Guner, Nezih, Remzi Kaygusuz, and Gustavo Ventura (2023). "Rethinking the Welfare State." *Econometrica*, 91(6), pp. 2261–2294. URL https://ideas.repec.org//a/wly/emetrp/v91y2023i6p2261-2294.html.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller (2025). "Artificial Intelligence and the Labor Market." *NBER Working Papers*. URL https://ideas.repec.org//p/nbr/nberwo/33509.html, number: 33509.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2017). "Optimal Tax Progressivity: An Analytical Framework." *The Quarterly Journal of Economics*, 132(4), pp. 1693–1754. URL https://www.jstor.org/stable/26372716.

Hernandez Martinez, Victor (2025). "Capital-skill complementarity in manufacturing: Lessons from the US shale boom." *European Economic Review*, 178, p. 105,072. doi:10.1016/j.euroecorev.2025.105072. URL https://www.sciencedirect.com/science/article/pii/S0014292125001229.

Holter, Hans A., Dirk Krueger, and Serhiy Stepanchuk (2019). "How do tax progressivity and household heterogeneity affect Laffer curves?" *Quantitative Economics*, 10(4), pp. 1317–1356. URL https://ideas.repec.org//a/wly/quante/v10y2019i4p1317-1356.html.

Jerbashian, Vahagn (2026). "On the elasticity of substitution between labor and ICT and IP capital and traditional capital." *Review of Economic Dynamics*, 60, p. 101,332. doi:10.1016/j.red.2026.101332. URL https://www.sciencedirect.com/science/article/pii/S1094202526000116.

Keane, Michael P. (2011). "Labor Supply and Taxes: A Survey." *Journal of Economic Literature*, 49(4), pp. 961–1075. doi:10.1257/jel.49.4.961. URL https://www.aeaweb.org/articles?id=10.1257/jel.49.4.961.

Kindermann, Fabian and Dirk Krueger (2022). "High Marginal Tax Rates on the Top 1 Percent? Lessons from a Life-Cycle Model with Idiosyncratic Income Risk." *American Economic Journal: Macroeconomics*, 14(2), pp. 319–366. doi:10.1257/mac.20150170. URL https://www.aeaweb.org/articles?id=10.1257/mac.20150170.

Korinek, Anton and Donghyun Suh (2024). "Scenarios for the Transition to AGI." doi:10.3386/w32255. URL https://www.nber.org/papers/w32255.

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.

Kuhn, Moritz and José-Víctor Ríos-Rull (2025). "Income and Wealth Inequality in the United States: An Update Including the 2022 Wave." doi:10.3386/w33823. URL https://www.nber.org/papers/w33823.

Luduvice, André Victor Doherty (2024). "The macroeconomic effects of universal basic income programs." *Journal of Monetary Economics*, 148, p. 103,615. doi:10.1016/j.jmoneco.2024.103615. URL https://www.sciencedirect.com/science/article/pii/S0304393224000680.

Marinescu, Ioana (2025). "Resilient by Design: Dual Safety Nets for Workers in the AI Economy." Publication Title: The Digitalist Papers.

McDaniel, Cara (2007). "Average tax rates on consumption, investment, labor and capital in the {OECD} 1950-2003." URL https://www.caramcdaniel.com/research.

Office Of Management And Budget (2023). "Historical Tables, Budget of the U.S. Government." URL https://www.whitehouse.gov/omb/information-resources/budget/historical-tables/.

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].

Thuemmel, Uwe (2023). "Optimal Taxation of Robots." *Journal of the European Economic Association*, 21(3), pp. 1154–1190. doi:10.1093/jeea/jvac062. URL https://EconPapers.repec.org/RePEc:oup:jeurec:v:21:y:2023:i:3:p:1154-1190.

Trabandt, Mathias and Harald Uhlig (2011). "The Laffer curve revisited." *Journal of Monetary Economics*, 58(4), pp. 305–327. doi:10.1016/j.jmoneco.2011.07.003. URL https://www.sciencedirect.com/science/article/pii/S030439321100064X.

Webb, Michael (2020). "The Impact of Artificial Intelligence on the Labor Market." URL https://conference.nber.org/confer/2020/YSAIf20/Required%2011.pdf.

Zeira, Joseph (1998). "Workers, Machines, and Economic Growth." *The Quarterly Journal of Economics*, 113(4), pp. 1091–1117. doi:10.1162/003355398555847. URL https://EconPapers.repec.org/RePEc:oup:qjecon:v:113:y:1998:i:4:p:1091-1117.

Appendix

“AI-Augmented Capital-Skill Complementarity”

André Victor D. Luduvice Roberto Pinheiro

A Empirical Analysis: Data and Measurement

A.1 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.[1] While there were changes to NIPA Table 5.6, we start by 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:

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

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 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 (A.1) we assume only one lag, we follow a mixture of Akaike and Schwarz information criteria in order to select the optimal lag order 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 between computers and peripherals and other office and accounting machinery. In contrast, office, computing, and accounting machinery (OCAM) was a single category in Gordon (1990).

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 the NIPA index for the hedonic-based quality-adjusted price index

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

for computers and peripherals for the period 1959–2024. In order to mend these two series (which have different bases), we start by rebasing Gordon's quality-adjusted index such that the base year is 1958 (=100 in 1958). Then, we construct an adjustment factor $AdjFact = \frac{\text{Gordon}_{t=1958}}{\text{NIPA}_{t=1958}}$. We then 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 have also 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 either agricultural machinery or construction machinery in most cases, while in Gordon's quality-adjusted price index we do have that distinction. Finally, the 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 Törnqvist 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 $j \in \{1, \ldots, \mathcal{J}\}$ 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:

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

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,[2] pre-packaged, and custom software.[3] Following Cummins and Violante (2002), we use the NIPA's quality-adjusted software price index by itself.

A.2 AI Price Index and Capital

The AI price deflator is estimated from the hedonic regression (1) in Section 2, following the methodology described in Chapter 6 of Gordon (1990). Because our database records release

[2]Software developed internally by firms themselves.

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

dates and release prices, the regression is run on newly released processors and therefore traces the technology frontier, the preferred methodology according to Fisher et al. (1985). Gordon's Computers and Peripherals hedonic index is constructed on the same principle.

The database (Ai, 2025) 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. Only a subset of these processing units are commercialized and have price data. We combine these prices with Epoch AI's price data accessed through airtable and discussed by Pilz et al. (2025). In the end, we have price data on 61 processing units over the period starting in 2008, 56 starting in 2013.

The three quality characteristics we control for – *ML OP/s*, *memory size per board*, and thermal design power (*TDP*) – are defined in Section 2.[4] We use them not only because they are relevant for explaining variation in prices, but also because there are few missing observations in these variables, boosting our sample size.

We present the regression results for the full (pre-trimming) sample 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 processors, i.e., the ones above the 90th percentile.[5] We then re-estimate specification (3) on the trimmed sample. The published index is computed from this re-estimated regression, whose year-dummy coefficients therefore differ from the pre-trimming ones shown in 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.

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

[5]Gordon (1990) drops the processors in which the log-ratio is higher than 1.5 times the standard error. Given our small sample size, we adopt a more lenient threshold.

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

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

A.3 Additional Data Figures

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

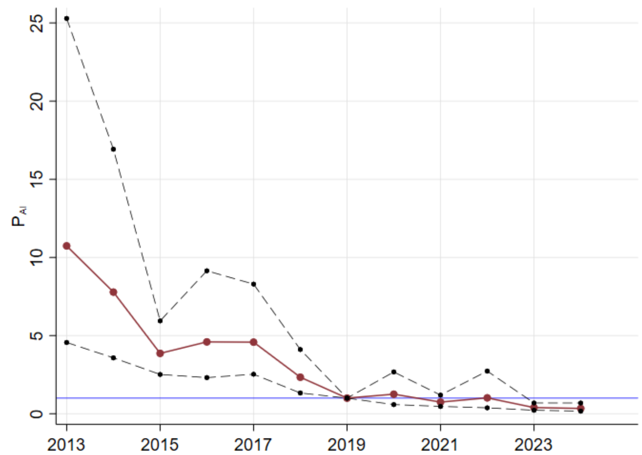

Notes: Quality-adjusted AI price from the hedonic regression in equation (1) and Epoch AI data for 2013–2024. Dashed lines: 95% confidence interval. Index normalized to 1 in 2019.

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

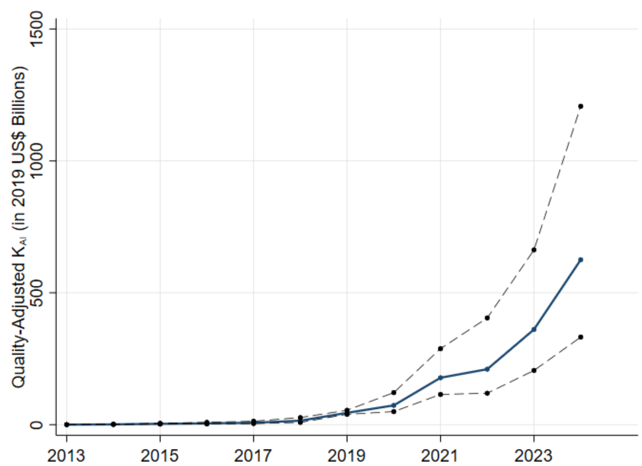

Notes: Quality-adjusted AI capital stock via the perpetual inventory method, based on the hedonic AI price index. Dashed lines: 95% confidence interval. Price index normalized to 1 in 2019.

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

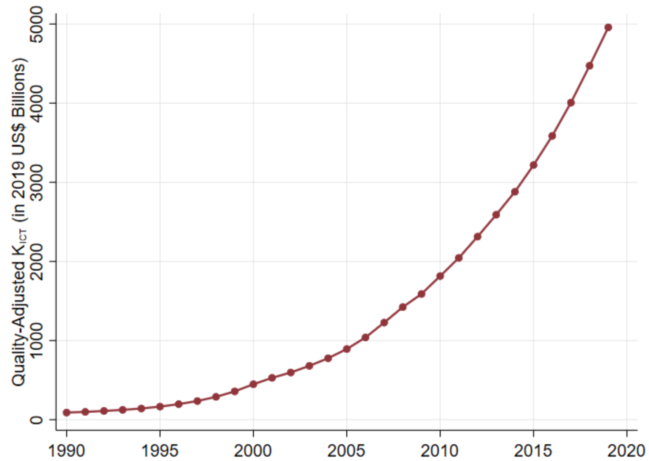

Notes: Quality-adjusted ICT capital stock. For details on the data construction, see Ohanian et al. (2023). Price index normalized to 1 in 2019.

B Model

B.1 Market Clearing Conditions

Product market clearing condition:

$$c_t + x_{st} + \frac{x_{et}}{q_t} + \frac{x_{AIt}}{\zeta_t} + G_t = A_t G(k_{s,t}, k_{e,t}, k_{AI,t}, m_{L,t}, m_{H,t})$$

using the capital laws of motion, we have:

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

Labor market clearing condition:

$$M_j L_{j,t} = m_{j,t}, \;\; \forall j \in \{L, H\} \text{ and } t \in \mathbb{N}$$

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

B.2 Steady State

We set constant technology parameters $A_t = A$, $q_t = q$, and $\zeta_t = \zeta$ for all $t \in \mathbb{N}$. Tax rates $(\tau_c, \tau_L, \tau_H, \tau_k)$ are constant, and the per-capita transfer T_{pc} and government purchases G are determined by the government budget constraint.

Then, from consumer $j \in \{L, H\}$, the FOCs give:

$$\begin{cases} \Lambda_j = \frac{1}{(1+\tau_c)\, C_j} & (FOC_{C_j}) \\ BL_j^{1/\vartheta} = \Lambda_j\, (1-\tau_j)\, w_j & (FOC_{L_j}) \\ (1+\tau_c)\, C_j = (1-\tau_j)\, w_j L_j + \mathcal{I}_j + T_{pc} & (BC_j) \end{cases}$$

where $\mathcal{I}_j$ denotes non-labor income: $\mathcal{I}_L = 0$ for low-skill workers and $\mathcal{I}_H = (1-\tau_k)\, \theta_H\, \Omega^{\text{net}}/M_H$ for high-skill workers.

Because the capital income tax applies to ownership income net of economic depreciation, the entrepreneur's Euler equations equate the after-tax net return per unit of capital *value* across the three capital types. Writing $P_s = 1$, $P_e = 1/q$, and $P_{AI} = 1/\zeta$ for the investment-good prices, the steady-state conditions are

$$\begin{cases} (1-\tau_k)\left(\dfrac{r_i}{P_i} - \delta_i\right) = \dfrac{1}{\beta} - 1, \qquad i \in \{s, e, AI\} & (FOC_{k_i}) \\ \Lambda_E = \frac{1}{(1+\tau_c)\, C_E} & (FOC_C) \end{cases}$$

Equivalently, the per-unit user costs – the rental rates r_i that the firm's first-order conditions must deliver, measured in consumption units per unit of physical capital – are

$$r_i = P_i\left(\delta_i + \frac{\frac{1}{\beta} - 1}{1-\tau_k}\right), \qquad i \in \{s, e, AI\}$$

so that $r_s = \delta_s + \frac{1/\beta - 1}{1-\tau_k}$, $r_e = \frac{1}{q}\left(\delta_e + \frac{1/\beta-1}{1-\tau_k}\right)$, and $r_{AI} = \frac{1}{\zeta}\left(\delta_{AI} + \frac{1/\beta-1}{1-\tau_k}\right)$. It is convenient to define the net rental rate $r_i^{\text{net}} \equiv r_i/P_i - \delta_i$, which by the conditions above is common to all three capital types and equal to $(1/\beta - 1)/(1-\tau_k)$. Faster-depreciating capital therefore commands a higher *gross* rental rate per unit of capital value, but not a higher net return.

Market clearing conditions give us:

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

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.

We now solve for steady-state labor supply. Combining (FOC_{C_j}) and (FOC_{L_j}) yields the intratemporal condition:

$$BL_j^{1/\vartheta}\left(1+\tau_c\right)C_j = \left(1-\tau_j\right)w_j$$

Unlike the no-tax case, progressive taxation $(\tau_H > \tau_L)$ implies that $L_L \neq L_H$ in general, because the after-tax wage rate differs across skill types even after controlling for the level of consumption. Specifically, from the intratemporal condition we have:

$$L_j = \left(\frac{\left(1-\tau_j\right)w_j}{B\left(1+\tau_c\right)C_j}\right)^{\vartheta}$$

From the entrepreneur's Euler equations we have the user-cost conditions above, together with $C_E = 1/\left(\left(1+\tau_c\right)\Lambda_E\right)$. Denote by $\tilde{G}$ the skill-bundle aggregate of the production function, that is, the braced term in the market clearing condition above:

$$\tilde{G} \equiv \left[A_L\left(m_L^{\gamma_L} + k_e^{\gamma_L}\right)^{\frac{\lambda_L}{\gamma_L}} + \Psi_{AI}k_{AI}^{\lambda_L}\right]^{\frac{\phi}{\lambda_L}} + \left[A_H\left(m_H^{\gamma_H} + k_e^{\gamma_H}\right)^{\frac{\lambda_H}{\gamma_H}} + \Psi_{AI}k_{AI}^{\lambda_H}\right]^{\frac{\phi}{\lambda_H}}$$

Substituting into the goods-market clearing condition yields:

$$\left\{\begin{array}{c} M_L C_L + M_H C_H + M_E C_E + G \\ +\delta_s\, k_s + \frac{\delta_e}{q}\, k_e + \frac{\delta_{AI}}{\zeta}\, k_{AI} \end{array}\right\} = Ak_s^{\alpha_1}\tilde{G}^{\frac{\alpha_2}{\phi}} \tag{A.3}$$

The steady-state allocation is pinned down by the firm's first-order conditions, the worker and entrepreneur optimality conditions, the aggregate resource constraint in (A.3), and the government budget constraint.

C Calibration and Estimation

C.1 GMM Implementation Details

Sample. The estimation uses annual U.S. data for 1963–2024. The moment conditions involve one-period leads and the instrument set one-period lags, so the first and last years drop out, leaving 60 usable observations. The Krusell et al. (2000) moments – the wage-bill ratio, the labor share, and the equipment vs. structures no-arbitrage condition – are evaluated over all of them, as are the Katz–Murphy relative-demand condition, the output level, and the capital-demand conditions.

The two AI-specific moments – the AI vs. structures no-arbitrage condition and the AI user-cost condition – enter the objective only from the AI start year onward, which is 2013 in the baseline estimation. From that year the equipment series excludes computers and peripherals so that the AI proxy is not counted twice (see Appendix C.3).

Orthogonality conditions and degrees of freedom. Each moment condition is interacted with the subset of instruments suited to it, with instruments standardized to mean zero and unit variance: every condition carries the constant; the four Krusell et al. (2000)-style conditions additionally carry the structures stock and the lagged equipment price; the output level carries the lagged equipment price; the Katz–Murphy condition is instrumented by relative skill supply; and the two gated AI conditions additionally carry the hedonic AI log price of Section 2. This yields 22 orthogonality conditions for 7 estimated parameters – with (α_1, α_2) counted as calibrated, not estimated – so the overidentification test has $22 - 7 = 15$ degrees of freedom. As several instruments have strong trends, the asymptotic reference distributions for the standard errors and the J-statistic rest on stationarity conditions that are difficult to verify in this sample. We therefore read the J-test as a specification diagnostic and complement it with the profile and bootstrap evidence below.

Weighting and inference. We use two-step GMM. The first step minimizes the objective under an identity weighting matrix and the second step uses the inverse of a heteroskedasticity- and autocorrelation-consistent (Newey–West, 3 lags, uncentered) estimate of the moment covariance evaluated at the first-step solution, inverted directly without regularization. Standard errors are computed from the usual sandwich formula, with the Jacobian of the moment conditions obtained by numerical differentiation. The J-statistic is the sample size times the minimized second-step objective with distribution $\chi^2(15)$ under the null hypothesis. The reported standard errors are conditional on the calibrated (α_1, α_2); redrawing the calibration inputs within their sampling uncertainty widens the standard error of λ_H by roughly a factor of four, and the block bootstrap below, which re-calibrates α_1 on every resampled dataset, measures the same total uncertainty directly.

Optimization and parameter space. The objective is minimized with a sequential-quadratic-programming routine from 48 dispersed starting points – twelve in each sign quadrant of (γ_H, λ_H), with the quadrants enforced through the bounds so no run can drift across the excluded band – and the quadrant solutions are compared under a common weighting matrix. The parameter space imposes wide box bounds: $\phi \in [0.05, 1.60]$, $\mu_H \in [0.50, 1.00]$, $\xi_H \in [0.05, 0.999]$, $\gamma_H \in [-4.50, 0.90]$, and $\lambda_H \in [-1.50, 0.90]$; (α_1, α_2) are pinned at their calibrated values. In addition, we exclude a small neighborhood of zero for λ_H, $|\lambda_H| \geq 0.05$: at $\lambda_H = 0$ the Cobb–Douglas limit of the inner nest is itself finite, but the implementation evaluates exponents of the form ϕ/λ_H, so the objective is numerically ill-behaved in a neighborhood of zero. The parameter space admits substitutability as well as complementarity. The point estimate $\lambda_H = -0.12$ is interior to every constraint. Single-moment ablations locate where the sign is carried: dropping any one of the AI-specific conditions leaves λ_H essentially unchanged, while dropping the wage-bill or labor-share condition leaves the optimizer pinned at the excluded-band edge, so the complementarity is identified by the AI conditions jointly with the labor-market moments rather than by any single condition.

The calibrated exponents. The structures exponent is calibrated inside the estimation as the sample mean of the structures income share evaluated at the externally declared required return $\bar{\rho}$ from **?**, with $\alpha_2 = 1 - \alpha_1 - \bar{\pi}$ and the profit share $\bar{\pi}$ declared at its NIPA value, so returns to scale are measured rather than estimated. The recipe is not reverse-engineered: applied to the arm of the data that uses the textbook structures depreciation rate, the same formula delivers $\alpha_1 = 0.115$, within two percent of the value of 0.117 that Krusell et al. (2000) calibrate from the same income-share logic. The baseline arm, which carries the BEA structures depreciation series, delivers $\alpha_1 = 0.094$.

Profile of the objective function around λ_H. Figure A.4 traces the GMM objective as a function of λ_H: at each of 24 grid points spanning $[-0.50, +0.50]$ we fix λ_H and re-estimate the remaining parameters. The figure plots the first-step objective, computed under one common identity weighting matrix and therefore directly comparable across grid points; the two-step objective, whose weighting matrix re-optimizes at every point, is nearly flat by construction and carries no visual information. The comparable profile has a clear interior minimum at $\lambda_H = -0.115$, immediately adjacent to the unconstrained estimate of -0.12 and far from both the box bounds and the excluded neighborhood of zero, with a steep penalty on the far-complements side. Without the band, the optimizer converges to exactly the baseline estimate – $\lambda_H = -0.118$, $\gamma_H = -3.17$, $J = 15.3$ – so the exclusion band never binds at the optimum.

Figure A.4: Profile of the GMM objective in λ_H

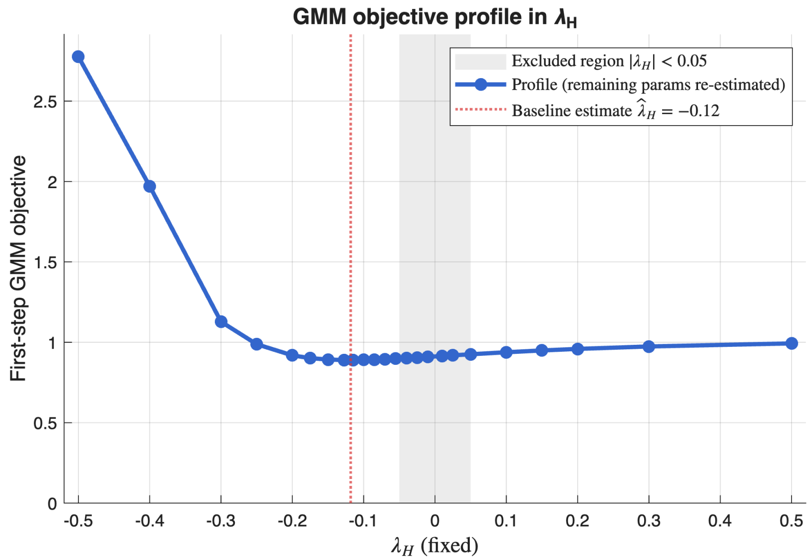

Notes: First-step GMM objective (common identity weighting matrix, directly comparable across grid points) with λ_H fixed at each grid point and the remaining parameters re-estimated. The shaded band marks the excluded neighborhood of the Cobb–Douglas singularity, $|\lambda_H| < 0.05$. The dotted line marks the unconstrained baseline estimate $\hat{\lambda}_H = -0.12$.

C.2 Block Bootstrap of the GMM Estimates

Due to the short time series, we assess the sampling uncertainty of the GMM estimates with a block bootstrap that preserves the serial dependence of the data. We draw 500 bootstrap samples using overlapping blocks of length 5 and re-run the full estimation recipe on each sample, including the in-run calibration of α_1 from the resampled data, so the resulting intervals measure total – calibration plus estimation – uncertainty. Our reported intervals use a time-ordered scheme: the resampled blocks are re-ordered so that observations appear in their original calendar order and are then assigned synthetic sequential years. The moment conditions exploit the monotone time-series structure of the data, in particular the trending quality-adjusted AI price and capital series and the lagged instruments k_{t-1} and $P_{e,t-1}$. Schemes that concatenate blocks in random order do not preserve those trends and lead–lag relationships. The cost is that the time-ordered resamples are closer to jittered subsamples than to independent redraws, so we also report two standard schemes below as a stress test.

In every scheme, the gating of the two AI-specific moments follows each resampled observation's original calendar year rather than the synthetic year it is assigned, so a replication activates the AI moments on exactly the observations that were drawn from 2013 onward. Under the time-ordered scheme, all 500 replications converge and 408 yield interior solutions for (λ_H, γ_H). The remainder end near a constraint, reflecting the limited information in the short AI sample. Table A.2 reports point estimates alongside asymptotic 95 percent confidence intervals and bootstrap percentile intervals for the two complementarity parameters and their implied elasticities.

Table A.2: Block bootstrap: asymptotic confidence intervals and percentile intervals

Parameter	**Point Est.**	**Asymptotic 95% CI**	**Bootstrap percentile interval**
λ_H	-0.118	$[-0.138,\ -0.098]$	$[-0.315,\ -0.050]$
γ_H	-3.174	$[-3.365,\ -2.983]$	$[-4.295,\ -1.932]$
$\sigma_\lambda = \frac{1}{1-\lambda_H}$	0.894	$[0.879,\ 0.910]$	$[0.761,\ 0.952]$
$\sigma_\gamma = \frac{1}{1-\gamma_H}$	0.240	$[0.229,\ 0.251]$	$[0.189,\ 0.341]$

As a stress test, we re-run the bootstrap under two standard schemes that do not restore the time ordering: a moving-block bootstrap, which concatenates the sampled blocks in the order drawn, and a circular-block bootstrap, which additionally allows blocks to wrap around the end of the sample. Table A.3 compares the three schemes. As expected, scrambling the block order destroys the trends that the moment conditions exploit. The share of interior solutions falls from 82 percent to about 19–38 percent under the standard schemes, with λ_H splitting between the -0.05 edge of the excluded Cobb–Douglas band and the -1.50 box bound. Percentile intervals computed over all replications therefore degenerate toward the width of the parameter space itself.

Table A.3: Block bootstrap: comparison across resampling schemes

Scheme	**Interior**	**λ_H 95% CI (interior)**	**σ_λ 95% CI (interior)**	**$\lambda_H > 0$**
Time-ordered	408/500	$[-0.317,\ -0.065]$	$[0.759,\ 0.939]$	0/500
Moving block	93/500	$[-1.435,\ -0.093]$	$[0.411,\ 0.915]$	0/500
Circular block	191/500	$[-1.471,\ -0.192]$	$[0.405,\ 0.839]$	0/500

Notes: 500 replications per scheme, block length 5. "Interior" counts replications in which (λ_H, γ_H) is away from every constraint. Percentile confidence intervals are computed over interior replications. Over all replications, the moving- and circular-block intervals degenerate to the bound range $[-1.50, -0.05]$. The last column counts replications in the substitutes region.

In none of the 1,500 replications across the three schemes does λ_H enter the substitutes region, so σ_λ remains below one in every replication; γ_H crosses zero in a single replication under the reported scheme, and only in a minority of the scrambled-scheme replications whose intervals are degenerate anyway.[6] Table A.2 reports the time-ordered intervals. Those intervals are somewhat wider than the asymptotic ones, particularly for γ_H, but they remain entirely within the complementarity region.

[6]The optimizer is started from an interior point in the complements branch, so the local search does not explore the positive-λ_H branch directly, and the count should be read as conditional on that branch. The profile in Figure A.4, which does cover that branch, shows that it fits the data substantially worse.

C.3 Avoiding double counting of AI capital

In our estimation, we treat computers and peripherals as the empirical proxy for AI capital, K_{AI}. Thus, the same assets must be removed from the equipment aggregate K_E over the period in which they are counted as AI. We therefore construct, for each AI start year, an equipment series that excludes computers and peripherals from that year onward, rebuilding the no-computer equipment stock from the remaining detailed asset categories and splicing it onto the standard (mixed) equipment series at the cutoff. The corresponding no-computer price index is a Törnqvist aggregate over those categories, with expenditure shares computed within the no-computer basket.

C.4 AI Start-Year Robustness

As we work with a short AI capital series, the baseline gates the AI-specific moments from 2013, the start of the AI-specific era in our hedonic price data. Here we re-estimate the full parameter vector while instead gating these moments from each of the alternative start years $\{1995, 2000, 2005\}$. For each start year we use that year's own equipment series, constructed to exclude computers and peripherals from the corresponding cutoff.[7] Table A.4 reports the implied elasticities and the overidentification test.

Table A.4: AI start-year robustness: implied elasticities and fit

AI start year	σ_λ	σ_γ	J	p-value
2013 (baseline)	0.894	0.240	15.3	0.427
2000	0.924	0.238	15.3	0.431
2005	0.902	0.233	15.5	0.419
1995	0.894	0.270	15.3	0.433

Notes: Two-step IV-GMM re-estimated with the AI-specific moments gated from each start year, each using its own equipment series excluding computers and peripherals from the corresponding cutoff (15 degrees of freedom). $\sigma_\lambda = 1/(1-\lambda_H)$ and $\sigma_\gamma = 1/(1-\gamma_H)$.

The AI–skill complementarity is stable across start years: σ_λ ranges narrowly over $[0.89, 0.92]$ and stays below one for every start year, as does σ_γ. The overidentifying restrictions are satisfied at every start year. At the 2000 gate the point estimate sits closest to the excluded band – within two standard errors of its edge – so that specification pins down the sign of the complementarity more firmly than its magnitude.

Dropping the wage-bill or the labor-share moment leaves λ_H pinned at the edge of the excluded band and the AI–skill elasticity unidentified, while dropping any single AI-specific condition

[7] For each year, we rebuild the price index by Törnqvist aggregation, with the no-computer price chained onto the mixed level at the cutoff.

moves the estimate little: the complementarity is identified by the AI conditions jointly with the labor-market moments.

D Quantitative Results

D.1 Transitory TFP Shock

Figure A.5: Transitory impulse response: 1% TFP shock ($\rho = 0.90$)

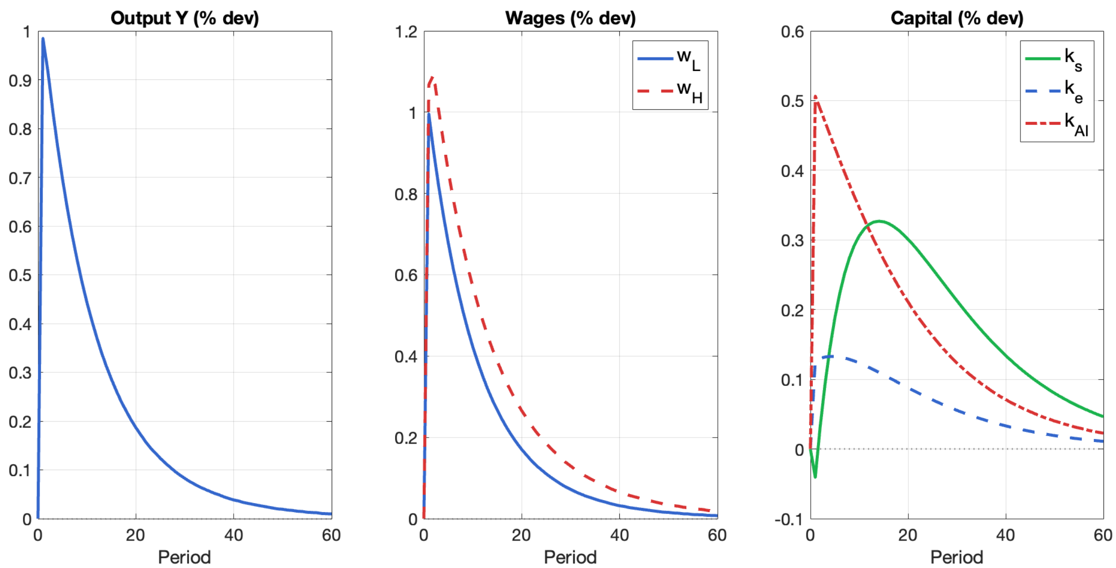

Notes: Perfect-foresight impulse responses (percent deviation from the initial steady state) to a transitory 1% increase in neutral productivity A_t, following an AR(1) process with persistence $\rho = 0.90$.

D.2 Permanent-Shock Transition Dynamics

Figure A.6: Transition dynamics: Capital Stocks by Type

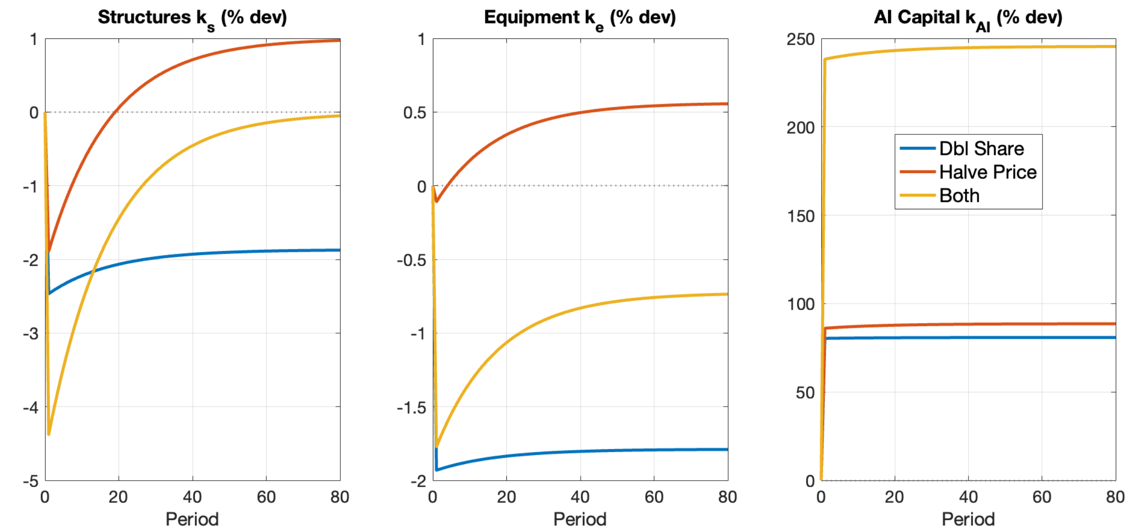

Notes: Perfect-foresight transition paths of the three capital stocks – structures k_s, equipment k_e, and AI capital k_{AI} – in percent deviation from the initial steady state, for the three permanent AI experiments.

Figure A.7: Transition dynamics: Consumption by Agent Type

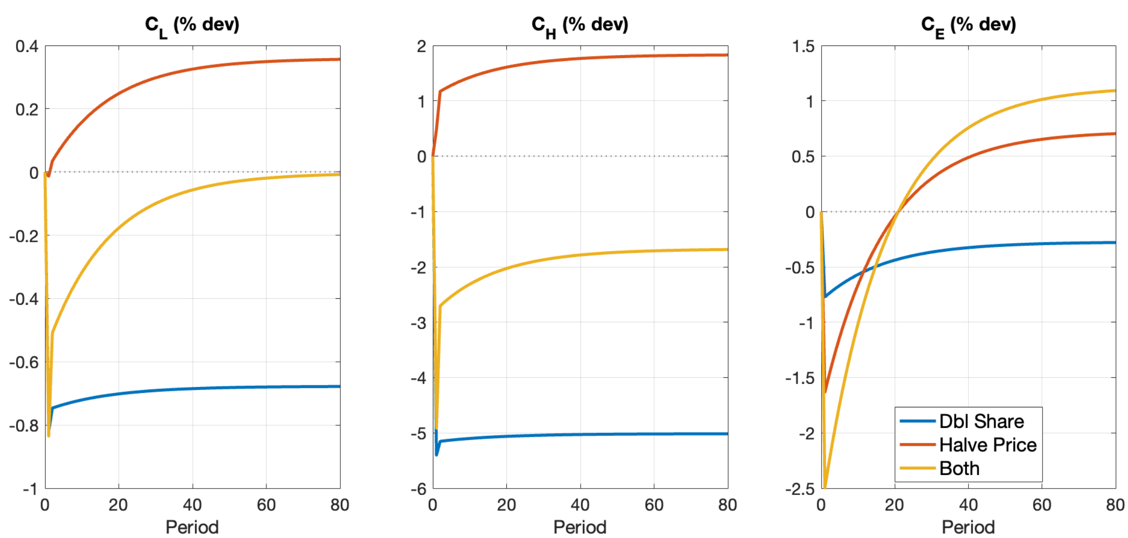

Notes: Perfect-foresight transition paths of consumption by agent type – low-skill C_L, high-skill C_H, and entrepreneur C_E – in percent deviation from the initial steady state, for the three permanent AI experiments.

Figure A.8: Transition dynamics: Skill Premium and Output

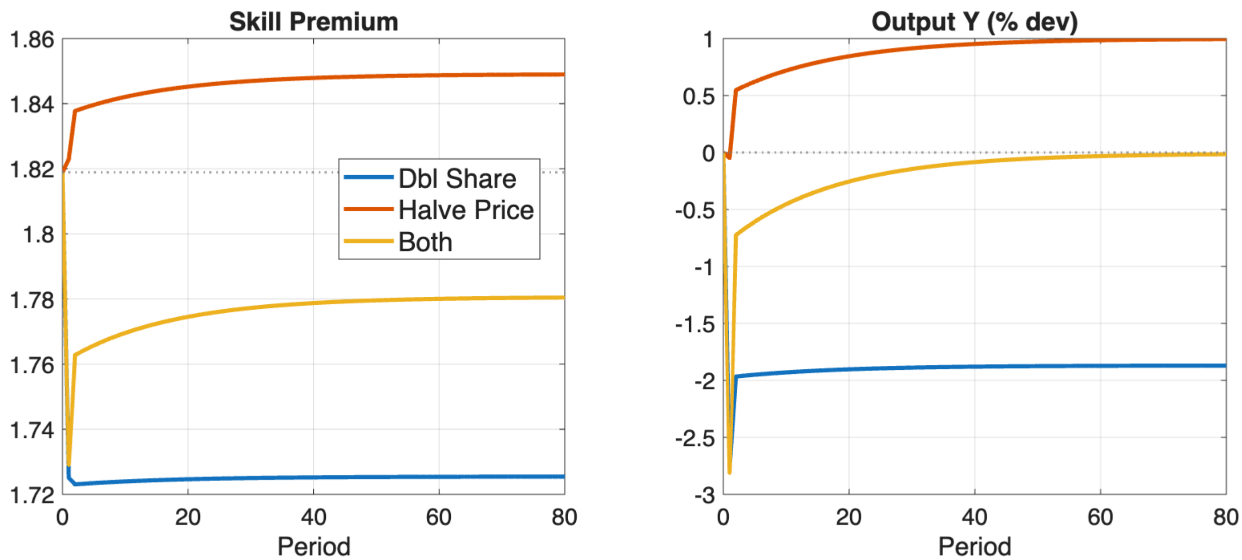

Notes: Perfect-foresight transition paths of the skill premium w_H/w_L (in levels) and output (percent deviation from the initial steady state) for the three permanent AI experiments.

D.3 Individual Experiments: Level Figures

Figure A.9: Transition dynamics (levels): Output, Consumption, and Aggregate Capital

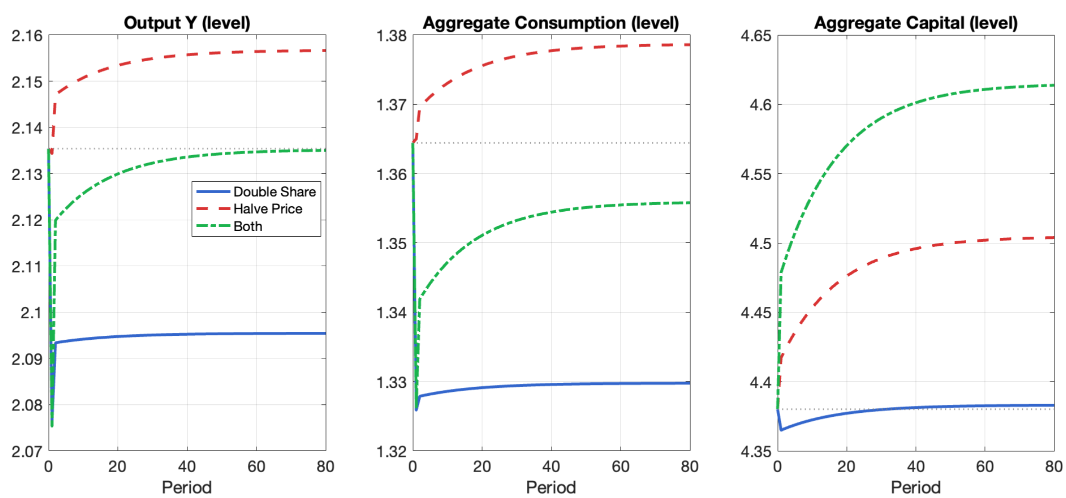

Notes: Transition paths of output, aggregate consumption, and the aggregate capital stock in levels for the three permanent AI experiments. The percent-deviation counterparts are Figures A.8, A.7, and A.6.

Figure A.10: Transition dynamics (levels): Capital Stocks by Type

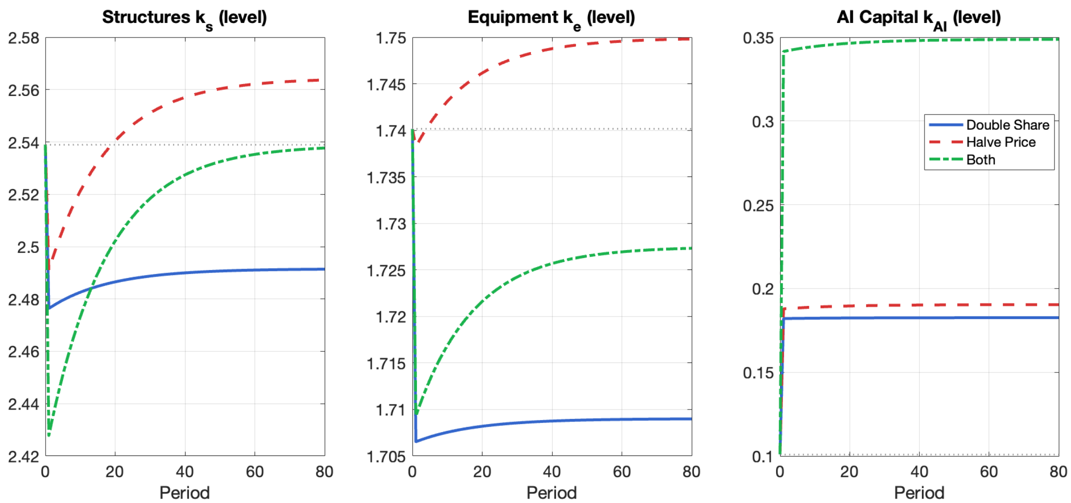

Notes: Transition paths of the three capital stocks in levels for the three permanent AI experiments. The percent-deviation counterpart is Figure A.6.

Figure A.11: Transition dynamics (levels): Consumption by Agent Type

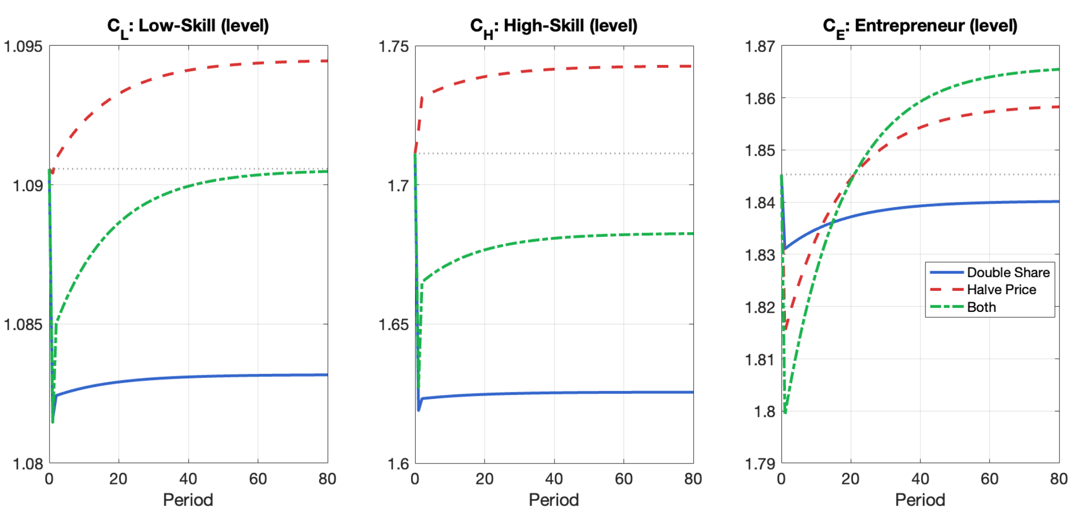

Notes: Transition paths of consumption by agent type in levels for the three permanent AI experiments. The percent-deviation counterpart is Figure A.7.

Figure A.12: Transition dynamics (levels): Wages

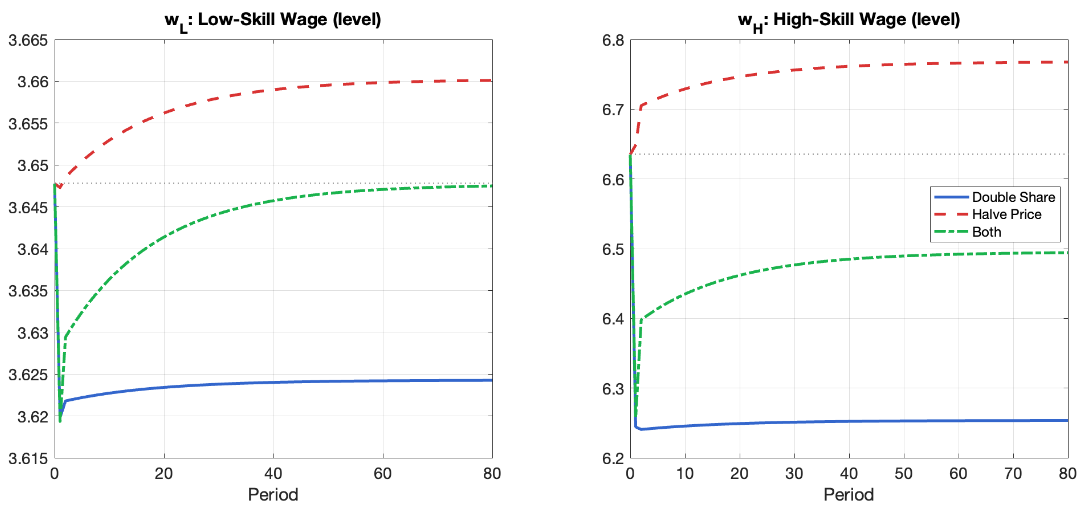

Notes: Transition paths of low- and high-skill wages in levels for the three permanent AI experiments. The skill-premium counterpart in percent deviations is Figure A.8.

D.4 Combined Experiment: Additional Figures

Figure A.13: Combined experiment: Consumption by Agent Type (levels)

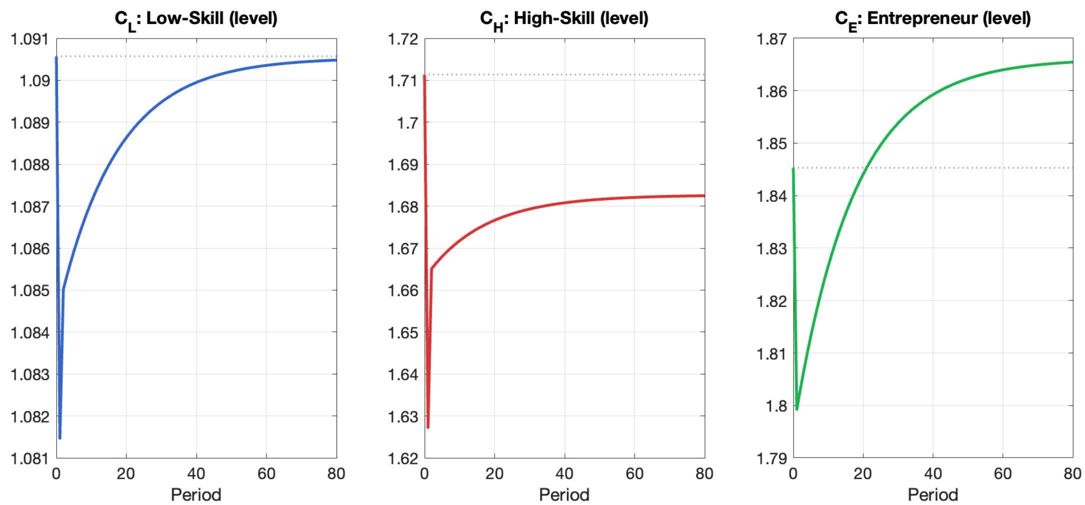

Notes: Transition paths of consumption by agent type in levels under the combined AI shock.

Figure A.14: Combined experiment: Income Distribution

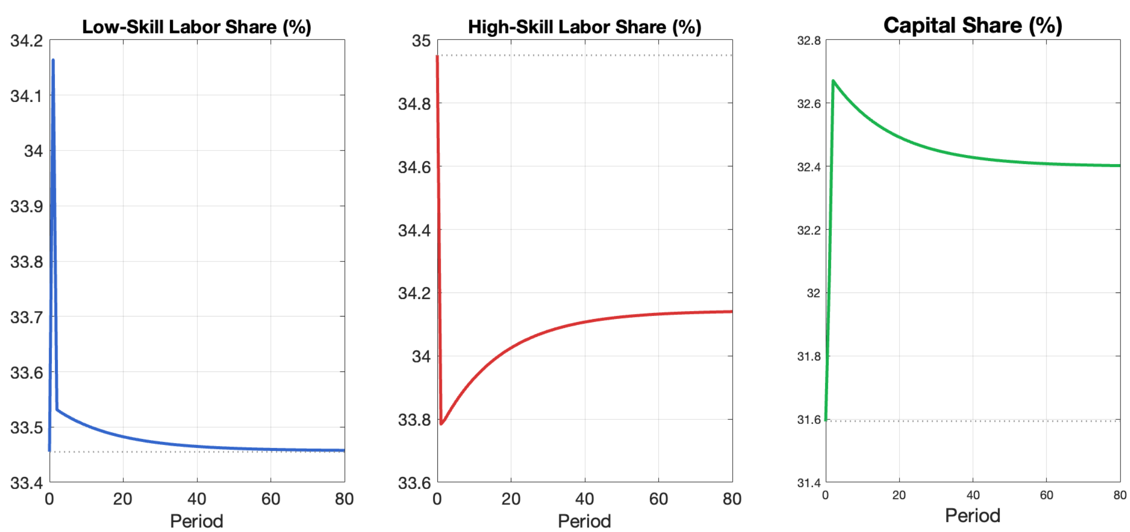

Notes: Functional income shares – low-skill labor, high-skill labor, and capital (rental income plus profits) – along the perfect-foresight transition under the combined AI shock. High-skill workers also receive a portion of ownership income through θ_H, hence these factor shares do not map one-for-one into the incomes of the three agent groups.

D.5 Robustness

D.5.1 Varying σ_λ

To assess the sensitivity of our results to the estimated complementarity between AI and the high-skill–equipment composite, we vary λ_H to produce three alternative values of σ_λ: the baseline estimate (0.89, complements), moderate substitutes (1.20), and strong substitutes (1.49). Table A.5 reports the results.

Table A.5: Robustness: Results across σ_λ

	$\sigma_\lambda = 0.89$ (complements)	$\sigma_\lambda = 1.20$ (mod. subs.)	$\sigma_\lambda = 1.49$ (strong subs.)
Output Y (% Δ)			
Double Share	−1.87	−2.04	−2.08
Halve Price	+1.00	+0.60	+0.36
Both	−0.01	−0.66	−1.07
Super-additivity (pp)			
Both − (Exp1+Exp2)	+0.86	+0.77	+0.64
Skill Premium w_H/w_L (level)			
Initial SS	1.819	1.864	1.888
Both	1.781	1.809	1.829

Notes: The first column reproduces the baseline calibration, so it matches Tables 5 and 6 exactly. The two substitutes columns re-solve the steady state at the indicated λ_H, which is why the initial skill premium differs across columns.

The share increase is always contractionary, and the contraction deepens with higher substitutability. With higher σ_λ, AI substitutes more easily for the high-skill–equipment composite, so the economy loses more from reweighting away from that composite. The price drop is always expansionary, but less so 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 replaces it instead of amplifying it.

The level of the combined effect is not robust across the grid: output is essentially flat at the estimated elasticity and falls increasingly as σ_λ rises. Whether the combined shock leaves the economy unchanged or contracts it therefore depends on the elasticity. The interaction gain does not inherit that sensitivity. It stays between 0.64 and 0.86 percentage points across the grid, declining only mildly in σ_λ, and remains positive where AI and the high-skill–equipment composite are substitutes. Figure A.15 displays the pattern.

Figure A.15: Super-additivity across σ_λ

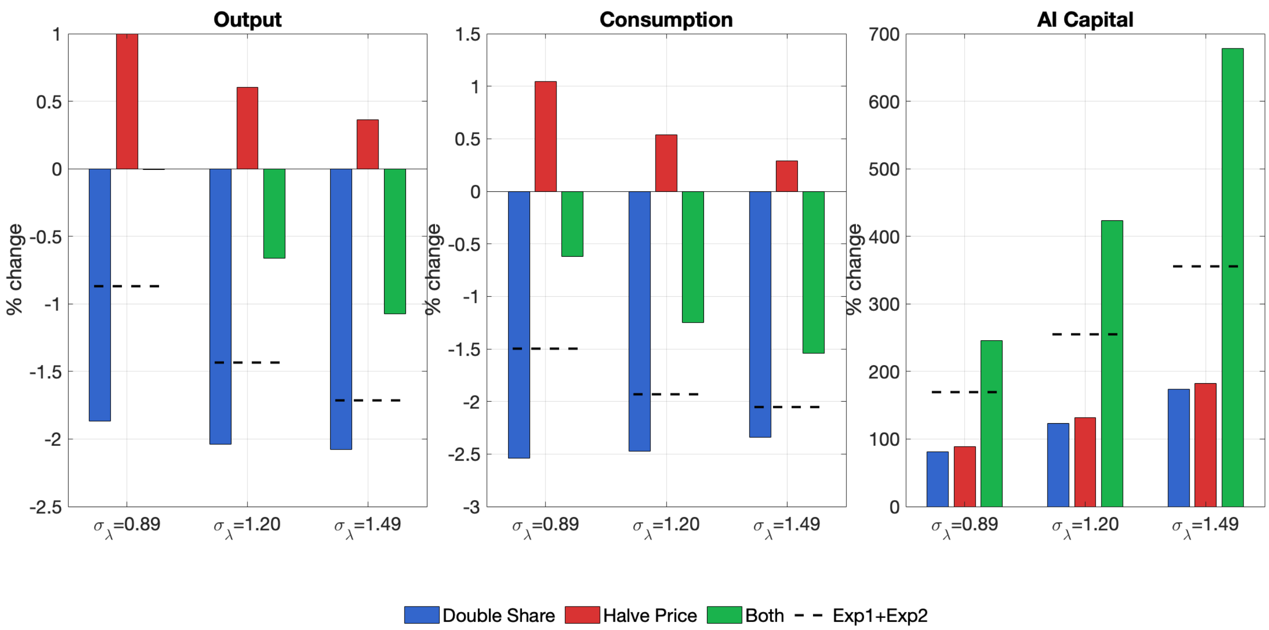

Notes: Super-additivity of the combined AI shock across three values of the AI–skill elasticity σ_λ (baseline, moderate substitutes, strong substitutes. Super-additivity is defined as the gain from both shocks together minus the sum of the two individual output effects.

Finally, the skill premium moves only modestly across all three σ_λ values in the combined experiment, as the share-increase and price-drop effects on high-skill wages systematically offset each other. Appendix Figure A.17 shows the consumption dynamics across the three elasticity values.

D.5.2 Varying δ_{AI}

We also assess the sensitivity of our results to δ_{AI}. Table A.6 reports the steady-state effects for three values: $\delta_{AI} = 0.10$ (low depreciation), $\delta_{AI} = 0.156$ (baseline), and $\delta_{AI} = 0.25$ (high depreciation).

Table A.6: Robustness: Results across δ_{AI}

	$\delta_{AI} = 0.10$ (low)	$\delta_{AI} = 0.156$ (baseline)	$\delta_{AI} = 0.25$ (high)
Output Y (% Δ)			
Double Share	−1.48	−1.87	−2.34
Halve Price	+0.97	+1.00	+1.04
Both	+0.33	−0.01	−0.42
Skill Premium w_H/w_L (level)			
Initial SS	1.833	1.819	1.802
Both	1.805	1.781	1.751
Super-additivity (pp)			
Both − (Exp1+Exp2)	+0.84	+0.86	+0.88

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. More depreciation also slightly increases the expansionary effect of the price reduction, since the user cost $r_{AI} = \left(\delta_{AI} + \frac{1/\beta - 1}{1-\tau_k}\right)/\zeta$ falls by more in absolute terms when δ_{AI} is high and ζ doubles. Similarly to the σ_λ robustness, the combined output effect is the sensitive object – ranging from +0.33 percent at $\delta_{AI} = 0.10$ down to −0.42 percent at $\delta_{AI} = 0.25$ – while super-additivity is nearly constant across all three values.

Figure A.16 displays the super-additivity pattern across δ_{AI} values.

D.5.3 Robustness: Additional Figures

E Policy Experiments

E.1 Revenue and Iso-Revenue Surface

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

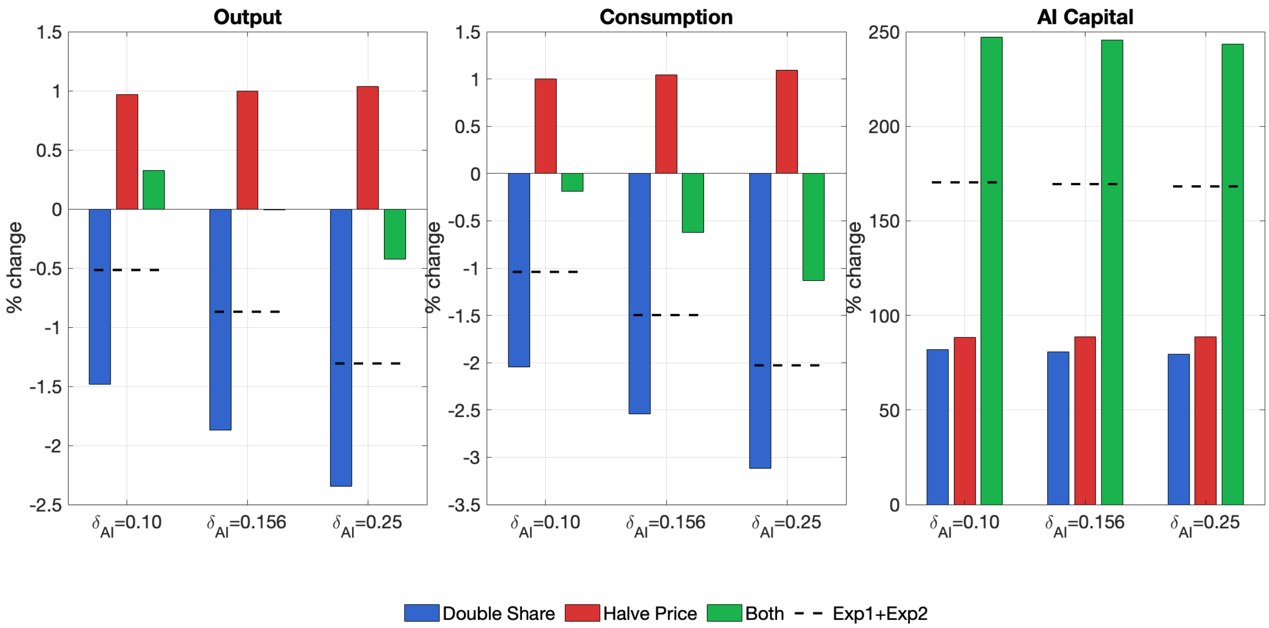

Notes: Super-additivity across three values of the AI depreciation rate δ_{AI} (0.10, 0.156, 0.25). Super-additivity is defined as the output from both shocks together minus the sum of the two individual output effects.

Figure A.17: Consumption across σ_λ

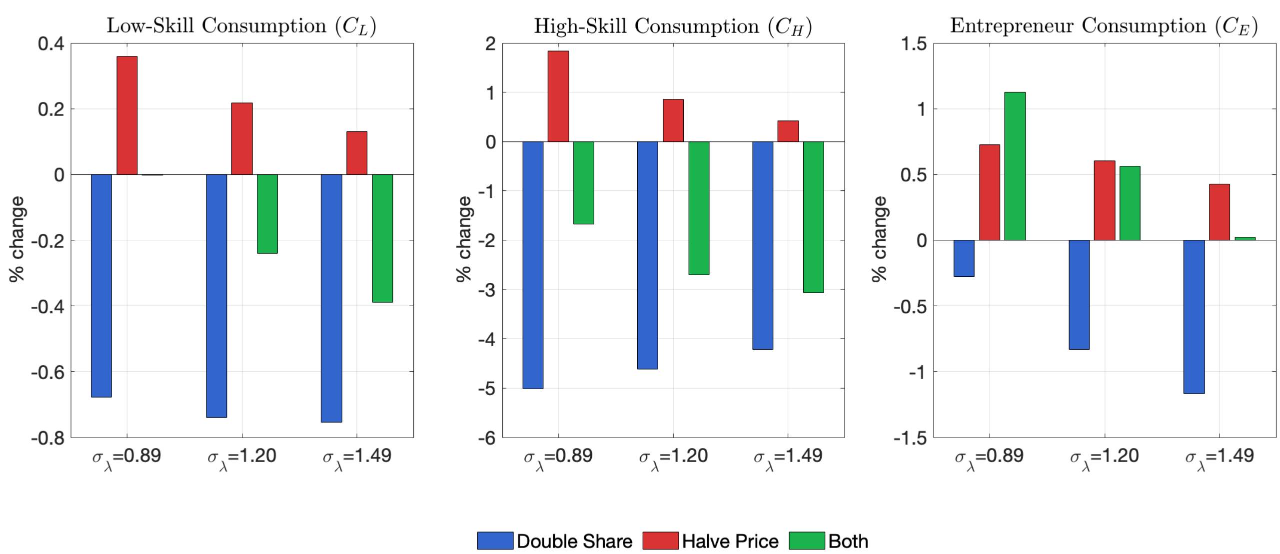

Notes: New steady-state consumption by agent type for each of the three permanent AI experiments (Double Share, Halve Price, Both), across three values of the AI–skill elasticity σ_λ (baseline, moderate substitutes, strong substitutes). Bars are percent changes from the initial steady state.

Figure A.18: Robustness: Skill Premium across σ_λ

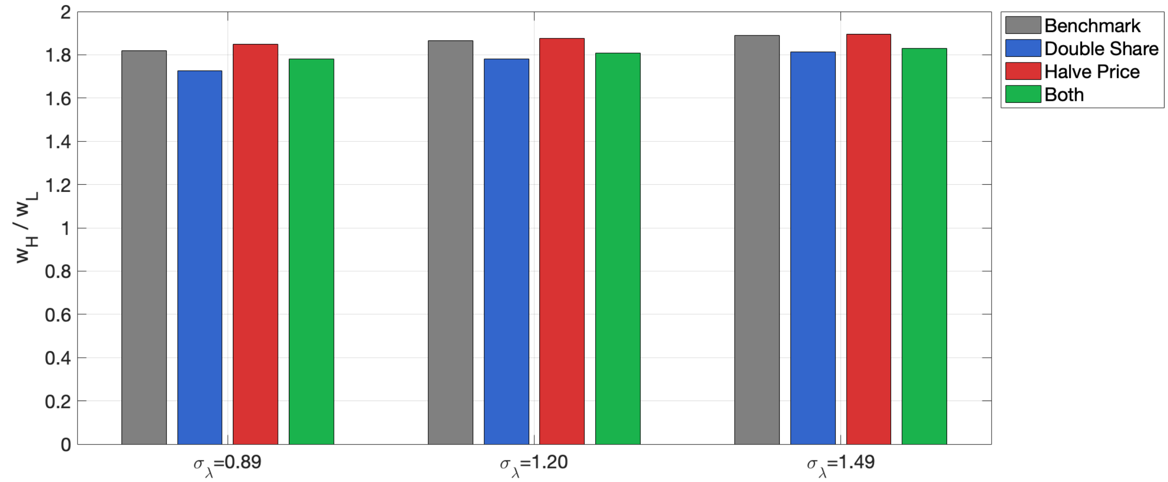

Notes: Skill premium w_H/w_L (in levels) across the three permanent AI experiments and the initial steady state, for three values of the AI–skill elasticity σ_λ (baseline, moderate substitutes, strong substitutes).

Figure A.19: Robustness: Wages across σ_λ

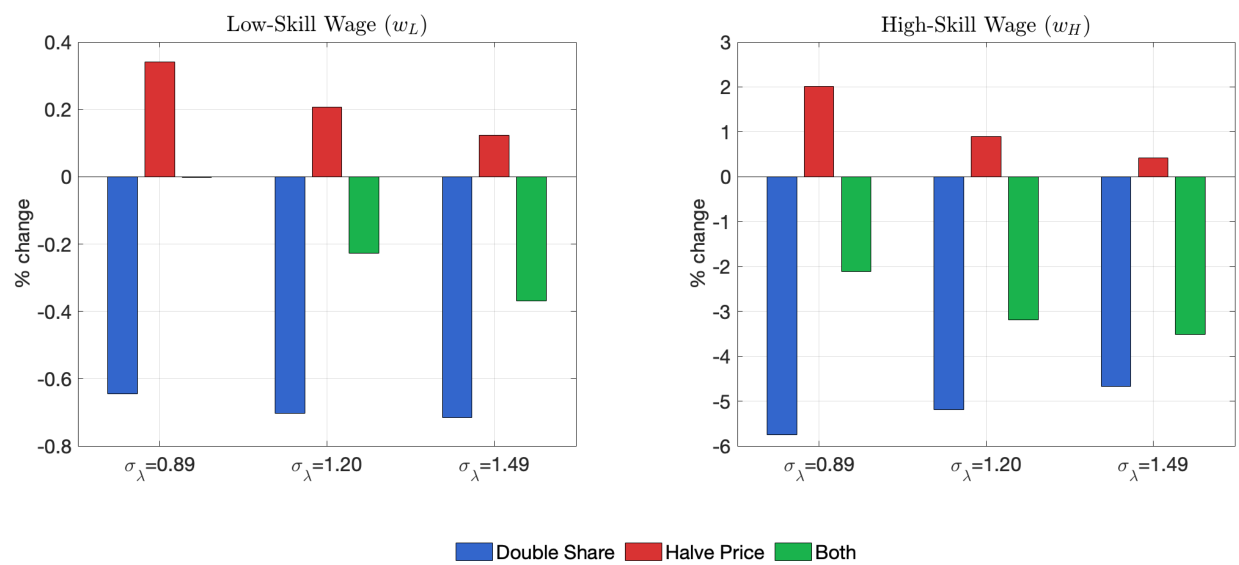

Notes: Low- and high-skill wages across the three permanent AI experiments and the initial steady state, for three values of the AI–skill elasticity σ_λ (baseline, moderate substitutes, strong substitutes).

Figure A.20: AI capital income tax: revenue Laffer curve

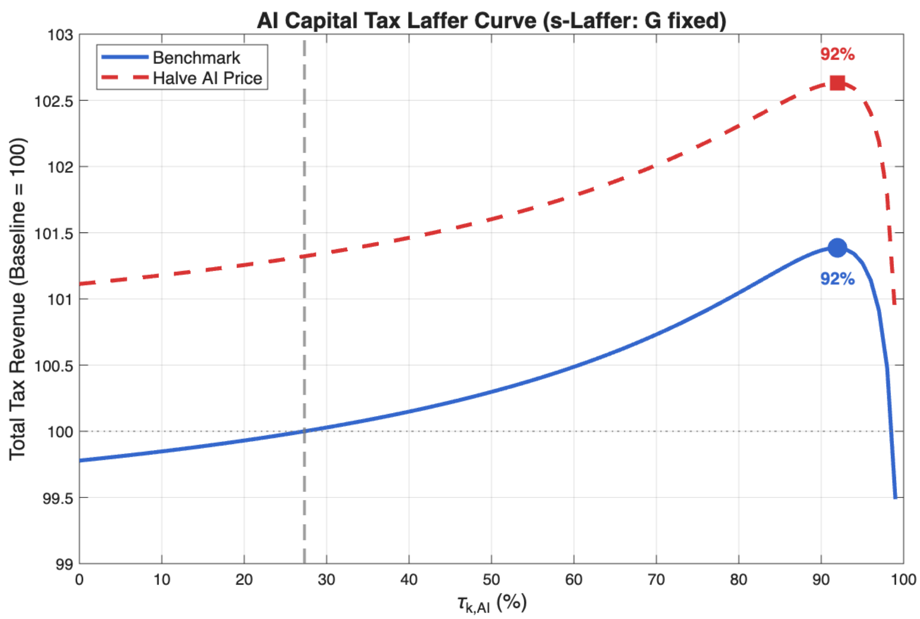

Notes: Total tax revenue (baseline = 100) as a function of the AI capital income tax $\tau_{k,AI}$, with government purchases G held fixed. Solid: baseline economy (labeled "Benchmark" in the panel). Dashed: halved AI price. Revenue peaks near $\tau_{k,AI} \approx 92\%$ in both economies.

Figure A.21: Iso-revenue surface and the Laffer ridge (halved AI price)

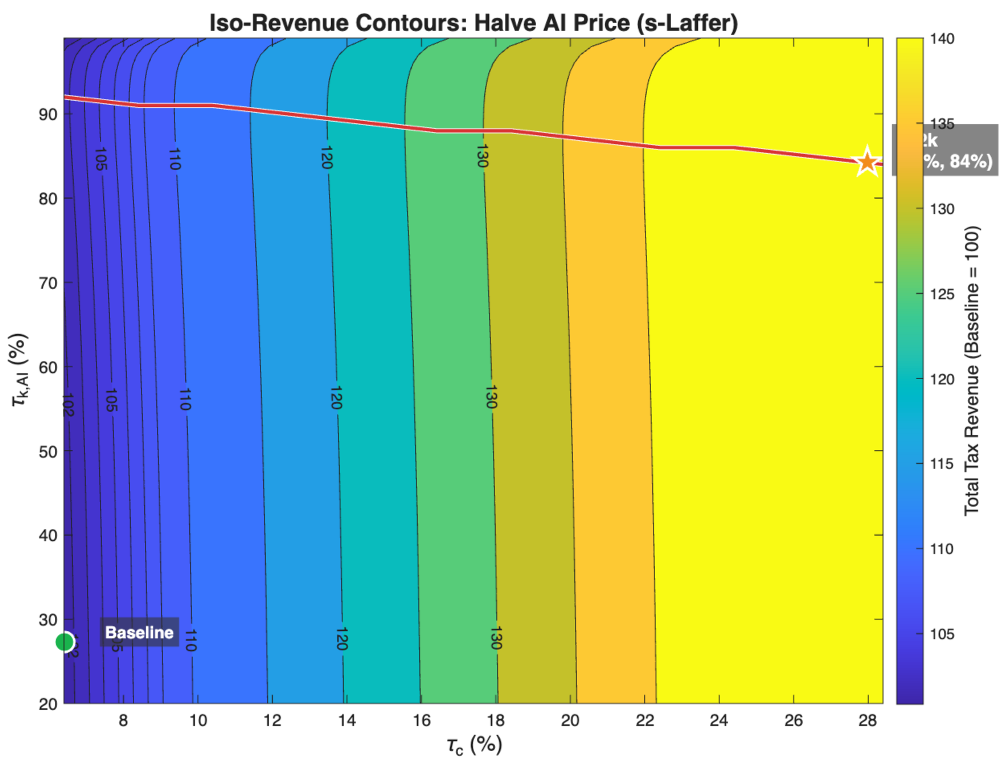

Notes: Total tax revenue (baseline = 100) over combinations of the consumption tax τ_c and the AI capital income tax $\tau_{k,AI}$, under the halved AI price. The red line is the Laffer ridge – the revenue-maximizing $\tau_{k,AI}$ at each τ_c. The star marks the Policy (iii) allocation that funds the \$12,000 transfer, $(\tau_c, \tau_{k,AI}) \approx (28\%, 84\%)$, the dot marks the baseline.