AI and Human Capital Accumulation: Aggregate and Distributional Implications*

Yang K. Lu† HKUST

Eunseong Ma‡ Yonsei

June 2026

Abstract

This paper investigates how human capital responses to anticipated advances in artificial intelligence (AI) reshape aggregate and distributional consequences of AI. We develop an incomplete-markets model with endogenous human capital and asset accumulation in general equilibrium, featuring three skill sectors and uninsurable idiosyncratic risk. AI enters as an anticipated, sector-biased shock that narrows middle-skill wage premiums and boosts returns to top expertise. We find that human capital responses to AI (i) drive *voluntary job polarization*, shifting workers from the middle toward both lower and higher skill sectors; (ii) magnify AI's positive effects on aggregate output and consumption, while dampening its impact on employment; and (iii) alter inequality: even as polarization increases disparities in income and consumption, precautionary saving by middle-sector households reduces the rise in wealth inequality. In an extension (AI+), where AI raises the human-capital threshold for high-skill jobs, additional training and saving become concentrated among high-sector households, further increasing wealth inequality.

Keywords: AI, Human Capital, Job Polarization, Inequality

*We thank Sewon Hur, Ryan Kim, Yena Park, and Jiheum Yeon, as well as seminar and conference participants at Sogang, SNU, the 2025 BSE Summer Forumn, KDI, and the 2026 NAESM. All errors are our own. We acknowledge financial support from the Yonsei Institute for AI and Social Innovation.

†Corresponding author. Economics Department, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong. Email: yanglu@ust.hk.

‡School of Economics, Yonsei University, 50 Yonsei-ro, Seodaemun-gu, Seoul 03722, South Korea. Email: masilver@yonsei.ac.kr.

1 Introduction

Artificial intelligence is reshaping labor markets in ways that are both profound and uneven. Recent evidence shows that AI affects labor demand and productivity heterogeneously: entry-level opportunities are declining, productivity gains depend on worker skill and experience, and firm-level effects vary widely across occupations and firms (Revelio Labs, 2025; Calvino et al., 2025; Souza, 2025; Asam and Heller, 2025). However, the macroeconomic consequences of these changes depend not only on how AI directly reshapes jobs, but also on how households respond. These responses are already visible in human capital decisions: confidence in traditional higher education is eroding, while demand for reskilling and generative-AI training is surging (Higher Education Strategy Associates, 2022; Pew Research Center, 2024; Public Agenda, 2022; Oliver Wyman Forum, 2024). Yet most existing studies of AI's labor-market effects abstract from these endogenous human capital responses. This paper addresses that gap by examining how AI-driven changes in skill returns influence human capital investment and, in turn, reshape aggregate and distributional outcomes throughout the transition.

We build an incomplete-markets model in the tradition of Aiyagari (1994), extended to incorporate endogenous human capital investment and a three-sector labor market. Households are heterogeneous in assets, human capital, and idiosyncratic labor productivity, and face uninsured income risk. Human capital determines sector access – low, middle, or high – with each sector offering a different productivity premium, so that investment grants households the possibility of a discrete wage increase in the future. Workers can accumulate human capital either while employed or out of work. Investment is a discrete choice among three effort levels: none, part-time, and full-time that requires non-employment in the current period. This competition between working and learning is the key mechanism in our model, generating rich interactions among human capital investment, labor supply, and saving decisions.

AI is introduced as an anticipated sector-biased shock that raises productivity in the low and high sectors while leaving the middle unchanged. This is intended to capture the empirical pattern that AI complements high-skill workers, substitutes for routine middle-skill tasks, and augments some lower-skill work. In a second experiment denoted AI+, AI additionally raises the human-capital threshold for high-sector entry, consistent with evidence that AI is reducing junior hiring even as demand for experienced workers rises.

To isolate the core mechanisms, we first study a two-period version of the model, which delivers three analytical results. *First*, uninsured idiosyncratic risk discourages human capital investment and strengthens self-insurance motives: relative to a certainty benchmark, households save more, work more, and invest less in human capital. *Second*, human capital investment has heterogeneous effects on saving. By raising expected future income, investment crowds out saving through a consumption-smoothing channel, but it also raises future income risk, crowding in saving through a precautionary motive. We show that the precautionary motive dominates for high-income households while the consumption-smoothing motive dominates for low-income ones, and that this asymmetry carries over to full-time training when idiosyncratic risk is large enough. *Third*, an anticipated AI shock generates endogenous divergence in investment incentives: households below the middle-skill threshold face weaker incentives to invest toward the middle sector, while those near or in the middle face stronger incentives to reach the high sector. Because full-time training directly competes with labor supply, these divergent investment responses translate into opposing labor supply effects even before AI arrives – some households work more while others voluntarily exit employment to retrain. In the AI+ case, a *churning effect* additionally emerges near the new high-sector cutoff: households sufficiently close invest more aggressively to retain access, while those too far below reduce investment as the goal becomes unattainable, further amplifying the divergence in human capital accumulation and saving behavior.

We then embed these mechanisms in a fully calibrated infinite-horizon general equilibrium model matched to key U.S. moments, including observed inequality measures and sectoral distributions. We study the transition dynamics following an AI adoption that is anticipated ten years in advance, so the economy begins reallocating well before the productivity gains are realized and continues adjusting until it converges to a new steady state.

On the aggregate side, the divergent investment responses identified in the two-period model aggregate into what we call voluntary job polarization. In response to the reduced premium on middle-sector employment, many middle-sector households endogenously exit the labor market to invest in human capital and reposition themselves for the post-AI economy, even before AI is actually implemented. As a result, the population share of the middle sector declines substantially in the new steady state. Aggregate employment also declines, both along the transition path and in the long run, as some households remain out of work while retraining. Human capital investment also becomes more concentrated among a smaller but more selected group of households, with average human capital among investors rising substantially in the new steady state. Comparing the benchmark economy against a counterfactual in which human capital is held fixed reveals that endogenous adjustment raises long-run output, consumption, and investment by reallocating labor more effectively toward the sectors that benefit most from AI. In the AI+ experiment, the churning effect near the new high-sector cutoff generates additional reallocation across sectors, and the gap between the benchmark and fixed-human-capital counterfactual widens further, underscoring that endogenous human capital adjustment becomes even more important when AI reshapes not only sectoral productivity but also the skill requirements for entry into top occupations.

On the distributional side, income and earnings inequality rise due to polarization, as households moving to the low sector experience earnings losses while those reaching the high sector enjoy gains. Wealth inequality, however, behaves differently. Consistent with the saving mechanism identified in the two-period model, middle-sector households undertaking intensive retraining increase saving through the precautionary motive, strengthening asset accumulation in the middle of the wealth distribution and substantially mitigating the rise in wealth inequality relative to the fixed-human-capital counterfactual. In the AI+ experiment, this pattern reverses in an important way. Unlike in the baseline, where the precautionary-saving channel compresses wealth inequality, the churning effect now concentrates additional training among relatively high-income households near the new cutoff. Because these households also save more through the same precautionary mechanism, asset accumulation becomes more top-heavy, pushing wealth inequality above its baseline-AI level. The comparison with the fixed-human-capital counterfactual confirms that these contrasting distributional outcomes are driven by endogenous human capital responses rather than by the direct productivity effects of AI alone.

More broadly, these findings show that abstracting from endogenous human capital adjustment leads to a fundamental mischaracterization of both the aggregate and distributional consequences of AI.

1.1 Related Literature

This paper relates to the literature on how technological change, including AI and robotics, drives job polarization and affects the demand and supply of labor. Studies find that rising employment in both high- and low-wage occupations – at the expense of middle-skill jobs – characterizes job polarization across the UK, US, and Western Europe (Goos and Manning, 2007; Autor and Dorn, 2013; Goos et al., 2014). Robots and automation have also been shown to reduce employment and wages across US regions (Acemoglu and Restrepo, 2020), with automation-induced job losses and declining labor force participation especially concentrated among vulnerable workers in highly automated sectors (Lerch, 2021; Faber et al., 2022). Wang and Wong (2025) model AI as a learning-by-using technology and predicts large productivity gains and employment loss in the long-run. Suh (2023) develops a hierarchical production model in which the distributional effects of automation depend on the complexity of tasks that machines can perform: once technology becomes sufficiently capable, it can amplify top labor incomes and generate "superstar" dynamics. Korinek and Suh (2024) model the transition toward AGI as an expanding frontier of automation over tasks ranked by complexity, showing that macro and wage outcomes depend on how far automation progresses and how it interacts with capital accumulation. Acemoglu and Loebbing (2026) show that automation can generate employment and wage polarization in an assignment model where tasks of intermediate complexity are performed by capital.

Technological disruption also influences human capital accumulation. Faced with employment risks caused by automation, many affected workers invest in further education as a form of self-insurance, rather than relying solely on increases in the college wage premium (Atkin, 2016; Beaudry et al., 2016). Consistent with this, Di Giacomo and Lerch (2023) and Dauth et al. (2021) find that the adoption of industrial robots in the U.S. and Germany, respectively, has led to increased college and university enrollments.

Building on this literature, our paper develops a model that explicitly allows for a trade-off between labor supply and human capital investment. In our framework, job polarization emerges as a voluntary response to AI advancements: households in the middle sector may choose to either downskill to the low sector or upskill to the high sector, while an increasing number of middle-sector households opt for full-time training to accelerate their upskilling.

This paper also relates to the literature that studies human capital and physical capital in a unified framework. Chanda (2008) shows that the rise in returns to education reduces household savings. Waldinger (2016) finds that human capital is much more important than physical capital for innovation in both the short and long-run. Huggett et al. (2011) develops a risky human capital model with incomplete markets to estimate the source of lifetime inequality. Park (2018) investigates whether capital and human capital are over-accumulated in an incomplete market economy. Our model is most similar to Huggett et al. (2011) in that human capital is risky and there is a trade-off between human capital investment and labor supply. Our analysis sheds light on the effect of AI-induced human capital adjustments on household labor supply and saving.

A growing body of literature suggests that AI and automation may contribute to rising inequality across income, consumption, and wealth (e.g., Sachs and Kotlikoff, 2012; Berg et al., 2018; Prettner and Strulik, 2020; Hémous and Olsen, 2022). Our model confirms that AI advancements indeed increase inequality in all three dimensions. However, we find that the endogenous human capital responses to AI amplify the rise in income and consumption inequality, while at the same time mitigating the increase in wealth inequality.

The remainder of the paper proceeds as follows. Section 2 introduces the model environment. Section 3 analyzes household decision-making using a two-period version of the model, clarifying the main mechanisms. Section 4 presents the fully specified quantitative model and its calibration to key features of the U.S. economy, including household heterogeneity. Section 5 examines the aggregate and distributional impacts of AI in the baseline experiment, while Section 6 considers the effects of AI when sectoral skill requirements change. Section 7 concludes.

2 Model Environment

Time is discrete and infinite. There is a continuum of households. Each household is endowed with one unit of time and faces an idiosyncratic labor productivity shock, z, and an idiosyncratic learning-ability shock, y. The labor productivity shock follows an AR(1) process in logs:

$$\ln z' = \rho_z \ln z + \varepsilon_z, \varepsilon_z \overset{\text{iid}}{\sim} N(0, \sigma_z^2)$$

The learning-ability shock follows an AR(1) process in logs:

$$\ln y' = \rho_y \ln y + \varepsilon_y, \varepsilon_y \overset{\text{iid}}{\sim} N(0, \sigma_y^2)$$

Households observe (z_t, y_t) at the beginning of each period before making decisions. The asset market is incomplete following Aiyagari (1994), and the physical capital, a, is the only asset available to households to insure against idiosyncratic risks. Households can also invest in human capital, h, which allows them to work in sectors with different human capital requirements.

2.1 Production Technology

The production technology in the economy is a constant-returns-to-scale Cobb-Douglas production function:

$$F(K, L) = K^{1-\alpha} L^{\alpha}$$

K represents the total physical capital accumulated by households, while L denotes the total effective labor supplied by households, aggregated across three sectors: low, middle, and high. Given the real wage rate per effective labor, w, and the real interest rate, r, the representative firm chooses labor and capital to maximize current profits. Capital depreciates at rate δ_k.

These sectors differ in their technologies for converting labor into effective labor units and in the levels of human capital required for employment. The middle sector employs households with human capital above h_M and converts one unit of labor to one effective labor unit. The high sector, requiring human capital above h_H, converts one unit of labor to $1+\lambda$ effective units, while the low sector, with no human capital requirement, converts one unit into $1-\lambda$ effective units. This implies a sectoral labor productivity $x(h)$ that is a step function in human capital:

$$x(h) = \begin{cases} 1-\lambda & \text{low sector if } h < h_M \\ 1 & \text{middle sector if } h_M \le h < h_H \\ 1+\lambda & \text{high sector if } h \ge h_H \end{cases}$$

A household i who chooses to work n_i hours contributes $z_i x(h_i) n_i$ units of effective labor, where z_i is his idiosyncratic productivity. The aggregate labor is

$$L = \int z_i x(h_i) n_i di,$$

assuming perfect substitutability of effective labor across the three sectors.

2.2 Household's Problem

Households derive utility from consumption, incur disutility from labor and effort of human capital investment. A household maximizes the expected lifetime utility by optimally choosing consumption, saving, labor supply and human capital investment each period, based on his idiosyncratic shocks (z_t, y_t):

$$\max_{\{c_t, a_{t+1}, n_t, e_t\}_{t=0}^{\infty}} E_0 \left[\sum_{t=0}^{\infty} \beta^t (\ln c_t - \chi_n n_t - \chi_e e_t) \right]$$

where c_t represents consumption, a_{t+1} represents saving, n_t is labor supply, and e_t is the effort of human capital investment. Following Chang and Kim (2006), we assume that labor supply is indivisible, so that each household chooses either not to work or to work a fixed number of hours, $\bar{n}$.

If a household decides to work in period t, he will be employed into the appropriate sector according to his human capital h_t and receive labor income $w_t z_t x(h_t)$. The household's budget constraint is

$$c_t + a_{t+1} = w_t z_t x(h_t) n_t + (1 + r_t) a_t$$
$$c_t \geq 0 \text{ and } a_{t+1} \geq 0$$

We prohibit households from borrowing $a_{t+1} \geq 0$ to simplify analysis.

Human capital investment can take three levels of effort: $\{0, e_L, e_H\}$. A non-working household is free to choose any of the three effort levels but a working household cannot devote the highest level of effort e_H, reflecting a trade-off between working and human capital investment. Hence:

$$e_t \in \{0, e_L, (1 - n_t) e_H\}.$$

Its contribution to next-period human capital is subject to the learning-ability shock:

$$h_{t+1} = y_t e_t + (1 - \delta) h_t$$

where δ is human capital's depreciation rate. We interpret $y_t e_t$ as effective human-capital investment, with y_t capturing a learning-ability shock.[1]

[1] The learning ability shock y_t is likely to be correlated with the labor productivity shock z_t in reality. In our quantitative analysis, we assume they are perfectly correlated to reduce the dimensionality of the state space. For readers interested in the two-period model under this perfect-correlation assumption, Appendix A.3 displays the corresponding decision rule diagram.

2.3 AI Shock

AI stands apart from previous waves of automation or computerization because it is capable of performing complex, cognitive, and non-routine tasks that previously required significant education and expertise. While the full impact of AI on the labor market remains to be seen, there is a growing consensus around two major effects.

2.3.1 Evidences for AI's effects on labor market

First, AI complements high-skill workers, substitutes for routine middle-skill tasks, and augments some lower-skill work. Recent experimental evidence surveyed by Calvino et al. (2025) shows that workers' productivity gains from AI depend on their skill levels and experience. On simpler tasks where AI performs well, the technology can narrow the productivity gap between experienced and less experienced workers. However, for more complex tasks that AI cannot yet perform effectively, those with greater digital proficiency or task-specific experience achieve higher productivity gains, as successful use of AI in these settings requires more advanced skills and experience that involves understanding AI's capabilities and limitations.

Firm-level evidence reveals similar patterns. Aghion et al. (2019) documents that the average worker in low-skilled occupations receives a significant wage premium when employed by a more innovative firm. Souza (2025) finds that the adoption of AI in Brazilian firms increases employment for low-skilled production workers but reduces employment and wages for middle-wage office workers. Asam and Heller (2025) report that GitHub Copilot enables software startups to raise initial funding 19% faster with 20% fewer developers, and that these productivity gains disproportionately benefit startups with more experienced founders.

Second, AI may raise the effective skill threshold for high-sector entry, thereby reducing the set of workers who qualify for high-sector employment. Recent labor market data highlight the disproportionate impact of AI on entry-level employment opportunities. Bloomberg (2025) reports that, in the words of Matt Sigelman, president of the Burning Glass Institute, "Demand for junior hires in many college-level roles is already declining, even as demand for experienced hires in the same jobs is on the rise." According to Revelio Labs (2025), postings for entry-level jobs in the US declined by about 35% since January 2023, with roles more exposed to AI experiencing even steeper reductions. Klein Teeselink and Carey (2026) find similar patterns in job posting data across 39 countries. Moreover, they show that occupations where AI exposure is concentrated in low-expertise tasks experience an increase in advertised salaries, reflecting rising expertise requirements for these occupations.

We model these effects as two components of an anticipated AI shock occurring in a future period: a baseline component that changes the skill premium and an additional component that reduces accessibility to the high sector.

2.3.2 Baseline component: AI changes skill premium

Suppose an AI shock is expected to boost labor productivity in both the low and high sectors in a future period, while leaving the middle sector unaffected. We capture this impact on sectoral productivity through the parameter $\gamma \in (0,1)$:

$$x(h') = \begin{cases} 1-\lambda+\gamma\lambda & \text{low sector if } h' < h_M \\ 1 & \text{middle sector if } h_M \leq h' < h_H \\ 1+\lambda+\gamma\lambda & \text{high sector if } h' \geq h_H \end{cases}$$

where h' denotes human capital in the future period. In this baseline component, the AI shock raises average labor productivity, diminishes the earnings premium for the middle sector, and increases the earnings premium for the high sector relative to the middle sector.

2.3.3 Additional component: AI raises the bar for high-sector entry

Suppose an AI shock also raises the threshold for high-sector employment from h_H to h'_H in a future period. Future sectoral productivity is thus given by:

$$x^{AI+}(h') = \begin{cases} 1-\lambda+\gamma\lambda & \text{low sector if } h' < h_M \\ 1 & \text{middle sector if } h_M \leq h' < h'_H \\ 1+\lambda+\gamma\lambda & \text{high sector if } h' \geq h'_H \end{cases}$$

As a result, some households whose human capital falls between the old and new thresholds, that is, $h_H < h' < h'_H$, will be reclassified from high-sector to middle-sector employment.

3 Household Decisions in a Two-Period Model

In this section, we solve the household's problem with two periods to gain intuition. In this analysis, we assume that $\bar{n} = 1$ for simplicity.

Period-2 decisions Households do not invest in human capital or physical capital in the last period. The only relevant decision is whether to work.

The household works $n = 1$ if and only if $z \geq \overline{z}(h,a)$, with $\overline{z}(h,a)$ defined as

$$\ln(w\overline{z}(h,a)x(h) + (1+r)a) - \chi_n = \ln((1+r)a)$$

The household faces a trade-off between earning labor income and incurring the disutility of working. Given the sector-specific productivity $x(h)$ specified in (2.1), the threshold for idiosyncratic productivity, $\overline{z}(h,a)$, takes on three possible values:

$$\overline{z}(h,a) = \begin{cases} \overline{z}(a)\frac{1}{1-\lambda} & \text{if } h < h_M \\ \overline{z}(a) & \text{if } h_M \leq h < h_H \\ \overline{z}(a)\frac{1}{1+\lambda} & \text{if } h \geq h_H \end{cases}$$
$$\text{where } \overline{z}(a) := \frac{(\exp(\chi_n)-1)(1+r)a}{w}$$

Households with higher human capital are more likely to work, whereas households with higher physical capital is less likely to work.

Period-1 decisions In addition to labor supply, period-1 decisions include saving and human capital investment, both of which are forward-looking and affected by the idiosyncratic risk associated with the productivity shock z'. Our model also features a trade-off between human capital investment and labor supply as a working household cannot devote the highest level of effort e_H in human capital investment. Therefore, human capital investment grants households the possibility of a discrete wage hike in the future but may entail a wage loss in the current period.

To see the implication of this trade-off and how it interacts with uninsured idiosyncratic risk, we proceed in two steps. We first derive the period-1 decisions without uncertainty by assuming that z' is known to the household at period 1 and z' is such that the household will work in period 2. We then reintroduce uncertainty in z' and compare the decision rules with the case without uncertainty.

3.1 Period-1 Labor Supply and Human Capital Investment

3.1.1 Consumption and saving without uncertainty

The additive separability of the household's utility implies that labor supply n and human capital investment e enter in consumption and saving choices only via the intertemporal budget constraint:

$$c + \frac{c'}{1+r'} = (1+r)a + wzx(h)n + \frac{w'z'x(h')}{1+r'}$$
$$\text{with } h' = ye + (1-\delta)h.$$

The log utility in consumption implies the optimality condition:

$$c' = \beta(1+r')c.$$

Combining it with the budget constraint, we obtain the optimal consumption as a function of labor supply n and human capital investment e:

$$c(n,e) = \frac{1}{1+\beta}\left[(1+r)a + wzx(h)n + \frac{w'z'x\left(h' = ye + (1-\delta)h\right)}{1+r'}\right].$$

3.1.2 Labor supply and human capital investment

The optimal consumption rules in (3.1.1) and (3.1.1) allow us to express the household's problem as the maximization of an objective function in labor supply n and human capital investment e:[2]

$$\max_{n,e}(1+\beta)\ln c(n,e) - \chi_n n - \chi_e e$$

This maximization depends critically on the household's current human capital and achievable next-period human capital. Accordingly, we partition households into three ranges of h: $[0, h_M(1-\delta)^{-1})$, $[h_M(1-\delta)^{-1}, h_H(1-\delta)^{-1})$, and $[h_H(1-\delta)^{-1}, h_{\max}]$.

[2] This follows since $c' = \beta(1+r')c$, so $\ln c' = \ln\beta + \ln(1+r') + \ln c$.

We now derive the decision rules for households $h \in [0, h_M(1-\delta)^{-1})$ in detail, as the decision rules for the other two ranges are similar. Conditional on a given learning ability y, we define two cutoffs in human capital:

$$\underline{h}_M(y) := \frac{h_M - ye_H}{1-\delta}, \quad \overline{h}_M(y) := \frac{h_M - ye_L}{1-\delta}$$

These cutoffs divide households into three groups based on their ability to be employed in the middle sector in the next period. **Non-learners** are households with $h < \underline{h}_M(y)$. They cannot achieve $h' > h_M$ with either e_L or e_H level of human capital investment today. **Full-time learners** are households with $h \in (\underline{h}_M(y), \overline{h}_M(y))$. These households can reach $h' > h_M$ in the next period only by investing $e = e_H$ today. **Part-time learners** are households with $h > \overline{h}_M(y)$. They can achieve $h' > h_M$ in the next period if they invest $e = e_L$ today.

Conditional on y and learner type, the optimal choices of labor supply and human capital investment are characterized by cutoffs in z. The cutoff formulas vary with households' current sectoral productivity, i.e, if $h < h_M$ or not.

Non-learners will not invest in human capital, $e = 0$, and their future sectoral productivity will be $x(h') = 1-\lambda$. Non-learners work $n = 1$ if and only if:

$$z \geq \overline{z}^L_{non}(a,h) = \frac{\left(\exp(\frac{\chi_n}{1+\beta}) - 1\right)\left[(1+r)a + \frac{w'z'(1-\lambda)}{1+r'}\right]}{w\left(1 - \mathbb{1}_{h<h_M}\lambda\right)}$$

Full-time learners choose between $e = 0$ and $e = e_H$, since $e = e_L$ incurs a cost without any future benefit. Full-time learners must trade off between working and human capital investment: choosing $e = e_H$ requires not working today ($n = 0$), while opting to work means forgoing investment in human capital ($n = 1, e = 0$).[3]

Full-time learners prefer $(n = 1, e = 0)$ to $(n = 0, e = e_H)$ if and only if:

$$z \geq \overline{z}^L_{full}(a,h) = \frac{(\exp(\frac{\chi_n - \chi_e e_H}{1+\beta}) - 1)[(1+r)a + \frac{w'z'}{1+r'}] + \lambda\frac{w'z'}{1+r'}}{w\left(1 - \mathbb{1}_{h<h_M}\lambda\right)}$$

Part-time learners do not exert high effort e_H in human capital investment. They choose among three options: $(n = 1, e = 0)$, $(n = 1, e = e_L)$, and $(n = 0, e = e_L)$,[4] with the decision rule as follows:

$$n(z,h,a), e(z,h,a) = \begin{cases} n = 1, e = 0 & \text{if } z \geq \overline{z}^L_{part}(a,h) \\ n = 1, e = e_L & \text{if } \underline{z}^L_{part}(a,h) \leq z < \overline{z}^L_{part}(a,h) \\ n = 0, e = e_L & \text{if } z < \underline{z}^L_{part}(a,h) \end{cases}$$

[3] The choice between $(n = 0, e = e_H)$ and $(n = 0, e = 0)$ does not depend on z. For e_H to be relevant, λ must be large enough so that $(n = 0, e = e_H)$ is preferred to $(n = 0, e = 0)$. See Appendix A.2 for details on the lower bound for λ.

[4] Similar to the case of full-time learners, the choice between $(n = 0, e = e_L)$ and $(n = 0, e = 0)$ does not depend on z. Moreover, since our model is set up so that $(n = 0, e = e_H)$ dominates $(n = 0, e = 0)$, it implies that $(n = 0, e = e_L)$ dominates $(n = 0, e = 0)$.

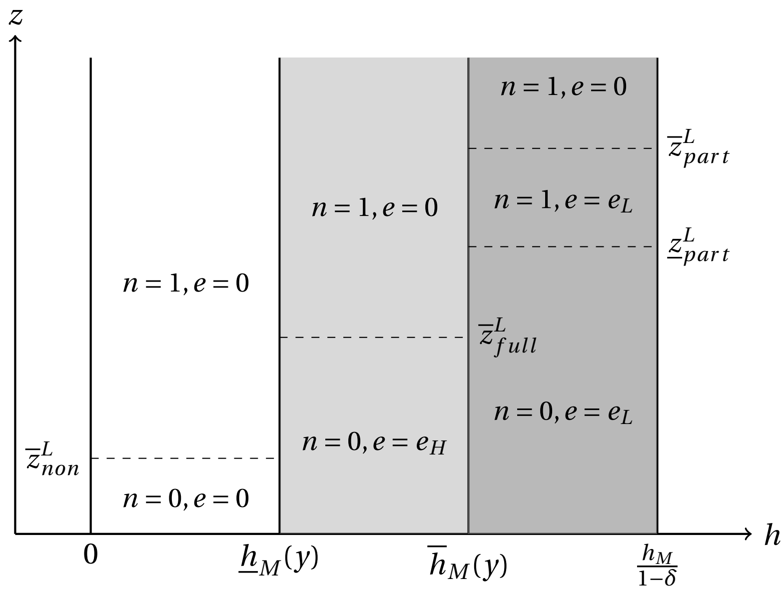

Figure 1: Decision Rule Diagram for $0 \le h < h_M(1-\delta)^{-1}$

Note: Conditional on a given learning ability y, the figure illustrates the decision rule (n, e) as a function of states (z, h, a) for households with $0 \le h < h_M \frac{1}{1-\delta}$. The human capital h changes along the horizontal axis and the labor productivity shock z changes along the vertical axis. The two vertical lines $\underline{h}_M(y)$ and $\overline{h}_M(y)$ defined in (3.1.2) separate the state space into non-learners, full-time learners, and part-time learners. Within each learner type, labor supply and investment choices vary with z.

where

$$\overline{z}^L_{part}(a,h) = \frac{\left\{\exp\left(\frac{\chi_e e_L}{1+\beta}\right)\lambda\left[\exp\left(\frac{\chi_e e_L}{1+\beta}\right)-1\right]^{-1}-1\right\}\frac{w'z'}{1+r'}-(1+r)a}{w\left(1-\mathbb{1}_{h<h_M}\lambda\right)}$$

$$\underline{z}^L_{part}(a,h) = \frac{\left(\exp\left(\frac{\chi_n}{1+\beta}\right)-1\right)\left[(1+r)a+\frac{w'z'}{1+r'}\right]}{w\left(1-\mathbb{1}_{h<h_M}\lambda\right)}$$

We set up our model so that $\overline{z}^L_{part}(a,h) > \underline{z}^L_{part}(a,h)$.[5]

Decision rule diagram: Figure 1 illustrates how the decision rules of labor supply n and human capital investment e vary with (z, h, a) for households with $h < h_M \frac{1}{1-\delta}$ conditional on a given y. The human capital h changes along the horizontal axis and the labor productivity shock z changes along the vertical axis. The two vertical lines $\underline{h}_M(y)$ and $\overline{h}_M(y)$ defined in (3.1.2) separate the state space into non-learners (unshaded area), full-time learners (lightly shaded area), and

[5]Appendix A.2 provides the parameter restrictions such that the condition for $(n=0, e=e_H)$ to dominate $(n=0, e=0)$ is sufficient for $\overline{z}^L_{part}(a,h) > \underline{z}^L_{part}(a,h)$.

part-time learners (darkly shaded area). Within each learner type, labor supply and investment choices vary with z and the dashed horizontal lines represents the cutoffs in z.

Households with $h_M \frac{1}{1-\delta} \le h < h_H \frac{1}{1-\delta}$ have a decision rule analogous to those with $h < h_M \frac{1}{1-\delta}$. Conditional on y, the cutoffs for learner type change to:

$$\underline{h}_H(y) := \frac{h_H - ye_H}{1-\delta}; \qquad \overline{h}_H(y) := \frac{h_H - ye_L}{1-\delta}$$

The four cutoffs in z that partition the decision regions change to $\overline{z}^M_{non}(a,h)$, $\overline{z}^M_{full}(a,h)$, $\underline{z}^M_{part}(a,h)$, and $\overline{z}^M_{part}(a,h)$ (see Appendix A.1 for the explicit formulae).

Households with $h \ge h_H \frac{1}{1-\delta}$ are always non-learners, since their human capital guarantees high-sector employment next period without further investment. For them, only the cutoff $\overline{z}^H_{non}(a)$ matters.

3.2 The Effects of Uninsured Idiosyncratic Risk

We now reintroduce the idiosyncratic risk to households in period 1 by assuming that z' follows a log-normal distribution with mean $\overline{z}'$ and variance σ_z^2.

Our previous analysis without uncertainty is a special case with $\sigma_z^2 = 0$. The effects of uninsured idiosyncratic risk can be thought as how households' decisions change when the distribution of z' undergoes a mean-preserving spread in the sense of second-order stochastic dominance.

From a consumption-saving perspective, the uncertain z' is associated with future labor income risk. It is well understood in the literature that idiosyncratic future income risk raises the expected marginal utility of future consumption for households with log utility and makes them save more. In our model, households can also supply more labor to mitigate the effect of idiosyncratic income risk on the marginal utility of consumption.

From the perspective of human capital investment, the uncertain z' is associated with risk in the return to human capital. Conditional on working, households' income increases with z': $c' = (1+r')a' + w'x(h')z'$. $\ln(c')$ is increasing and concave in z', and a higher $x(h')$ increases the concavity.[6] Consider two levels of h', $\overline{h'} > \underline{h'}$, a mean-preserving spread of z' distribution reduces the expected utility at both levels of h' but the reduction is larger for the higher level $\overline{h'}$. Hence, the expected utility gain of moving from $\underline{h'}$ to $\overline{h'}$ is smaller due to the idiosyncratic risk. Human capital investment is discouraged.

Taking into account endogenous labor supply reinforces the discouragement of human capital investment by the idiosyncratic risk. Recall from Section 3.1.2 that households with z'

[6] The marginal effect of z' on $\ln(c')$ is

$$\frac{\partial \ln(c')}{\partial z'} = \frac{w'x(h')}{(1+r')a' + w'x(h')z'} > 0$$

The second derivative is

$$\frac{\partial^2 \ln(c')}{(\partial z')^2} = -\left[\frac{w'x(h')}{(1+r')a' + w'x(h')z'}\right]^2 < 0$$

and is more negative if $x(h')$ is higher.

lower than a cutoff do not work. The endogenous labor supply therefore provides insurance against the lower tail risk of the idiosyncratic z'. Moreover, the cutoff in z' is lower for those with higher human capital h'. This makes households with higher h' more exposed to the lower tail risk than those with lower h', further reducing the gain of human capital investment.

Proposition 3.1. *The uninsured idiosyncratic risk in z' makes households in period 1 save more, work more and invest less in human capital.*

3.3 Period-1 Saving and Human Capital Investment

In this section, we study the impact of endogenous human capital investment on households' saving decisions. Specifically, we compare optimal saving behavior in two scenarios: one in which households can choose to invest in human capital, and an alternative scenario in which human capital is exogenously fixed. To facilitate the comparison, we assume in this section that there is no human capital depreciation.[7]

When the optimal choice of human capital investment is zero, optimal saving is identical in both scenarios. When the optimal human capital investment is either e_L or e_H, we compare the household's optimal saving to the case where human capital investment is exogenously fixed at zero, i.e., $(n = 1, e = 0)$.[8]

To make the human capital relevant, we assume that $n' = 1$ in period 2. The additive separability of work and human capital investment effort from consumption allows us to consider the optimal saving conditional on a given choice of labor supply and human capital investment.

In particular, the household maximizes expected lifetime utility:

$$\max_{a'} : \ln(c) + \beta \mathbb{E}_{z'}[\ln(c')],$$

subject to the budget constraints

$$\begin{aligned} c + a' &= (1+r)a + n(wzx(h)), \\ c' &= (1+r')a' + w'z'x(h'), \\ \text{with } h' &= ye + (1-\delta)h, e \in \{0, e_L, (1-n)e_H\} \end{aligned}$$

3.3.1 Effect of part-time training on saving

We now compare the optimal saving between $(n = 1, e = e_L)$ and $(n = 1, e = 0)$, where e_L allows households to move to a higher sector in period 2 with higher sectoral productivity $x(h')$.

[7] If depreciation is allowed, the analysis proceeds similarly but involves more comparison paris.

[8] Why not compare to $(n = 0, e = 0)$? Such a comparison is not meaningful when considering $(n = 1, e = e_L)$ because the two scenarios involve different state spaces. To see it, suppose conditions are such that $(n = 1, e = e_L)$ is optimal. If we were to fix $e = 0$ exogenously, the household's lifetime income would fall, and as a result the household would have a greater incentive to work. Thus, it is not possible for the household to deviate from choosing $n = 1$ when human capital is held fixed at $e = 0$. The comparison between $(n = 0, e = 0)$ and $(n = 0, e = e_L$ or $e_H)$ is similar to the comparison between $(n = 1, e = 0)$ to $(n = 1, e = e_L)$, since human capital investment does not affect period-1 labor income in either case.

To simplify the notation while maintaining the key economic forces, we normalize $(1+r) = (1+r') = 1$, $w = w' = 1$, the period-1 productivity shock $z = 1$, the period-1 learning-ability shock $y = 1$, and the period-2 productivity shock z' to $\ln z' \sim \mathcal{N}(0, \sigma_z^2)$. The budget constraints become:

$$c + a' = a + x, \quad c' = a' + txz'$$

where x is the household's period-1 labor income that reflects both productivity and skill. $t \geq 1$ represents the effect of human capital investment on period-2 income: $t > 1$ if $e = e_L$; $t = 1$ if $e = 0$.

The optimal saving is determined by the FOC:

$$\frac{1}{a + x - a'} = \beta \mathbb{E}_{z'}\Big(\frac{1}{a' + txz'}\Big)$$

Denoting the mean and variance of z' as μ and Σ, respectively:

$$\mu \equiv \mathbb{E}[z'] = e^{\sigma_z^2/2}, \quad \Sigma \equiv \mathrm{Var}(z') = e^{\sigma_z^2}(e^{\sigma_z^2} - 1).$$

The second-order approximate solution to the FOC is:

$$a'^{*}(x, a; t) = \underbrace{\frac{\beta(a+x) - tx\mu}{1+\beta}}_{\text{CE}} + \underbrace{\frac{t^2 x^2 \Sigma}{\beta(a + x + tx\mu)}}_{\text{Precautionary}}$$

The first term is the *certainty-equivalent* saving, which reflects the consumption smoothing motive, increasing in the period-1 resources $a + x$ and decreasing in the period-2 expected labor income $tx\mu$. The second term is the *precautionary* saving, which is increasing in the variance of period-2 labor income $t^2 x^2 \Sigma$ and decreasing in the expected total resources $a + x + tx\mu$.

The effect of part-time training on saving can be decomposed into two components:

$$\frac{\partial a'^{\star}}{\partial t}(x, a; t) = -\frac{x\mu}{1+\beta} + \frac{x^2 \Sigma}{\beta} \frac{t\,[2(a+x) + tx\mu]}{(a + x + tx\mu)^2}.$$

The first term being negative captures the *crowd-out* effect on saving via consumption-smoothing motive as part-time training increases the expected period-2 labor income $tx\mu$. The second positive term captures the *crowd-in* effect via precautionary saving motive as part-time training exposes households to larger future income risk.

To capture the overall impact of part-time training on saving, we define:

$$\Delta(x, a; t) = a'^{\star}(x, a; t) - a'^{\star}(x, a; 1) = \int_1^t \frac{\partial a'^{\star}}{\partial u}(x, a; u)\,du,$$

where $a'^{\star}(x, a; t)$ is the optimal saving when households undertake part-time training, and $a'^{\star}(x, a; 1)$ is the optimal saving when human capital is kept exogenously fixed.

Whether part-time training increases or decreases saving ultimately depends on the balance between the crowd-out effect (via higher expected future income) and the precautionary crowd-in effect (via heightened future income risk). The next proposition demonstrates that these effects can dominate differently depending on period-1 income x, so that the overall impact of part-time training on saving can differ between low- and high-income households.

Proposition 3.2. *If the idiosyncratic risk is large enough, i.e., $\frac{\Sigma}{\mu} > \sigma^*(t)$, part-time training crowds out saving for low-income households and crowds in saving for high-income households: for $x < x^*(a,t)$, $e = e_L$ lowers saving $\Delta(x,a;t>1) < 0$; for $x > x^*(a,t)$, $e = e_L$ raises saving $\Delta(x,a;t>1) > 0$.*

Proof. See Appendix C. ∎

3.3.2 Effect of full-time training on saving

We next compare the optimal saving between $(n = 0, e = e_L$ or $e_H)$ and $(n = 1, e = 0)$. Note that full-time training requires the households to give up their labor income in period 1, which is not the case for part-time training. Following the same normalization and notation as in the previous subsection, we can write the budget constraints with full-time training and without training as:

$$\begin{aligned} e = e_H: &\quad c + a' = a, \quad c' = a' + txz' \\ e = 0: &\quad c + a' = a + x, \quad c' = a' + xz' \end{aligned}$$

where $t > 1$ captures the effect of full-time training on period-2 income.

The second-order approximation to the optimal saving problem yields:

$$\begin{aligned} e = e_H: &\quad a'^*_{e_H}(x,a;t) = \underbrace{\frac{\beta a - tx\mu}{1+\beta}}_{\text{CE}} + \underbrace{\frac{t^2x^2\Sigma}{\beta(a+tx\mu)}}_{\text{Precautionary}} \\ e = 0: &\quad a'^*(x,a;1) = \underbrace{\frac{\beta(a+x) - x\mu}{1+\beta}}_{\text{CE}} + \underbrace{\frac{x^2\Sigma}{\beta(a+x+x\mu)}}_{\text{Precautionary}} \end{aligned}$$

The overall effect of full-time training on saving can be expressed as:

$$\begin{aligned} \Delta_{\text{full-time}}(x,a;t) &= a'^*_{e_H}(x,a;t) - a'^*(x,a;1) \\ &= \Delta(x,a;t) + \Delta_H(x,a;t) \\ \text{where } \Delta_H(x,a;t) &\equiv x\left[-\frac{\beta}{1+\beta} + \frac{\Sigma}{\beta}\frac{t^2x^2}{(a+x+tx\mu)(a+tx\mu)}\right] \end{aligned}$$

Here, $\Delta_H(x,a;t)$ captures the additional impact of full-time training on saving, over and above that of part-time training. The first term reflects a further reduction in saving due to the need to forgo period-1 labor income. The second term shows an increase in precautionary saving, as reduced current resources limit households' ability to self-insure against idiosyncratic risk in period 2.

The following lemma establishes some properties of $\Delta_H(x,a;t)$:

Lemma 3.1. *If $\frac{\Sigma}{\mu} < \hat{\sigma}(t)$, $\Delta_H(x,a;t) < 0$ and decreases in x. If $\frac{\Sigma}{\mu} > \overline{\sigma}(t)$, $\Delta_H(x,a;t) > 0$ if and only if $x > \hat{x}(a,t)$; moreover, for $x > \hat{x}(a,t)$, $\Delta_H(x,a;t)$ increases in x.*

Proof. See Appendix C. ∎

Taken together, Proposition 3.2 and Lemma 3.1 imply that, when the idiosyncratic risk is large enough, full-time training *crowds out* saving for low-income households, but *crowds in* saving for high-income households.

Proposition 3.3. *If the idiosyncratic risk is large enough, i.e.,* $\frac{\Sigma}{\mu} > \max\{\sigma^*(t), \hat{\sigma}(t)\}$, *full-time training crowds out saving for low-income households and crowds in saving for high-income households: for* $x < \min\{x^*(a,t), \hat{x}(a,t)\}$, $e = e_H$ *lowers saving* $\Delta_{\textit{full-time}}(x, a; t > 1) < 0$*; for* $x > \max\{x^*(a,t), \hat{x}(a,t)\}$, $e = e_H$ *raises saving* $\Delta_{\textit{full-time}}(x, a; t > 1) > 0$.

3.4 The Effects of an Anticipated Period-2 AI Adoption

We assume that an AI adoption is anticipated to occur in period 2 and study how it affects the household's decisions in period 1.

3.4.1 The effects of skill premium change

As specified in Equation (2.3.2), the baseline AI component lowers the relative earnings advantage of the middle sector in period 2 while increasing it for the high sector.

Human capital investment: The AI adoption lowers the incentive to work in the middle sector in period 2. Consequently, households with $h < h_M/(1-\delta)$ reduce their human capital investment, while those with $h > h_M/(1-\delta)$ increase it. The effect is formalized in the following proposition.

Proposition 3.4. *An anticipated AI adoption decreases human capital investment among households with* $h < h_M/(1-\delta)$, *but increases it among those with* $h > h_M/(1-\delta)$. *Specifically,* $\overline{z}^L_{full}$ *and* $\overline{z}^L_{part}$ *decrease with* γ, *while* $\overline{z}^M_{full}$ *and* $\overline{z}^M_{part}$ *increase with* γ.

Proof. See Appendix C. ■

The AI adoption influences labor supply and saving through two main channels: directly, by altering households' period-2 labor income, and indirectly, by affecting their human capital investment decisions. In what follows, we focus on the indirect effects, while the more standard direct channel is discussed in Appendix B.

Labor supply: Because full-time training and labor supply compete for time, households with $h < h_M/(1-\delta)$ shift away from full-time training and supply more labor in the first period, while households with $h > h_M/(1-\delta)$ become more likely to engage in full-time training and thus reduce period-1 labor supply.

Saving: As established in Propositions 3.2 and 3.3, investing in human capital tends to crowd out saving for low-income households, but crowds in saving for high-income households. Therefore, when the AI adoption leads low-income households ($h < h_M/(1-\delta)$) to stop investing, and encourages high-income households ($h > h_M/(1-\delta)$) to start investing, optimal saving *increases* for both groups.

3.4.2 The effects of raising the high-sector cutoff

The additional component of the AI adoption raises the high-sector cutoff from h_H to $h'_H > h_H$ in period 2 (see Equation (2.3.3)), which in turn increases the period-1 human capital thresholds required to access the high sector with any effort level $e \in \{e_L, e_H\}$. Specifically, the upper bounds for non-learners ($\underline{h}_H(y)$), full-time learners ($\overline{h}_H(y)$), and part-time learners ($h_H \frac{1}{1-\delta}$) all rise by $(h'_H - h_H)\frac{1}{1-\delta}$.

As a result, some households that previously could have remained in the high sector with $e = 0$ under the original cutoff h_H (i.e., those with $h'(0) > h_H$) may now find themselves reclassified into the middle sector if $h'(0) \in (h_H, h'_H)$. For these households, the higher cutoff h'_H in period 2 reduces the attractiveness of choosing $e = 0$ and instead creates an incentive to invest in human capital to remain in the high sector – provided that attaining h'_H is feasible.

Conversely, for households for whom h'_H is unattainable even with the maximum effort ($e = e_H$), the raised cutoff diminishes the returns to training for the high sector, leading these households to opt for $e = 0$ in period 1.

Churning effect Beyond the effect of the baseline γ-only shock (Proposition 3.4), the introduction of the h'_H component generates a *churning* effect that causes some households to transition between different learner types.

Specifically, it shifts some high-sector households with $h_H \frac{1}{1-\delta} < h < h'_H \frac{1}{1-\delta}$ from being non-learners to becoming part-time or full-time learners, depending on whether h'_H is attainable through part-time or full-time investment.

Similarly, it reclassifies households with $h_M \frac{1}{1-\delta} \leq h < h_H \frac{1}{1-\delta}$ across different learner types: some part-time learners may become full-time learners if h'_H is attainable, while others—whether previously part-time or full-time learners—may become non-learners if achieving h'_H is no longer feasible.

Human capital investment and labor supply Conditional on a given learner type, however, the within-type cutoffs in z that determine labor supply and human capital investment remain unchanged from the baseline γ-only shock. The main effect of h'_H is therefore to modify which learner type applies to each household.

Saving The *churning* effect from the h'_H component also affects households' saving behavior, since changes in learner type lead to different human capital investment decisions and, consequently, to different saving outcomes. For those who increase their human capital investment due to the h'_H component, their saving changes by $\Delta(x, a; t)$, $\Delta_H(x, a; t)$ or $\Delta_{\text{full-time}}(x, a; t)$ with $t > 1$. For those who decrease their human capital investment, their saving changes by $-\Delta(x, a; t)$ or $-\Delta_{\text{full-time}}(x, a; t)$.

For households whose human capital investment remains unchanged, their saving responses largely mirror those under the baseline γ-only shock (Appendix B.2), with one notable exception: some households with $h_H \frac{1}{1-\delta} < h < h'_H \frac{1}{1-\delta}$ are reclassified to the middle sector absent any human capital investment. As a result, they anticipate lower period-2 income compared to the γ-only shock, leading to a change in their saving given by $\Delta(x, a; t < 1)$.

3.5 Limitations of the two-period model

Up to this point, our analysis has focused on how AI influences household-level decisions regarding human capital investment, labor supply, and saving within the framework of a two-period model. While this provides valuable insights into individual behavioral responses, understanding the broader, economy-wide implications of AI requires moving to a more comprehensive setting – a quantitative model with an infinite time horizon, endogenous asset accumulation, and general equilibrium feedback.

General equilibrium (GE) effects When households adjust their investment in human capital, labor supply, and savings in response to AI, these changes aggregate up to affect the total supply of effective labor and capital in the economy. As these aggregates shift, they exert downward or upward pressure on the wage rate and the interest rate, feeding back into each household's optimization problem. Thus, general equilibrium effects capture the intricate loop by which individual decisions shape, and are shaped by, the macroeconomic environment.

Composition effects Endogenizing human capital investment injects dynamism into how households sort themselves among the three skill sectors. When an AI adoption occurs, individuals may choose to retrain, upskill, or even move to lower-skilled work, reshaping the distribution of labor across sectors. This shifting composition changes the relative size of each sector, with significant consequences for both aggregate outcomes and the distributional effects of AI.

4 A Full Dynamic Quantitative Model

We now solve the full dynamic model with infinite horizon, endogenous asset accumulation, and general equilibrium. We calibrate the model to reflect key features of the U.S. economy, capturing reasonable household heterogeneity.

4.1 Calibration

We calibrate the model to match key features of the U.S. economy. Some parameters are set to standard values from the literature, while others are chosen to match targeted moments. Table I summarizes the parameter values used in the benchmark model.

The time discount factor, β, is calibrated to match an annual interest rate of 4 percent. We set χ_n to replicate an 80 percent employment rate. We calibrate χ_e to match the fact that around 30 percent of the population invests in human capital in the U.S. (OECD, 2025).

We calibrate parameters regarding labor productivity process as follows. Recall that z follows the AR(1) process in logs: $\log z' = \rho_z \log z + \epsilon_z$, where $\epsilon_z \sim N(0, \sigma_z^2)$. The shock process is discretized using the Tauchen (1986) method, resulting in a transition probability matrix with 11 grids. We set the persistence parameter to $\rho_z = 0.948$ and the standard deviation to $\sigma_z = 0.269$, following the estimates reported in Chang and Kim (2006). In the quantitative analysis, we assume that the learning ability shock y_t and the labor productivity shock z_t are perfectly correlated, i.e., $y_t = z_t$.

Table I: Parameters for the Calibration

Parameter	Value	Description	Target or Reference
β	0.91795	Time discount factor	Annual interest rate
ρ_z	0.948	Persistence of z shocks	Chang and Kim (2006)
σ_z	0.269	Standard deviation of z shocks	Chang and Kim (2006)
$\underline{a}$	0	Borrowing limit	See text
χ_n	2.47	Disutility from working	Employment rate
χ_e	1.48	Disutility from HC effort	See text
$\overline{n}$	1/3	Hours worked	Average hours worked
e_H	1/3	High level of effort	Average hours worked
e_L	1/6	Low level of effort	See text
h_M	0.41	Human capital cutoff for M	See text
h_H	0.96	Human capital cutoff for H	See text
λ	0.2	Skill premium	Earnings Gini
δ	0.1	HC depreciation rate	Standard value
α	0.36	Capital income share	Standard value
δ_k	0.1	Capital depreciation rate	Standard value

We deviate from the two-period model by assuming that the labor supply is a discrete choice between 0 and $\overline{n} = 1/3$. This change only rescales the two-period model without altering the trade-off facing the households. But such rescaling facilitates the interpretation that households are deciding whether to allocate one-third of their fixed time endowment to work. The high-level human capital accumulation effort, e_H is assumed to equal $\overline{n}$. The low-level effort, e_L is set to half of e_H. The skill premium across sectors, λ, is set at 0.2 to match the income Gini coefficient. Human capital cutoffs, h_M and h_H, are set so that the population shares in low, middle, and high sectors are, respectively, 20, 40, and 40 percent. This population distribution roughly matches the fractions of U.S. workers in 2014 who are employed in routine manual occupations (low sector), routine cognitive and non-routine manual (middle sector), and non-routine cognitive (high sector) (Cortes et al., 2017).

On the production side, we set the capital income share, α, to 0.36, and the depreciation rate, δ_k, to 0.1. For simplicity, we assume that human capital depreciates at the same rate, i.e., $\delta = 0.1$.

4.2 Model Fit and Steady-State Distribution

Table II compares key moments generated by the model with their empirical counterparts.[9] The calibration targets three main moments: the employment rate, the human capital investment ratio, and the income Gini coefficient. The model matches each of these moments closely. It reproduces the employment rate of 0.80 and the human capital investment ratio of 0.29 observed in the data. It also matches the empirical income Gini coefficient of 0.58.

[9] Employment and inequality measures are based on the 2007 Survey of Consumer Finances (SCF). Employment is defined as having positive labor income during the year. The human capital investment ratio is taken from OECD (2025).

Table II: Key Moments

Moment	Data	Model
Employment rate	0.80	0.80
Human capital investment ratio	0.29	0.29
Gini coefficient for earnings	0.64	0.62
Gini coefficient for income	0.58	0.58
Gini coefficient for wealth	0.82	0.76

Note: Employment and inequality measures are based on the 2007 SCF. The human capital investment ratio is from OECD (2025). Employment is defined as positive labor income in a given year. The employment rate, the human capital investment rate, and the earnings Gini coefficient are targeted moments.

The earnings and wealth Gini coefficients are not directly targeted in the calibration and therefore provide additional measures of the model's distributional performance. The model generates an earnings Gini coefficient of 0.62, compared with 0.64 in the data. It also produces a wealth Gini coefficient of 0.76, relative to an empirical value of 0.82. Although the model somewhat understates wealth inequality, it captures the overall degree of both earnings and wealth dispersion reasonably well.

Table III provides a more detailed assessment of the model's ability to reproduce the distributions of wealth and earnings. Importantly, none of the quintile shares reported in the table is directly targeted in the calibration. Beyond reproducing the overall degree of inequality summarized by the Gini coefficients, the model also captures the distribution of wealth and earnings across different parts of the population.

The model successfully generates the strong concentration of wealth at the top of the distribution. The top wealth quintile holds 78.1% of aggregate wealth in the model, compared with 83.4% in the data. Although the model shifts some wealth from the top quintile to the fourth quintile, it broadly reproduces the highly skewed shape of the empirical wealth distribution. The fit is particularly close for earnings: the model closely matches the earnings shares of both the middle and upper quintiles, including a top-quintile share of 63.0%, compared with 63.5% in the data. Overall, these untargeted moments show that the model captures not only aggregate measures of inequality but also the more detailed distributional patterns observed in the data.

4.3 AI Shock

The baseline component of AI is modeled as in (2.3.2), with $\gamma = 0.3$. Under this specification, AI increases the productivity of low-sector workers by 7.5% and high-sector workers by 5%, while leaving the middle sector unaffected. Given the employment distribution in the initial steady state, these changes imply that AI raises average labor productivity in the economy by 4%, assuming households do not adjust their behavior. This formulation captures the core idea that AI boosts overall labor productivity (Acemoglu and Restrepo, 2019), while at the same time narrowing the earnings premium for the middle sector and increasing it for the high sector relative to the middle sector.

The additional component of AI is a rise of human-capital threshold between the middle and high sectors as in Equation (2.3.3). We calibrate the new threshold h'_H such that the middle-sector population share in the initial steady state is increased from 40% to 45%. In other words, 5% population will be reclassified out of high-sector employment due to the AI adoption if their

Table III: Wealth and Earnings Distributions

	Quintile					Gini
	1st	2nd	3rd	4th	5th	
WEALTH DISTRIBUTION						
Data	−0.2	1.1	4.5	11.2	83.4	0.82
Model	0.1	0.2	3.7	17.9	78.1	0.76
EARNINGS DISTRIBUTION						
Data	−0.1	4.2	11.7	20.8	63.5	0.64
Model	0.0	4.1	10.8	22.1	63.0	0.62

Note: Quintile entries report the percentage of aggregate wealth or earnings held by each quintile. Households are ranked separately according to wealth and earnings. The model statistics are computed from the stationary distribution. The quintile shares are not directly targeted in the calibration.

human capital remains the same.

We now use the calibrated model to examine the effects of an AI adoption anticipated to arrive in 10 years, under the assumption that households have full information about its timing. This allows us to study how the economy adjusts both in anticipation of and in response to AI adoption.[10]

Our analysis centers on the baseline component, examining both transition dynamics and the differences between the initial and new steady states; we refer to this as the baseline AI experiment. We then consider the additional component only insofar as it qualifies the baseline results; we refer to this as the AI+ experiment.

5 The Baseline AI Experiment

In the baseline experiment, AI is expected to increase the productivity of the low and high sectors, but not the middle sector. As a result, the earnings premium for the middle sector decreases, while the earnings premium for the high sector increases.

5.1 Human Capital Adjustments

Changes in earnings premiums incentivize households to adjust their human capital investments.

Steady-state distribution: Figure 2 compares the initial and new steady-state human capital distributions, highlighting how households reallocate across sectors after the AI adoption. The x-axis reports human capital levels and the y-axis shows the mass of households at each level; the red and blue vertical lines indicate the low–middle and middle–high sector cutoffs, respectively.

[10] Later, we discuss robustness to alternative assumptions regarding both the magnitude and the timing of AI adoption.

Figure 2: Steady-state Human Capital Distribution

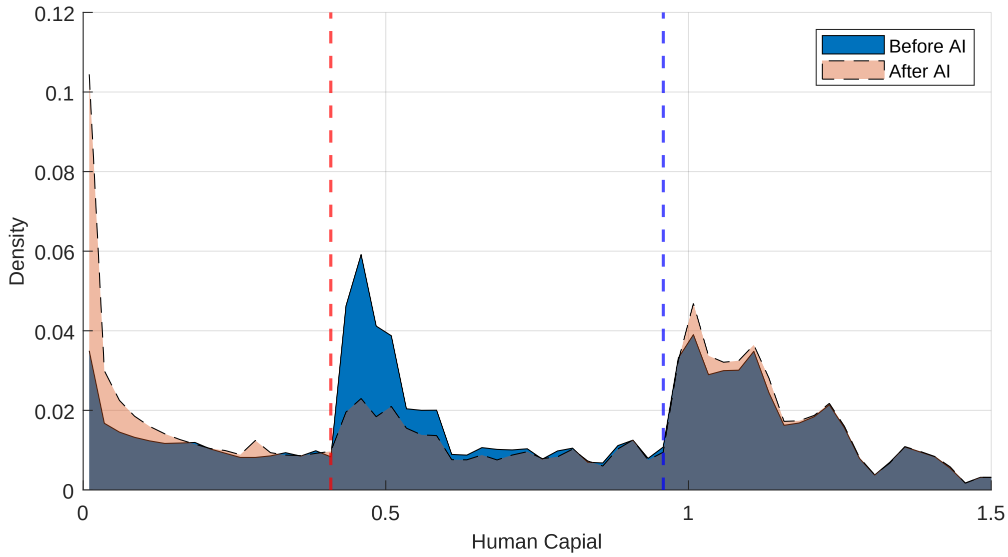

Note: The x-axis denotes the level of human capital, while the y-axis indicates the mass of households at each human capital level. The red vertical line marks the cutoff between the low and middle sectors, and the blue vertical line marks the cutoff between the middle and high sectors.

The gray shaded area indicates the overlap between the initial and new steady-state distributions. Within each sector, the distribution is left-skewed, consistent with human capital investment being concentrated among households near sectoral cutoffs; as illustrated by the decision rule in Figure 1, these households invest to move up or remain in more skilled sectors. The blue shaded area captures households exiting the middle sector after the AI adoption, and the pink shaded areas show where this mass is reallocated in the new steady state—toward the lower end of the low sector and the lower end of the high sector. This reallocation pattern highlights a divergence in human capital investment behavior across households. Proposition 3.4 provides the formal characterization of this mechanism.

Figure 3: Transition Path for Human Capital Investment

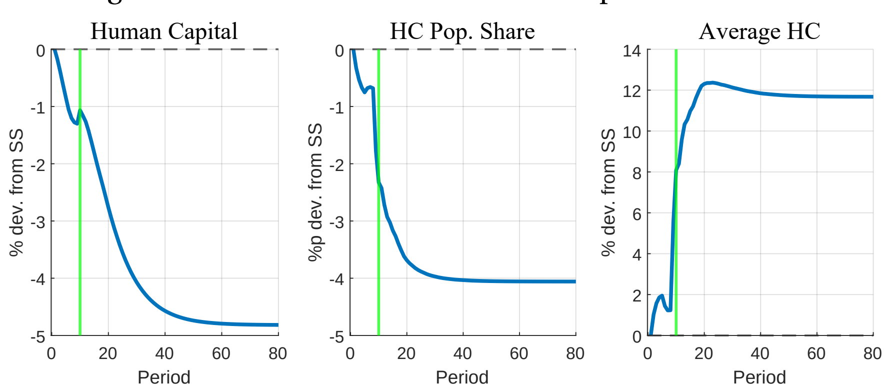

Note: The x-axis represents years, and the y-axis shows the percentage (or percentage point) deviation from the initial steady state. AI introduction is assumed to occur in period 10. "HC Pop. Share" denotes the fraction of households that make positive human capital investments, and "Average HC" denotes average human capital among those investing households.

Figure 4: Sectoral Population and Employment Transition

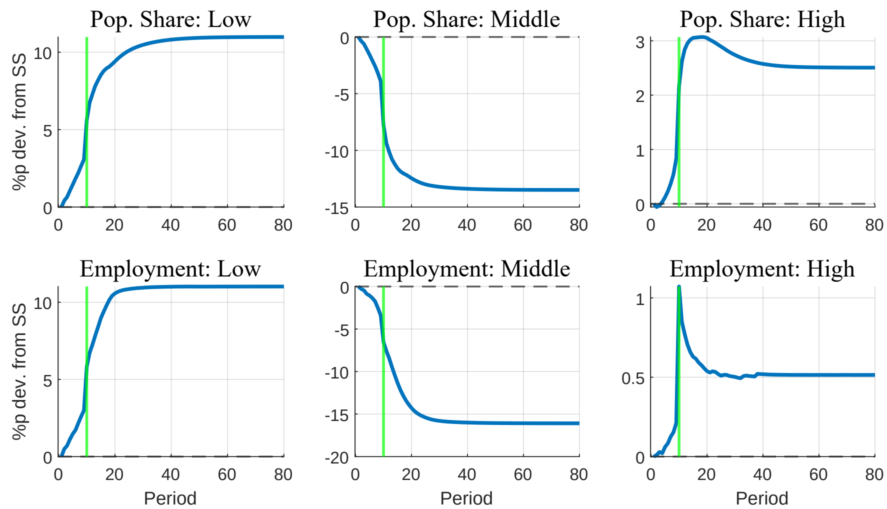

Note: The transition paths within each sector. The x-axis represents years, and the y-axis shows the percentage (or percentage point) deviation from the initial steady state. AI introduction is assumed to occur in period 10. "Pop. Share" denotes the population share within each sector. "Employment" is the percentage of households who are employed in each sector.

Transition path Figure 3 traces the transition of aggregate human capital from the initial to the new steady state, together with its two margins: the extensive margin (the share of households making positive human capital investments) and the intensive margin (average human capital among investing households).

As households move out of the middle sector and into the low and high sectors, aggregate human capital declines gradually along the transition path. This pattern is due to the fact that the increase in mass at the lower end of the low sector is larger than the increase in the high sector, as shown in Figure 2.

At the same time, human capital accumulation becomes more concentrated in a smaller group of households. The share of households making positive human capital investments falls steadily and settles about 4% below its initial steady-state level, while average human capital among investors rises and converges to roughly 12% above its initial steady-state level.[11]

5.2 Job Polarization

An important implication of the human-capital response to AI is job polarization. Figure 4 illustrates the transition paths of sectoral population shares and employment rates. The middle sector contracts sharply: its population share falls by about 13%, and its employment rate declines by 16%. By contrast, both the low and high sectors experience increases in population shares and employment. Taken together, these dynamics imply a reallocation of *workers* away from the middle sector and toward the low and high sectors after AI is introduced.

Voluntary job polarization This reallocation is consistent with the standard job-polarization pattern documented in the literature (Goos et al., 2014), in which AI and automation disproportionately displace tasks typically performed by middle-skill workers. Our model, however,

[11] The only exception to these patterns occurs in period 10, when the productivity effects of AI are realized.

adds a complementary channel. Households can begin leaving the middle sector even before AI is implemented by adjusting their human-capital investment decisions: many middle-sector workers choose non-employment in order to invest in skills that improve their position in the post-AI labor market.[12] This mechanism is formally characterized in Proposition 3.4 in the two-period model above.

Employment flows more toward the low sector A second result is that employment rises much more in the low sector than in the high sector. In the new steady state, the low-sector employment rate increases by 12%, whereas the high-sector employment rate rises by only 0.5%. This asymmetry indicates that reallocation out of the middle sector is tilted toward low-sector employment.

Two forces drive this pattern. First, AI raises productivity by 7.5% in the low sector, compared with 5% in the high sector. But this productivity gap alone is not enough to explain the difference in employment responses. The second force is heterogeneity in labor-supply elasticities: labor supply is more elastic in the low sector than in the high sector, so a given change in labor earnings generates a larger labor-supply response. Intuitively, because low-sector households consume less, their marginal utility of consumption is more sensitive to budget changes. As a result, a larger share of low-sector households is near the employment/non-employment margin (Chang and Kim, 2006).

5.3 Direct and Indirect Effects of AI

The aggregate and distributional effects of AI reflect two forces: AI's direct effects on sectoral productivity and households' endogenous human capital responses. When AI changes productivity across sectors, it shifts labor earnings and, through income effects, affects labor supply and savings decisions. These changes alter aggregate labor and capital supplies, and thus aggregate outcomes. As productivity gains differ across sectors, the same mechanism also generates distributional effects.

Sectoral productivity differences also reshape incentives to accumulate skills and move across sectors. As households reallocate, the composition of labor productivity changes, affecting both aggregate productivity and inequality. In particular, sectoral mobility changes households' relative positions in the income and asset distributions.

This section quantifies how important these human capital adjustments are for both transition dynamics and long-run outcomes. We compare the benchmark economy—where households adjust human capital endogenously—with a "No HC model," in which human capital is fixed at its initial steady-state level throughout the AI transition. In both economies, households still choose consumption, savings, and labor supply endogenously.

This comparison lets us decompose AI's effects. The No-HC model isolates the direct effects of AI on aggregate and distributional outcomes by shutting down human capital adjustment. The gap between the benchmark and No-HC paths then identifies the indirect effects that operate through endogenous human capital accumulation. Together, this decomposition clarifies the role of human capital dynamics in shaping the overall consequences of AI.[13]

[12] To highlight this mechanism, the model deliberately abstracts from any direct negative effect of AI on middle-sector workers.

[13] We also examine the general equilibrium channel of AI effects by comparing the benchmark model with a

Figure 5: Transition Path of Aggregate Variables: Benchmark vs. No HC Models.

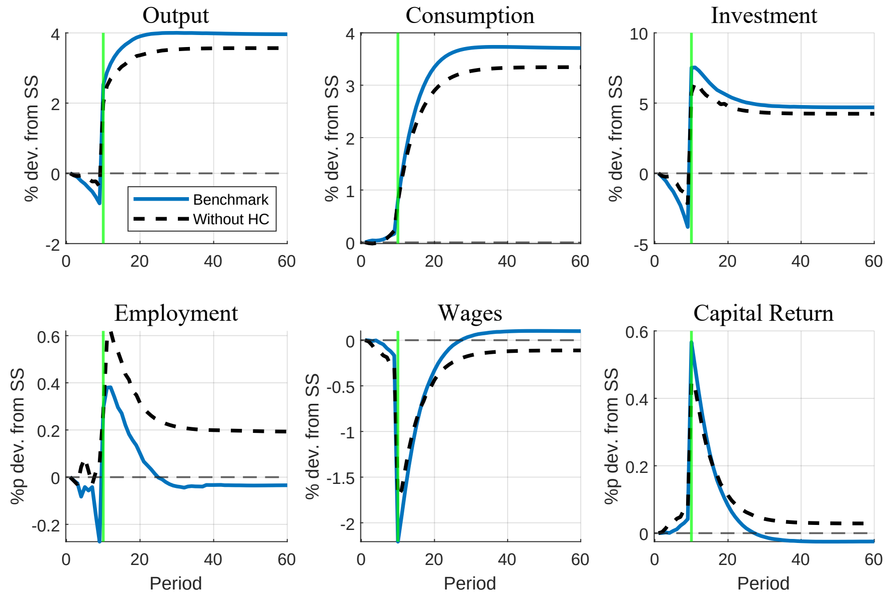

Note: The transition paths of aggregate variables: benchmark vs. No HC models. The x-axis represents years, and the y-axis shows the percentage deviation from the initial steady state. AI introduction is assumed to occur in period 10. The No HC model is an economy in which workers maintain their initial steady-state level of human capital throughout the AI implementation until the new steady state is reached.

5.4 Aggregate Outcomes

Figure 5 reports transition paths for output, consumption, investment, employment, and factor prices (the wage rate and the return to capital). Blue solid lines denote the benchmark economy with endogenous human capital adjustment, and black dashed lines denote the No-HC economy in which human capital is fixed.

5.4.1 AI's direct impacts via productivity changes

The No-HC model isolates the direct effects of AI. In the long run, AI raises output, consumption, investment, and employment. Before period 10, however, output and investment fall, while consumption and employment remain broadly stable.

The intuition is straightforward. Prior to AI implementation, sectoral productivities are unchanged; what changes is that households anticipate future productivity gains in the low and high sectors. Higher expected lifetime income leads households to save less and consume more, which lowers the aggregate capital stock. With less capital, output declines, the marginal product of labor falls, and the marginal product of capital rises. Employment changes little in this phase because current production technologies have not yet shifted.

Once AI arrives, productivity in the low and high sectors increases, raising labor income, employment, and output in those sectors. Because these gains are labor-augmenting, the supply of efficient labor units expands, putting downward pressure on wages and upward pressure on capital returns. Endogenous adjustments in employment and investment partially offset these

fixed-factor-prices model. As this channel is not our main focus and does not produce qualitative differences in results, we relegate the discussion to Appendix E.

factor-price movements. As a result, in the new steady state, wages are slightly below their initial level, while capital returns are slightly above it.

5.4.2 AI's indirect impacts via human capital responses

Comparing the benchmark and No-HC paths identifies the indirect effects that operate through endogenous human capital adjustment. These effects are most visible in employment dynamics.

Before period 10, employment declines in the benchmark economy because some households temporarily withdraw from work to invest in human capital and prepare for the post-AI environment.[14] Since productivity is unchanged during this anticipation phase, lower employment directly reduces output. Under consumption smoothing, most of this output loss is absorbed by lower investment. The employment decline also dampens the pre-AI movements in wages and capital returns found in the No-HC model.

After AI is introduced, employment recovers as productivity rises. Even so, continued human capital investment—especially among middle-sector households—keeps employment below the No-HC path, so the long-run net effect on employment is close to zero. At the same time, output, consumption, and investment are all higher in the benchmark economy, because endogenous skill accumulation reallocates labor toward the low and high sectors and improves the economy's ability to capture AI-driven productivity gains.

This reallocation also flips the long-run factor-price results: with endogenous human capital adjustment, the direct negative effect of AI on wages becomes a positive net effect, while the direct positive effect on capital returns becomes a negative net effect.

Taken together, these results show that aggregate outcomes are shaped not only by AI-induced productivity changes themselves, but also by households' endogenous adaptation through human capital investment.

5.5 Distributional Outcomes

Distributional outcomes are even more sensitive to endogenous human capital adjustment than aggregate outcomes. Figure 6 plots transition paths of Gini coefficients for earnings, total income, consumption, wealth (asset holdings), and human capital. Black dashed lines show the No-HC model (direct AI effects with fixed human capital), while blue solid lines show the benchmark model with endogenous human capital adjustment.

5.5.1 Income, earnings, and consumption inequalities

Comparing the No-HC and benchmark paths shows that human capital responses materially reshape AI's effects on earnings, income, and consumption inequality.

AI's direct impacts: With human capital fixed, distributional changes are driven mainly by sectoral productivity gains in the low and high sectors (7.5% and 5%, respectively). Accordingly, before period 10 there is little movement in the earnings and income Gini coefficients. After AI is implemented, both earnings and income inequality decline: higher productivity raises earnings

[14] Empirical studies, such as Lerch (2021) and Faber et al. (2022), support the short-term adverse effects of AI adoption on labor markets.

Figure 6: Transition Path of Inequality Measures: Benchmark vs. No HC Models.

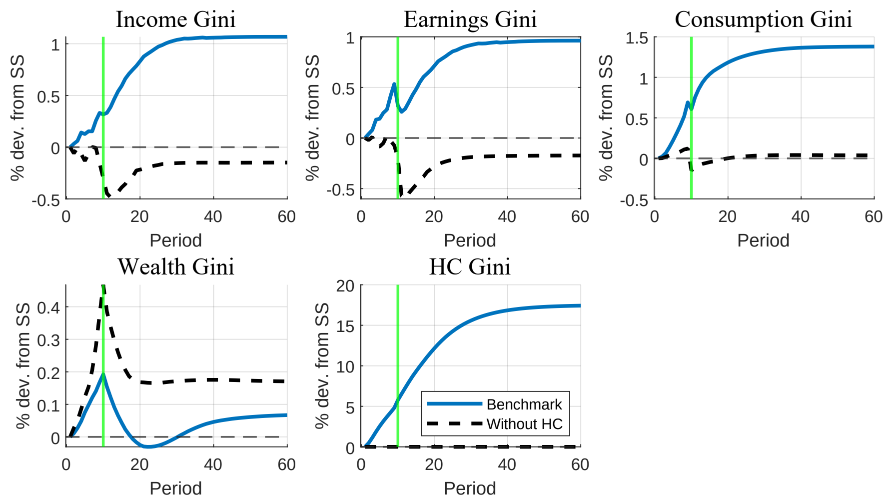

Note: The transition paths of inequality measures: benchmark vs. No HC models. The x-axis represents years, and the y-axis shows the percentage deviation from the initial steady state. AI introduction is assumed to occur in period 10. The No HC model is an economy in which workers maintain their initial steady-state level of human capital throughout the AI implementation until the new steady state is reached.

in the low sector, while lower wages in the middle sector compress the overall distribution. Consumption inequality remains broadly stable throughout the transition.

Effects of AI-induced human capital adjustments: Allowing human capital to adjust endogenously generates pronounced job polarization (Section 5.2). Households that would otherwise remain in middle-sector jobs reallocate toward either the low or high sector. Those moving down experience lower labor earnings, while those moving up gain higher earnings. This widening dispersion raises both earnings and income inequality, before and after AI adoption. As these income gaps expand, consumption inequality rises as well.

5.5.2 Human capital and wealth inequalities

Human-capital inequality (Gini) rises already during the anticipation phase and continues to increase after AI adoption. This pattern is a direct consequence of households reallocating from the middle sector toward the low and high sectors.

The dynamics of wealth inequality are more nuanced. The wealth Gini increases after AI adoption in both models, but the rise is smaller in the benchmark economy than in the No-HC economy. This comparison indicates that human capital responses mitigate, rather than amplify, AI's unequalizing effect on wealth.

To clarify the mechanism, it is useful to first describe the direct response of household saving in the No-HC economy and then explain how endogenous human-capital adjustment modifies it. Figure 7 shows the transition paths of average asset holdings by households in each sector.

AI's direct impacts: Without human-capital adjustment, AI affects household saving only through income effects. In the low and high sectors, households reduce saving before period 10 because they expect higher future labor income; after AI is implemented, they increase saving as labor income rises. In contrast, middle-sector households anticipate a negative income effect

Figure 7: Sectoral Average Asset Transition: Benchmark vs. No HC Models

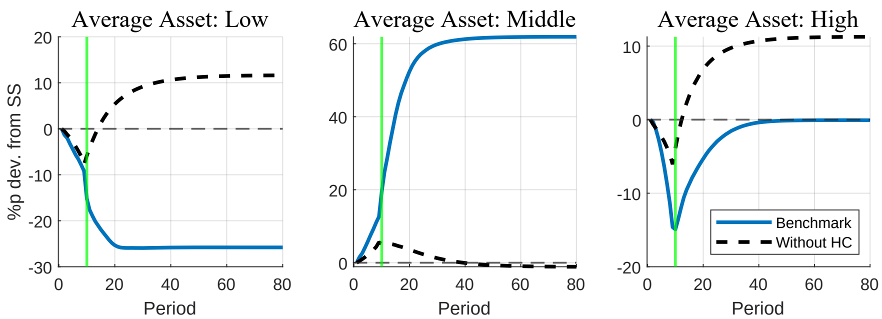

Note: The transition paths within each sector. "Average Asset" is defined as the total assets in a given sector divided by that sector's population share. The x-axis represents years, and the y-axis shows the percentage (or percentage point) deviation from the initial steady state. AI introduction is assumed to occur in period 10. The No HC model is an economy in which workers maintain their initial steady-state level of human capital throughout the AI implementation until the new steady state is reached.

from lower wages and therefore increase saving before period 10. Once AI is introduced and wages recover, they reduce saving.

Effects of AI-induced human capital adjustments: Allowing human capital to adjust significantly changes average household saving in all three sectors. The most pronounced shift occurs in the middle sector, where average saving rises markedly; this is the main reason why endogenous human-capital responses mitigate AI's unequalizing effect on wealth.

Two channels drive this result. First, a composition channel: as households move out of the middle sector, less productive households disinvest in human capital and sort into the low sector, while more productive households increase human capital investment and move up to the high sector. Because the share moving down exceeds the share moving up, the remaining middle-sector population is positively selected on productivity. Second, a saving-crowding-in channel: as implied by Propositions 3.2 and 3.3, human capital investment crowds in saving for high-income households. Because the composition channel raises the average productivity of the middle sector, more middle-sector households now fall into the range where human capital investment crowds in saving, boosting asset accumulation in the middle of the distribution. Together, these two channels narrow the gap between the middle and the top.[15]

5.6 Robustness

In the benchmark economy, AI is introduced so that γ reaches 0.3 in period 10. We now show that the main results are robust to alternative assumptions regarding both the magnitude and the timing of AI adoption. First, we consider alternative terminal values of γ, namely $\gamma = 0.2$ and $\gamma = 0.4$. We next consider a gradual-adoption specification in which AI is phased in over time and becomes fully implemented by period 10. Specifically, γ is assumed to increase linearly over time until it reaches its new long-run value in period 10.

[15] At the same time, average assets in the high sector decline as relatively low-wealth households move up from the middle sector. In the low sector, average assets also fall because the low-sector population is negatively selected on productivity.

Figure 8: High-Sector Population and Human Capital in AI+ Experiment

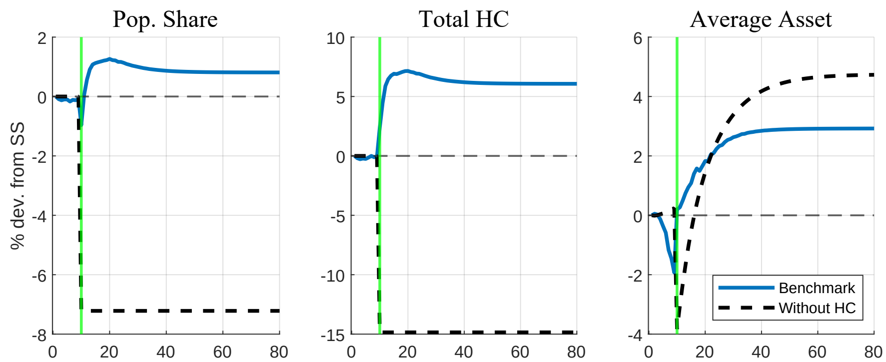

Note: Transition paths of the high-sector population share, total human capital, and average assets in the AI+ experiment: benchmark vs. No HC models. The x-axis represents years, and the y-axis shows the percentage deviation from the initial steady state. AI introduction is assumed to occur in period 10.

Figures F.1 and F.2 in the Appendix present the aggregate and distributional effects under $\gamma = 0.2$ and $\gamma = 0.4$. The main qualitative results remain unchanged. Relative to the model without endogenous human capital accumulation, the model with human capital generates a larger increase in GDP and a larger decline in employment, as workers reduce labor supply in order to invest in human capital during the transition. These endogenous human capital responses also continue to generate voluntary job polarization, leading to higher income and consumption inequality. Wealth inequality also remains lower in the benchmark model than in the model without human capital accumulation.

Figure F.3 in the Appendix reports the results for the gradual-adoption case. The qualitative implications are again very similar to those in the benchmark.

Overall, these robustness exercises confirm that our main qualitative results are not sensitive to either the magnitude of the AI shock or the timing of its adoption.

6 The AI+ Experiment

In the AI+ experiment, AI changes more than the skill premium: it also changes what firms classify as "high-skill" work by increasing the cutoff for entering the high sector. This tighter threshold implies that some workers who would previously have remained in the high sector are reclassified into the middle sector unless they increase their human capital investment.

6.1 High-sector Human Capital and Saving Adjustments

A higher high-sector cutoff generates a **churning effect** (Section 3.4.2). Workers close to the new threshold have stronger incentives to invest in human capital to remain in high-sector employment, whereas those far below it may scale back learning because moving into the high sector becomes less attainable.

Quantitatively, this churning effect leads to a significant increase in human capital investment among high-sector households relative to the baseline AI experiment. Figure 8 reports the

Figure 9: Aggregate Variables in AI+ experiment: Benchmark vs. No HC Models.

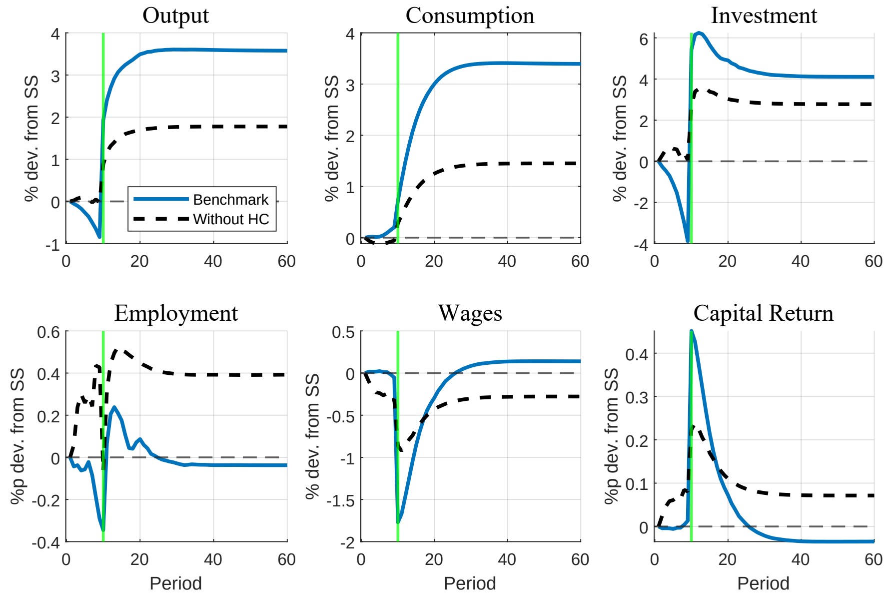

Note: The transition paths of aggregate variables in the AI+ experiment: benchmark vs. No HC models. The x-axis represents years, and the y-axis shows the percentage deviation from the initial steady state. AI introduction is assumed to occur in period 10.

transition dynamics of the high-sector population share and total human capital.[16] Although a higher cutoff would mechanically lower both variables, the stronger skill-investment response is large enough to reverse that effect, so both the high-sector population share and total human capital increase along the transition.

This concentration of skill investment among high-sector – and therefore high-productivity – households also crowds in precautionary saving, as highlighted by Propositions 3.2 and 3.3. The right panel of Figure 8 shows the transition path of average assets in the high sector. Relative to the baseline AI experiment, high-sector average assets are about 3% higher in the new steady state, while they are essentially unchanged in the baseline.

6.2 Aggregate and Distributional Implications

Because the churning effect induces a sharper increase in human-capital investment and precautionary saving among high-sector households than in the baseline AI experiment, the AI+ economy differs from the baseline in two important ways: the indirect human-capital channel accounts for a larger share of AI's aggregate effects, and both human-capital and wealth inequality rise more strongly.

Figure 9 shows the transition dynamics of aggregate variables. Qualitatively, the pattern is similar to the baseline AI experiment: AI adoption still raises real activity over the transition. Quantitatively, however, the underlying composition changes. Two offsetting forces are at work. On the one hand, the higher high-sector cutoff mechanically reclassifies some workers from the high sector to the middle sector, muting AI's direct productivity effect. On the other hand, this reclassification amplifies the human capital channel via the churning mechanism,

[16]See Figure **??** for the corresponding transition dynamics across all sectors in the AI+ experiment.

Figure 10: Inequality Measures in AI+ experiment: Benchmark vs. No HC Models.

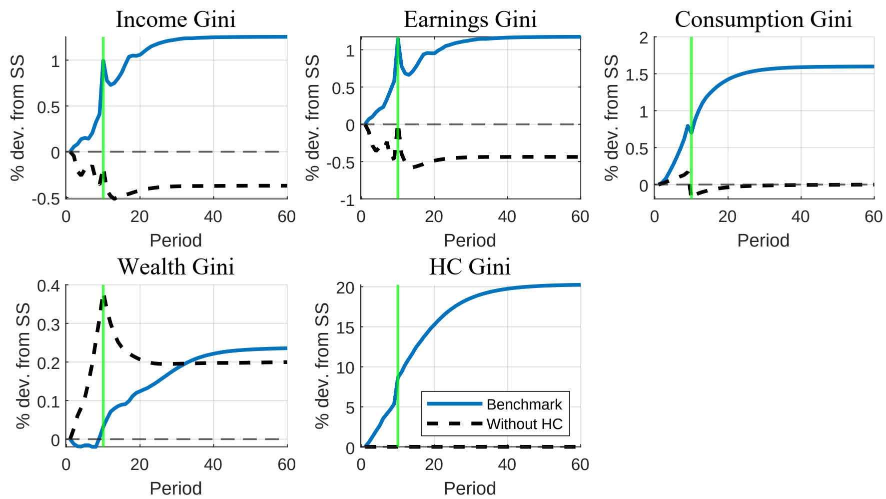

Note: The transition paths of inequality measures in the AI+ experiment: benchmark vs. No HC models. The x-axis represents years, and the y-axis shows the percentage deviation from the initial steady state. AI introduction is assumed to occur in period 10.

prompting households to invest more aggressively in human capital to meet the higher threshold (Section 6.1). As a result, aggregate activity remains broadly comparable to the baseline, but a larger portion of the total effect operates through endogenous human-capital accumulation.

Figure 10 reports the transition dynamics of inequality measures. Income, earnings, and consumption inequality evolve much as in the baseline. By contrast, human-capital and wealth inequality respond more strongly. The churning effect widens dispersion in human capital investments: relatively productive households invest more and regain high-sector access, while others fall behind. The skill investment among high-productivity households also crowds in precautionary saving, making asset accumulation more top-heavy and amplifying wealth inequality relative to the baseline. In the No-HC model, however, wealth-inequality responses are broadly similar across the two AI specifications, indicating that the additional increase under AI+ is driven by endogenous human-capital adjustment rather than by the direct productivity effect.[17]

7 Conclusion

This paper studies how anticipated advances in artificial intelligence reshape the economy when households can respond through human capital accumulation. We develop an incomplete-markets general equilibrium model with endogenous human capital investment, labor supply, and saving, and use it to examine how AI-driven changes in sectoral skill returns affect aggregate dynamics and inequality. Our main message is that the consequences of AI depend not only on its direct effects on productivity, but also on how households adapt in advance through retraining, labor supply adjustment, and precautionary saving.

In the baseline experiment, AI raises productivity in the low- and high-skill sectors while

[17] As an additional robustness check, we also consider cases in which the initial steady-state share of the middle-sector population is increased by 7 percent and 10 percent. The main results remain qualitatively similar. The corresponding results are available upon request.

leaving the middle sector unchanged, thereby reducing the relative attractiveness of middle-skill work. This generates endogenous divergence in human capital investment incentives and leads to what we call voluntary job polarization. Many middle-sector households begin to adjust before AI is implemented, with some moving toward the low sector and others withdrawing from work temporarily to invest in skills aimed at high-sector employment. These human capital responses amplify the long-run gains in output, consumption, and investment relative to a counterfactual economy with fixed human capital, but they also dampen employment during the transition and raise income and consumption inequality through stronger polarization.

At the same time, the distributional effects of AI are more nuanced than a simple polarization narrative would suggest. In our benchmark economy, endogenous human capital adjustment mitigates the rise in wealth inequality. The reason is that middle-sector households facing retraining risk increase saving for precautionary reasons, strengthening asset accumulation in the middle of the wealth distribution. In the AI+ experiment, however, AI also raises the human-capital threshold for high-sector employment. This produces a churning effect near the new cutoff: households close to the threshold invest more aggressively to retain access, while those further below reduce investment. Because the additional training and precautionary saving become concentrated among relatively high-productivity households, wealth accumulation becomes more top-heavy and wealth inequality rises more strongly than in the baseline experiment.

Taken together, our findings imply that analyses that abstract from endogenous human capital adjustment may mischaracterize both the aggregate and distributional consequences of AI.

There are several promising directions for future research. One important extension is to enrich the way AI is modeled. In our framework, AI operates through sector-specific productivity shifts and, in the AI+ experiment, through a higher human-capital threshold for high-skill employment. This structure is useful for isolating the human-capital adjustment channel, but future work could consider richer forms of technological change in which AI both substitutes for and complements tasks along a broader task spectrum, for example by modeling tasks ranked by complexity as in Korinek and Suh (2024). Relatedly, allowing for more flexible substitution patterns or sector-specific wages could provide a more realistic account of how AI reshapes relative earnings and retraining incentives across different parts of the labor market.

A second important direction is to study policy design under alternative AI scenarios. If AI changes both the returns to skills and the incentives to accumulate them, then the effects of redistribution, training subsidies, social insurance, and public insurance against displacement risk may depend crucially on the nature of the technological change itself. Examining these interactions would help clarify which policies are most effective in supporting adjustment, limiting inequality, and improving welfare during the transition to an AI-intensive economy. Finally, it would also be valuable to relax perfect foresight and allow households to learn gradually about the timing, scale, and incidence of AI adoption.

References

Acemoglu, Daron and Loebbing, Jonas (2026). "Automation and Polarization", *Journal of Political Economy*, Vol. 134 No. 3, pp. 1017–1072.

Acemoglu, Daron and Restrepo, Pascual (2019). “8. Artificial Intelligence, Automation, and Work”, *The Economics of Artificial Intelligence*. Ed. by Agrawal, Ajay, Gans, Joshua, and Goldfarb, Avi. University of Chicago Press: Chicago, pp. 197–236.

— (2020). “Robots and Jobs: Evidence from US Labor Markets”, *Journal of Political Economy*, Vol. 128 No. 6, pp. 2188–2244.

Aghion, Philippe et al. (2019). *The Innovation Premium to Soft Skills in Low-Skilled Occupations*, Working papers No. 739. Banque de France.

Aiyagari, S. Rao (1994). “Uninsured Idiosyncratic Risk and Aggregate Saving”, *The Quarterly Journal of Economics*, Vol. 109 No. 3. Publisher: Oxford University Press, pp. 659–684.

Asam, Dominik and Heller, David (2025). *Generative AI and Firm-level Productivity: Evidence from Startup Funding and Employment Dynamics*, Rochester, NY.

Atkin, David (2016). “Endogenous Skill Acquisition and Export Manufacturing in Mexico”, *American Economic Review*, Vol. 106 No. 8, pp. 2046–85.

Autor, David H. and Dorn, David (2013). “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market”, *American Economic Review*, Vol. 103 No. 5, pp. 1553–97.

Beaudry, Paul, Green, David A., and Sand, Benjamin M. (2016). “The Great Reversal in the Demand for Skill and Cognitive Tasks”, *Journal of Labor Economics*, Vol. 34 No. S1, S199–S247.

Berg, Andrew, Buffie, Edward F., and Zanna, Luis-Felipe (2018). “Should we fear the robot revolution? (The correct answer is yes)”, *Journal of Monetary Economics*, Vol. 97, pp. 117–148.

Bloomberg (2025). *AI’s Takeover of Entry-Level Tasks Is Making College Grads’ Job Hunt Harder*, accessed 02 Dec. 2024.

Calvino, Flavio, Reijerink, Jelmer, and Samek, Lea (2025). *The effects of generative AI on productivity, innovation and entrepreneurship*, OECD Artificial Intelligence Papers. Edition: 39 Series: OECD Artificial Intelligence Papers.

Chanda, Areendam (2008). “The rise in returns to education and the decline in household savings”, *Journal of Economic Dynamics and Control*, Vol. 32 No. 2, pp. 436–469.

Chang, Yongsung and Kim, Sun-Bin (2006). “From Individual to Aggregate Labor Supply: A Quantitative Analysis based on a Heterogeneous Agent Macroeconomy”, *International Economic Review*, Vol. 47 No. 1, pp. 1–27.

Cortes, Guido Matias, Jaimovich, Nir, and Siu, Henry E (2017). “Disappearing routine jobs: Who, how, and why?”, *Journal of Monetary Economics*, Vol. 91. Publisher: North-Holland, pp. 69–87.

Dauth, Wolfgang et al. (2021). “The Adjustment of Labor Markets to Robots”, *Journal of the European Economic Association*, Vol. 19 No. 6, pp. 3104–3153.

Di Giacomo, Giuseppe and Lerch, Benjamin (2023). “Automation and Human Capital Adjustment”, *Journal of Human Resources*,

Faber, Marius, Sarto, Andres P, and Tabellini, Marco (2022). *Local Shocks and Internal Migration: The Disparate Effects of Robots and Chinese Imports in the US*, Working Paper No. 30048. National Bureau of Economic Research.

Goos, Maarten and Manning, Alan (2007). “Lousy and Lovely Jobs: The Rising Polarization of Work in Britain”, *The Review of Economics and Statistics*, Vol. 89 No. 1, pp. 118–133.

Goos, Maarten, Manning, Alan, and Salomons, Anna (2014). “Explaining Job Polarization: Routine-Biased Technological Change and Offshoring”, *American Economic Review*, Vol. 104 No. 8, pp. 2509–26.

Hémous, David and Olsen, Morten (2022). “The Rise of the Machines: Automation, Horizontal Innovation, and Income Inequality”, *American Economic Journal: Macroeconomics*, Vol. 14 No. 1, pp. 179–223.

Higher Education Strategy Associates (2022). *World Higher Education: Institutions, Students and Funding*, accessed 02 Dec. 2024.

Huggett, Mark, Ventura, Gustavo, and Yaron, Amir (2011). “Sources of Lifetime Inequality”, *The American Economic Review*, Vol. 101 No. 7, pp. 2923–2954.

Klein Teeselink, Bouke and Carey, Daniel (2026). *AI, Automation, and Expertise*, Rochester, NY.

Korinek, Anton and Suh, Donghyun (2024). *Scenarios for the Transition to AGI*, NBER Working Paper No. 32255. Also available as arXiv:2403.12107. National Bureau of Economic Research.

Lerch, Benjamin (2021). *Robots and Nonparticipation in the US: Where Have All the Workers Gone?*, IdEP Economic Papers No. 2003. USI Università della Svizzera italiana.

OECD (2025). *Trends in Adult Learning: New Data from the 2023 Survey of Adult Skills*, Getting Skills Right. OECD Publishing.

Oliver Wyman Forum (2024). *How Generative AI is Transforming Business and Society*,

Park, Yena (2018). “Constrained Efficiency in a Human Capital Model”, *American Economic Journal: Macroeconomics*, Vol. 10 No. 3, pp. 179–214.

Pew Research Center (2024). *Public Views on the Value of a College Degree*, accessed 02 Dec. 2024.

Prettner, Klaus and Strulik, Holger (2020). “Innovation, automation, and inequality: Policy challenges in the race against the machine”, *Journal of Monetary Economics*, Vol. 116, pp. 249–265.

Public Agenda (2022). *Young Americans Without College Degrees Are Most Skeptical About the Value of Higher Education*, accessed 02 Dec. 2024.

Revelio Labs (2025). *Is AI Responsible for the Rise in Entry-Level Unemployment?*, accessed 02 Dec. 2024.

Sachs, Jeffrey D. and Kotlikoff, Laurence J. (2012). *Smart Machines and Long-Term Misery*, NBER Working Papers No. 18629. National Bureau of Economic Research, Inc.

Souza, Gustavo de (2025). *Artificial Intelligence in the Office and the Factory: Evidence from Administrative Software Registry Data*, Rochester, NY.

Suh, Donghyun (2023). *Machines and Superstars: Technological Change and Top Labor Incomes*, SSRN Working Paper. Date written: November 21, 2023; posted December 1, 2023.

Tauchen, George (1986). “Finite state markov-chain approximations to univariate and vector autoregressions”, *Economics Letters*, Vol. 20 No. 2, pp. 177–181.

Waldinger, Fabian (2016). “Bombs, Brains, and Science: The Role of Human and Physical Capital for the Creation of Scientific Knowledge”, *The Review of Economics and Statistics*, Vol. 98 No. 5, pp. 811–831.

Wang, Ping and Wong, Tsz-Nga (2025). “Artificial Intelligence and Technological Unemployment”, NBER Working Paper No. 33867.

Appendix

A Household Decision Rule Cutoffs

A.1 Additional cutoffs formulae for households

$$\overline{z}^M_{non}(a,h) := \frac{(\exp(\frac{\chi_n}{1+\beta})-1)[(1+r)a+\frac{w'z'}{1+r'}]}{w\left(1+\mathbb{1}_{h>h_H}\lambda\right)}$$

$$\overline{z}^M_{full}(a,h) := \frac{(\exp(\frac{\chi_n-\chi_e e_H}{1+\beta})-1)[(1+r)a+\frac{w'z'(1+\lambda)}{1+r'}]+\lambda\frac{w'z'}{1+r'}}{w\left(1+\mathbb{1}_{h>h_H}\lambda\right)}$$

$$\underline{z}^M_{part}(a,h) := \frac{(\exp(\frac{\chi_n}{1+\beta})-1)[(1+r)a+\frac{w'z'(1+\lambda)}{1+r'}]}{w\left(1+\mathbb{1}_{h>h_H}\lambda\right)}$$

$$\overline{z}^M_{part}(a,h) := \frac{\left\{\lambda\left[\exp(\frac{\chi_e e_L}{1+\beta})-1\right]^{-1}-1\right\}\frac{w'z'}{1+r'}-(1+r)a}{w\left(1+\mathbb{1}_{h>h_H}\lambda\right)}$$

$$\overline{z}^H_{non}(a) := \frac{(\exp(\frac{\chi_n}{1+\beta})-1)[(1+r)a+\frac{w'z'(1+\lambda)}{1+r'}]}{w(1+\lambda)}$$

A.2 Parameter restrictions for cutoffs ranking

To guarantee that $(n=0, e=e_H)$ dominates $(n=0, e=0)$, we need a lower bound for λ. The full-time learners prefer $(n=0, e=e_H)$ if and only if

$$(1+\beta)\ln c(n=0, e=e_H) - \chi_e e_H \geq (1+\beta)\ln c(n=0, e=0)$$

or equivalently:

$$\lambda \geq \underline{\lambda}_1 := \frac{(1+r)a+\frac{w'z'}{1+r'}}{\frac{w'z'}{1+r'}}\left(1-\frac{1}{\exp(\frac{\chi_e e_H}{1+\beta})}\right) \text{ if } h < h_M\frac{1}{1-\delta}$$

$$\lambda \geq \underline{\lambda}_3 := \frac{(1+r)a+\frac{w'z'}{1+r'}}{\frac{w'z'}{1+r'}}\left(\exp(\frac{\chi_e e_H}{1+\beta})-1\right) \text{ if } h \geq h_M\frac{1}{1-\delta}$$

To avoid $(n=1, e=e_L)$ from being a dominated choice, we need another lower bound for λ. To see it, recall that $(n=1, e=0)$ is better than $(n=1, e=e_L)$ if $z > \overline{z}_{part}$, and $(n=1, e=e_L)$ is better than $(n=0, e=e_L)$ if $z > \underline{z}_{part}$. $(n=1, e=e_L)$ is therefore the best choice over the interval $(\underline{z}_{part}, \overline{z}_{part})$. For such an interval to exist, it must be the case that when $z = \underline{z}_{part}$, $z < \overline{z}_{part}$. $z = \underline{z}_{part}$ means that the part-time learners are indifferent between $(n=1, e=e_L)$ and $(n=0, e=e_L)$ so that

$$(1+r)a + wzx(h) + \frac{w'z'}{1+r'} = \exp(\frac{\chi_n}{1+\beta})\left[(1+r)a+\frac{w'z'}{1+r'}\right] \text{ if } h < h_M\frac{1}{1-\delta}$$

$$(1+r)a + wzx(h) + \frac{w'z'(1+\lambda)}{1+r'} = \exp(\frac{\chi_n}{1+\beta})\left[(1+r)a+\frac{w'z'(1+\lambda)}{1+r'}\right] \text{ if } h \geq h_M\frac{1}{1-\delta}$$

For the part-time learners to prefer $(n=1, e=e_L)$ over $(n=1, e=0)$, we need

$$(1+\beta)\ln\frac{c(n=1, e=e_L)}{c(n=1, e=0)} \geq \chi_e e_L$$

If $h < h_M \frac{1}{1-\delta}$, inequality (A.2) is:

$$(1+\beta)\ln \frac{(1+r)a + wzx(h) + \frac{w'z'}{1+r'}}{(1+r)a + wzx(h) + \frac{w'z'(1-\lambda)}{1+r'}} \geq \chi_e e_L$$

Evaluating the left-hand-side at $z = \underline{z}_{part}$ yields:

$$\lambda \geq \underline{\lambda}_2 := \frac{(1+r)a + \frac{w'z'}{1+r'}}{\frac{w'z'}{1+r'}} \left(1 - \frac{1}{\exp(\frac{\chi_e e_L}{1+\beta})}\right) \exp(\frac{\chi_n}{1+\beta})$$

If $h > h_M \frac{1}{1-\delta}$, inequality (A.2) is:

$$(1+\beta)\ln \frac{(1+r)a + wzx(h) + \frac{w'z'(1+\lambda)}{1+r'}}{(1+r)a + wzx(h) + \frac{w'z'}{1+r'}} \geq \chi_e e_L$$

Evaluating the left-hand-side at $z = \underline{z}_{part}$ yields:

$$\lambda \geq \underline{\lambda}_4 := \frac{(1+r)a + \frac{w'z'}{1+r'}}{\frac{w'z'}{1+r'}} \frac{\left(\exp(\frac{\chi_e e_L}{1+\beta}) - 1\right)\exp(\frac{\chi_n}{1+\beta})}{\exp(\frac{\chi_e e_L}{1+\beta}) + \exp(\frac{\chi_n}{1+\beta}) - \exp(\frac{\chi_e e_L + \chi_n}{1+\beta})}$$

We have that $\underline{\lambda}_1 > \underline{\lambda}_2$ and $\underline{\lambda}_3 > \underline{\lambda}_4$ if

$$\exp(\frac{\chi_e e_H}{1+\beta}) > \frac{\exp(\frac{\chi_e e_L}{1+\beta})}{\exp(\frac{\chi_e e_L}{1+\beta}) + \exp(\frac{\chi_n}{1+\beta}) - \exp(\frac{\chi_e e_L + \chi_n}{1+\beta})}$$

Therefore, the inequality above implies that the conditions (A.2) and (A.2) are sufficient for the conditions (A.2) and (A.2). Furthermore, $\lambda_3 \geq \lambda_1$ so that the condition (A.2) is sufficient for the condition (A.2).

We can then conclude that the conditions (A.2) and (A.2) are sufficient for 1) the slower learners always prefers $(n = 0, e = e_H)$ over $(n = 0, e = 0)$, and 2) $\overline{z}_{part} > \underline{z}_{part}$, i.e., there exists state space where $(n = 1, e = e_L)$ is optimal.

A.3 Labor supply and human capital investment when y perfectly correlats with z

When the learning ability y is perfectly correlated with the idiosyncratic productivity z, the decision rule (n, e) is a function of states (z, h, a) instead of (z, h, a, y).

Figure A.1 illustrates the decision rule (n, e) for households with $h_M \leq h < h_M \frac{1}{1-\delta}$. The human capital h changes along the horizontal line and the idiosyncratic productivity z changes along the vertical line. The two diagonal lines, $\overline{z}_M(h)$ and $\underline{z}_M(h)$ are defined as:

$$\underline{z}_M(h) := \frac{h_M - (1-\delta)h}{e_H}; \qquad \overline{z}_M(h) := \frac{h_M - (1-\delta)h}{e_L}$$

They separate the state space into three areas: the unshaded area represents the non-learners, the lightly-shaded area represents the full-time learners, and the darkly-shaded area represents the part-time learners. The areas are divided by dashed horizontal lines associated with cutoffs $\overline{z}^L_{non}(a)$, $\overline{z}^L_{full}(a)$, $\underline{z}^L_{part}(a)$, and $\overline{z}^L_{part}(a)$ defined in (3.1.2), (3.1.2), (3.1.2), and (3.1.2).

This decision rule diagram is representative for households in other four ranges of human capital. Figure A.2 illustrates the regions in which households make positive human capital investments. Striped shading highlights where investment occurs, with dark areas denoting part-time learners and light areas representing full-time learners.

For households with $h < h_M$, $\overline{z}_M(h)$ and $\underline{z}_M(h)$ continue to be the boundaries that separate non-learners, full-time learners and part-time learners, but the four cutoffs are $\overline{z}^L_{non}\frac{1}{1-\lambda}$, $\overline{z}^L_{full}\frac{1}{1-\lambda}$, $\underline{z}^L_{part}\frac{1}{1-\lambda}$, and $\overline{z}^L_{part}\frac{1}{1-\lambda}$.

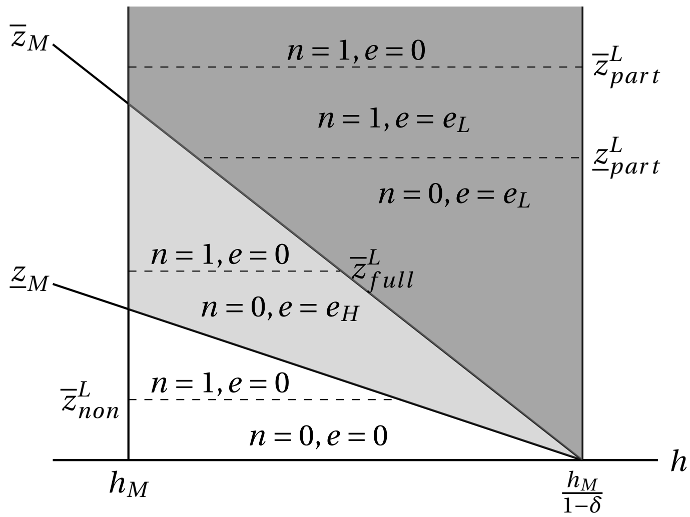

Figure A.1: Decision Rule Diagram for $h_M \le h < h_M(1-\delta)^{-1}$

The human capital h changes along the horizontal axis and idiosyncratic productivity z changes along the vertical axis. The two diagonal lines, $\overline{z}_M(h)$ and $\underline{z}_M(h)$, separate the state space into three areas: the unshaded area represents non-learners, the lightly shaded area full-time learners, and the darkly shaded area part-time learners. The areas are divided by four dashed horizontal lines associated with cutoffs $\overline{z}^L_{non}$, $\overline{z}^L_{full}$, $\underline{z}^L_{part}$, and $\overline{z}^L_{part}$, which are functions of asset holdings a.

For households with $h_M \frac{1}{1-\delta} \le h < h_H \frac{1}{1-\delta}$, the boundaries for state space division change to $\overline{z}_H(h)$ and $\underline{z}_H(h)$:

$$\underline{z}_H(h) := \frac{h_H - (1-\delta)h}{e_H}; \qquad \overline{z}_H(h) := \frac{h_H - (1-\delta)h}{e_L}$$

The four cutoffs of z are $\overline{z}^M_{non}$, $\overline{z}^M_{full}$, $\underline{z}^M_{part}$, and $\overline{z}^M_{part}$ defined in Appendix A.1

Households with $h \ge h_H \frac{1}{1-\delta}$ are non-learners, and have their decision exactly the same as in the benchmark model.

B Effects of baseline AI adoption in two-period model

The baseline AI shock is defined in Equation (2.3.2). In this section, we examine how this shock directly affects labor supply and saving through its effect on period-2 labor income.

B.1 Effects on labor supply

The AI adoption raises period-2 labor income for households who will work in the low or high sector, leading to a positive income effect that reduces their labor supply in period 1.

B.2 Effects on saving when human capital investment is unchanged

The effect of AI adoption on their saving can be analyzed as if we are varying the parameter t in the functions $\Delta(x, a; t)$, defined in Equation (3.3.1), and $\Delta_H(x, a; t)$, defined in Equation (3.3.2).

Proposition B.1. *$\Delta_H(x, a; t)$ is increasing in t. $\Delta(x, a; t)$ is convex in t:*

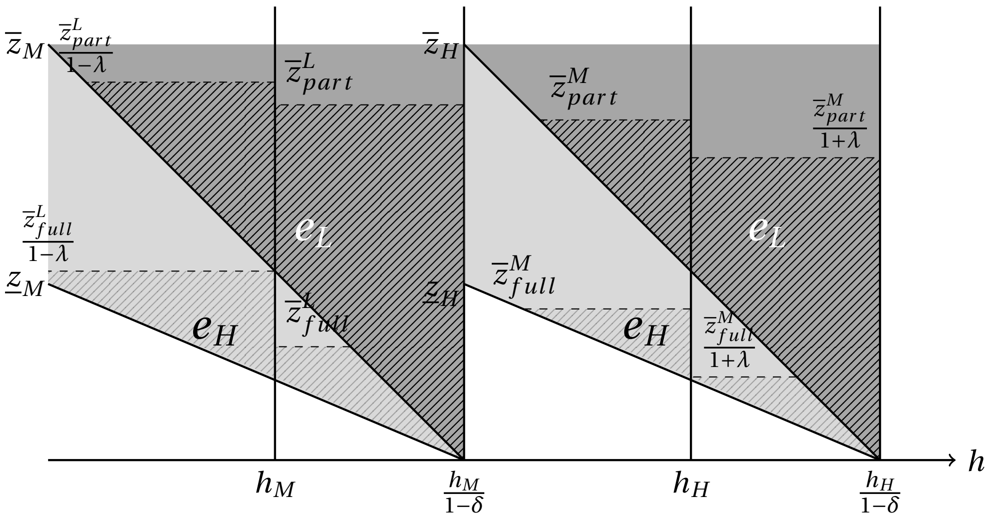

Figure A.2: State Space for Human Capital Investment

The darkly-shaded striped areas indicate the state space for human capital investment equal to e_L by the part-time learners. The lightly-shaded striped areas indicate the state space for human capital investment equal to e_H by the full-time learners.

- *If* $\Delta(x, a; t > 1) > 0$, $\Delta(x, a; t') > \Delta(x, a; t)$ *for* $t' > t > 1$.
- *If* $\Delta(x, a; t < 1) > 0$, $\Delta(x, a; t') < \Delta(x, a; t)$ *for* $1 > t' > t$.

Proof. See Appendix C. ■

Those who stay in the same sector For middle-sector households, the AI adoption leaves both their incomes and saving unchanged.

Low-sector and high-sector households experience an increase in period-2 labor income, x', as a result of the AI shock. If they can remain in the same sector via part-time training or without any additional human capital investment, the effect of the AI shock on their saving mirrors that of part-time training, $\Delta(x, a; t)$, since the AI shock raises x' (i.e., $t > 1$). As established in Proposition 3.2, $\Delta(x, a; t)$ has opposite signs for low-income and high-income households. Thus, the AI shock *crowds out* saving for low-sector households while *crowding in* saving for high-sector households.

For households who must undertake full-time training to remain in the high sector, $\Delta_H(x, a; t)$ captures the additional effect of such training on saving. In this case, a higher x'—brought about by the AI shock—corresponds to an increase in t, further boosting $\Delta_H(x, a; t)$ (Proposition B.1). Consequently, the AI shock *crowds in* saving for high-sector households in this scenario as well.

Those who upskill For low-sector households, saving behavior remains unchanged, as the AI shock does not affect their future productivity after upskilling.

For the middle-sector households who upskill via part-time training, their saving changes by $\Delta(x, a; t)$ with $t > 1$. The AI shock boosts their future productivity gain from λ to $(1+\gamma)\lambda$, which corresponds to increasing t to $t' > 1$. According to Proposition B.1, if the pre-shock effect of part-time training on saving is positive, the AI shock will *raise* saving. If this effect is negative, the overall impact of the AI shock on saving becomes ambiguous.

For the middle-sector households who upskill via full-time training, there is an *additional positive effect* of the AI shock on their saving, because a higher x' increases $\Delta_H(x, a; t)$ (Proposition B.1).

Those who downskill Downskilling is caused by human capital depreciation. For high-sector households who transition downward, the AI shock leaves their future productivity – and thus their saving – unchanged.

For middle-sector households who downskill to the low sector, their saving changes by $\Delta(x,a;t)$ with $t<1$. The AI shock mitigates their future productivity loss by reducing it from λ to $(1-\gamma)\lambda$, effectively increasing t to a new value $t'<1$. According to Proposition B.1, if the pre-shock effect $\Delta(x,a;t)$ is positive, the AI shock will *reduce* saving. If this effect is negative, however, the overall impact of the AI shock on saving is ambiguous.

C Proof of Proposition

C.1 Proof of Proposition 3.2

The derivative of saving with respect to t is

$$\frac{\partial a'^{\star}}{\partial t}(x,a;t) = -\frac{x\mu}{1+\beta} + \frac{x^2\Sigma}{\beta}\,\frac{t\,[2(x+a)+tx\mu]}{\left[(x+a)+tx\mu\right]^2}.$$

The total effect of part-time training on saving is

$$\Delta(x,a;t) = a'^{\star}(x,a;t) - a'^{\star}(x,a;1) = \int_1^t \frac{\partial a'^{\star}}{\partial u}(x,a;u)\,du.$$

Define

$$F(x,a;u) \equiv x\,\frac{u[2(x+a)+ux\mu]}{[(x+a)+ux\mu]^2}, \quad \overline{F}(x,a;t) \equiv \frac{1}{t-1}\int_1^t F(x,a;u)\,du.$$

Then equation Equation (C.1) can be written as

$$\Delta(x,a;t) = x(t-1)\Big[\frac{\Sigma}{\beta}\overline{F}(x,a;t) - \frac{\mu}{1+\beta}\Big].$$

Differentiating $F(x,a;u)$ with respect to x gives

$$\frac{\partial F(x,a;u)}{\partial x} = \frac{2u\,a\,(a+x)}{(a+(1+u\mu)x)^3} > 0,$$

so $\overline{F}(x,a;t)$ is strictly increasing in x.

The sign of $\Delta(x,a;t)$ is governed by

$$S(x,a;t) \equiv \frac{\Sigma}{\beta}\overline{F}(x,a;t) - \frac{\mu}{1+\beta}.$$

Because $\overline{F}(x,a;t)$ is strictly increasing, $S(x,a;t)$ increases monotonically with x.

For $x\to 0$, $F(x,a;u)\to 0$ and $\overline{F}(x,a;t)\to 0$ so that $S(x,a;t)\to -\frac{\mu}{1+\beta}<0$, implying $\Delta(x,a;t)<0$ for small x.

For $x\to\infty$, $F(x,a;u)\to \frac{u(2+u\mu)}{(1+u\mu)^2}$ and $\overline{F}(x,a;t)\to \overline{F}_\infty(t) \equiv \frac{1}{t-1}\int_1^t \frac{u(2+u\mu)}{(1+u\mu)^2}\,du$. If

$$\frac{\Sigma}{\mu} > \sigma^*(t) \equiv \frac{\beta}{1+\beta}\,\frac{1}{\overline{F}_\infty(t)}$$

then $S(x,a;t)>0$ for sufficiently large x, and hence $\Delta(x,a;t)>0$.

If idiosyncratic risk is large enough, i.e., condition (C.1) is satisfied, there exists a unique threshold $x^*(t)$ at which the sign flips:

$$\Delta(x,a;t)<0 \text{ for } x<x^*(a,t), \qquad \Delta(x,a;t)>0 \text{ for } x>x^*(a,t).$$

C.2 Proof of Lemma 3.1

Denote

$$G(x,a;t) \equiv \frac{t^2x^2}{(a+x+tx\mu)(a+tx\mu)}$$

the net additional effect of full-time training on saving can be rewritten as

$$\Delta_H(x,a;t) \equiv x\big[-\frac{\beta}{1+\beta}+\frac{\Sigma}{\beta}G(x,a;t)\big]$$

Differentiating $G(x,a;t)$ with respect to x gives

$$\frac{\partial G(x,a;t)}{\partial x} = \frac{t^2xa(2tx\mu+2a+x)}{(a+tx\mu)^2(a+x+tx\mu)^2} > 0,$$

so $G(x,a;t)$ is strictly increasing in x.

The limits of $G(x,a;t)$ are:

$$\begin{aligned} G(x,a;t) &\to 0 \quad (x\to 0), \\ G(x,a;t) &\to G_\infty(t) \equiv \frac{t}{\mu(1+t\mu)} \quad (x\to\infty), \end{aligned}$$

Therefore, $G(x,a;t) < G_\infty(t)$ for any x.

If

$$\frac{\Sigma}{\beta}G_\infty(t) < \frac{\beta}{1+\beta}, \text{ i.e. } \frac{\Sigma}{\mu} < \hat{\sigma}(t) \equiv \frac{\beta^2}{1+\beta}(\frac{1}{t}+\mu).$$

Then $\Delta_H(x,a;t) < x[-\frac{\beta}{1+\beta}+\frac{\Sigma}{\beta}G_\infty(t)] < 0$ for any x. Furthermore, with some tedious algebra, we can show that for any x

$$G(x,a;t)+x\frac{\partial G(x,a;t)}{\partial x} < G_\infty(t)$$

Hence, the inequality (C.2) also implies that

$$\frac{\partial \Delta_H(x,a;t)}{\partial x} = \frac{\Sigma}{\beta}[G(x,a;t)+x\frac{\partial G(x,a;t)}{\partial x}]-\frac{\beta}{1+\beta} < \frac{\Sigma}{\beta}G_\infty(t)-\frac{\beta}{1+\beta} < 0.$$

If

$$\frac{\Sigma}{\beta}G_\infty(t) > \frac{\beta}{1+\beta}, \text{ i.e. } \frac{\Sigma}{\mu} > \hat{\sigma}(t) \equiv \frac{\beta^2}{1+\beta}(\frac{1}{t}+\mu),$$

since $G(x,a;t)$ is strictly increasing in x, there exists a unique $\hat{x}(a,t)$ such that

$$\Delta_H(x,a;t) = x\big[-\frac{\beta}{1+\beta}+\frac{\Sigma}{\beta}G(x,a;t)\big] > 0 \Leftrightarrow x > \hat{x}(a,t)$$

Moreover, $\Delta_H(x,a;t) > 0$ implies that

$$\frac{\partial \Delta_H(x,a;t)}{\partial x} > 0.$$

C.3 Proof of Proposition 3.4

The relevant upper bounds of z for positive human capital investment are functions of γ (to the first order approximation):

$$\begin{aligned}
\bar{z}^L_{full}(a;\gamma) &= \bar{z}^L_{full}(a;\gamma=0) - \gamma\lambda\frac{w'z'}{w\left(1-\mathbb{1}_{h<h_M}\lambda\right)(1+r')} \\
\bar{z}^L_{part}(a;\gamma) &= \bar{z}^L_{part}(a;\gamma=0) - \gamma\lambda\frac{w'z'}{w\left(1-\mathbb{1}_{h<h_M}\lambda\right)(1+r')}\frac{\exp(\frac{\chi_e e_L}{1+\beta})}{\exp(\frac{\chi_e e_L}{1+\beta})-1} \\
\bar{z}^M_{full}(a;\gamma) &= \bar{z}^M_{full}(a;\gamma=0) + \gamma\lambda\frac{w'z'}{w\left(1+\mathbb{1}_{h>h_H}\lambda\right)(1+r')}\exp(\frac{\chi_n-\chi_e e_H}{1+\beta}) \\
\bar{z}^M_{part}(a;\gamma) &= \bar{z}^M_{part}(a;\gamma=0) + \gamma\lambda\frac{w'z'}{w\left(1+\mathbb{1}_{h>h_H}\lambda\right)(1+r')}\frac{1}{\exp(\frac{\chi_e e_L}{1+\beta})-1}
\end{aligned}$$

Therefore, an anticipated AI shock, $\gamma > 0$ makes those with $h < h_M \frac{1}{1-\delta}$ invest less human capital and those with $h > h_M \frac{1}{1-\delta}$ invest more human capital.

C.4 Proof of Proposition B.1

$$\Delta(x,a;t) = a'^{\star}(x,a;t) - a'^{\star}(x,a;1) = \int_1^t \frac{\partial a'^{\star}}{\partial u}(x,a;u)du.$$

differentiating with respect to t gives

$$\frac{d\Delta(x,a;t)}{dt} = \frac{\partial a'^{\star}}{\partial t}(x,a;t)$$

Since

$$\frac{\partial^2 a'^{*}(x,a;t)}{\partial t^2} = \frac{\partial}{\partial t}\left(-\frac{x\mu}{1+\beta} + \frac{x^2\Sigma}{\beta}\frac{t\,[2(x+a)+tx\mu]}{\left[(x+a)+tx\mu\right]^2}\right) = \frac{2x^2\Sigma(a+x)^2}{\beta(a+x+tx\mu)^3} > 0.$$

The slope $\frac{\partial a'^{\star}}{\partial t}(x,a;t)$ is strictly increasing in t. Hence $\Delta(x,a;t)$ is convex in t.

$$\Delta_H(x,a;t) = x\big[-\frac{\beta}{1+\beta} + \frac{\Sigma}{\beta}G(x,a;t)\big] \text{ with } G(x,a;t) = \frac{t^2x^2}{(a+x+tx\mu)(a+tx\mu)}$$

Differentiating $G(x,a;t)$ with respect to t gives

$$\frac{\partial G(x,a;t)}{\partial t} = \frac{tx^2(2a^2+2atx\mu+2ax+\mu tx^2)}{(a+tx\mu)^2(a+x+tx\mu)^2} > 0,$$

so $G(x,a;t)$ is strictly increasing in t, and so is $\Delta_H(x,a;t)$.

We now consider the comparison between $\Delta(x,a;t)$ and $\Delta(x,a;t')$ for $t' > t$. Given x and a, define

$$f(t) \equiv \frac{\partial a'^{*}}{\partial t}(x,a;t).$$

so $f'(t) > 0$, i.e. $f(t)$ is strictly increasing in t.

Case 1: $1 < t < t'$

Suppose $\Delta(x,a;t) > 0$. Then

$$\Delta(x,a;t) = \int_1^t f(u)\,du > 0.$$

Since f is increasing,

$$f(u) \le f(t) \quad \text{for all } u \in [1,t],$$

which implies

$$\Delta(x,a;t) = \int_1^t f(u)\,du \le (t-1)\,f(t).$$

Because $t > 1$, the inequality $\Delta(x,a;t) > 0$ forces $f(t) > 0$.

Now for any $t' > t$,

$$f(u) \ge f(t) > 0 \quad \text{for all } u \in [t,t'],$$

and therefore

$$\Delta(x,a;t') - \Delta(x,a;t) = \int_t^{t'} f(u)\,du > 0.$$

We then have that:

$$1 < t < t', \Delta(x,a;t) > 0 \quad \Longrightarrow \quad \Delta(x,a;t') > \Delta(x,a;t)$$

That is, once $\Delta(x,a;t)$ becomes positive for $t > 1$, it is strictly increasing in t thereafter.

Case 2: $t < t' < 1$

For $t < 1$,

$$\Delta(x,a;t) = \int_1^t f(u)\,du = -\int_t^1 f(u)\,du.$$

Suppose $\Delta(x,a;t) > 0$. Then

$$-\int_t^1 f(u)\,du > 0 \quad \Longrightarrow \quad \int_t^1 f(u)\,du < 0.$$

Since f is increasing

$$f(u) \geq f(t) \quad \text{for all } u \in [t,1],$$

which implies

$$\int_t^1 f(u)\,du \geq (1-t)\,f(t).$$

Because $t < 1$, the inequality $\Delta(x,a;t) > 0$ forces $f(t) < 0$.

Consider

$$\Delta(x,a;t') - \Delta(x,a;t) = \int_t^{t'} f(u)\,du$$

If $f(u) < 0$ for all $u \in [t,t']$, then $\int_t^{t'} f(u)\,du < 0$.

If there exists some $t_s \in [t,t']$ such that $f(t_s) = 0$, so $f(u) < 0$ for $u < t_s$ and $f(u) > 0$ for $u > t_s$. Then $f(u) > 0$ for $u \in [t',1]$. Hence,

$$\int_{t'}^1 f(u)\,du > 0$$

This implies that

$$\Delta(x,a;t') = -\int_{t'}^1 f(u)\,du < 0$$

Together with the inequality $\Delta(x,a;t) > 0$, we have that

$$\Delta(x,a;t') < \Delta(x,a;t)$$

We then have that

$$t < t' < 1,\ \Delta(x,a;t) > 0 \quad \Longrightarrow \quad \Delta(x,a;t') < \Delta(x,a;t).$$

Thus, for $t < 1$, whenever $\Delta(x,a;t) > 0$, increasing t toward 0 makes Δ strictly decrease.

D Computational Procedure for the Quantitative Model

D.1 Steady-state Equilibrium

In the steady-state, the measure of households, $\mu(a,h,z)$, and the factor prices are time-invariant. We find a time-invariant distribution μ. We compute the households' value functions and the decisions rules, and the time-invariant measure of the households. We take the following steps:

1. We choose the number of grid for the risk-free asset, a, human capital, h, and the idiosyncratic labor productivity, z. We set $N_a = 151$, $N_h = 151$, and $N_z = 11$ where N denotes the number of grid for each variable. To better incorporate the saving decisions of households near the borrowing constraint, we assign more points to the lower range of the asset and human capital.

2. Productivity z is equally distributed on the range $[-3\sigma_z/\sqrt{1-\rho_z^2}]$. As shown in the paper, we construct the transition probability matrix $\pi(z'|z)$ of the idiosyncratic labor productivity.

3. Given the values of parameters, we find the value functions for each state (a,h,z). We also obtain the decision rules: savings $a'(a,h,z)$, and $h'(a,h,z)$. The computation steps are as follow:

4. After obtaining the value functions and the decision rules, we compute the time-invariant distribution $\mu(a,h,z)$.

5. If the variables of interest are close to the targeted values, we have found the steady-state. If not, we choose the new parameters and redo the above steps.

D.2 Transition Dynamics

We incorporate the transition path from the status quo to the new steady state. We describe the steps below.

1. We obtain the initial steady state and the new steady state.

2. We assume that the economy arrives at the new steady state at time T. We set the T to 100. The unit of time is a year.

3. We initialize the capital-labor ratio $\{K_t/L_t\}_{t=2}^{T-1}$ and obtain the associated factor prices $\{r_t, w_t\}_{t=2}^{T-1}$.

4. As we know the value functions at time T, we can obtain the value functions and the decision rules in the transition path from $t = T-1$ to 1.

5. We compute the measures $\{\mu_t\}_{t=2}^{T}$ with the measures at the initial steady state and the decision rules in the transition path.

6. We obtain the aggregate variables in the transition path with the decision rules and the distribution measures.

7. We compare the assumed paths of capital and the effective labor with the updated ones. If the absolute difference between them in each period is close enough, we obtain the converged transition path. Otherwise, we assume new capital-labor ratio and go back to 3.

E Investigating the GE channel of AI's impact

Figures E.1a and E.1b compare the transition dynamics in the benchmark general-equilibrium model with those in a partial-equilibrium (PE) version of the model, where individual behavior responds to AI adoption but factor prices are held fixed at their initial steady-state values. The green vertical line marks the date of AI adoption.

On the aggregate side (Figure E.1a), both models deliver a long-run expansion in output, consumption, and investment after AI adoption. In the PE model, GDP responses are quite similar across the two models, but the composition of that response differs. In the PE model, consumption rises by less, while investment rises by more in the long run. The reason is that, in the benchmark, the long-run return on capital becomes negative (as shown in Figure 5), whereas in the PE model there is no such price effect. As a result, households in the PE environment have stronger incentives to save and invest, tilting the response toward investment rather than consumption. Even though aggregate human-capital dynamics do not differ much across the two environments, employment behaves very differently around the adoption date. In the PE model, employment rises sharply when AI is introduced because wages do not fall as employment increases.

Turning to inequality dynamics (Figure E.1b), the long-run behavior is similar across the two environments, but the impact responses differ markedly. As noted above, employment rises more on impact in the PE model even though output responses are similar. This implies that the additional employment mainly comes from low-productivity households. Consequently, the Gini coefficients for income, earnings, and consumption fall more on impact in the PE model but then move toward levels similar to those in the benchmark once job polarization and skill reallocation take hold. The human-capital Gini shows virtually no difference between the two models. By contrast, the wealth Gini exhibits very different transition dynamics. In the PE model, it displays a pronounced but short-lived spike early in the transition because poor households save less, as wages do not fall there in response to AI adoption, unlike in the benchmark economy. In the long run, however, the wealth Gini converges to a level similar to the benchmark, mainly because middle-sector households gradually increase their savings, as discussed in the main text.

Figure E.1: Transition Path of Aggregate Variables and Inequality Measures: Benchmark vs. PE Models

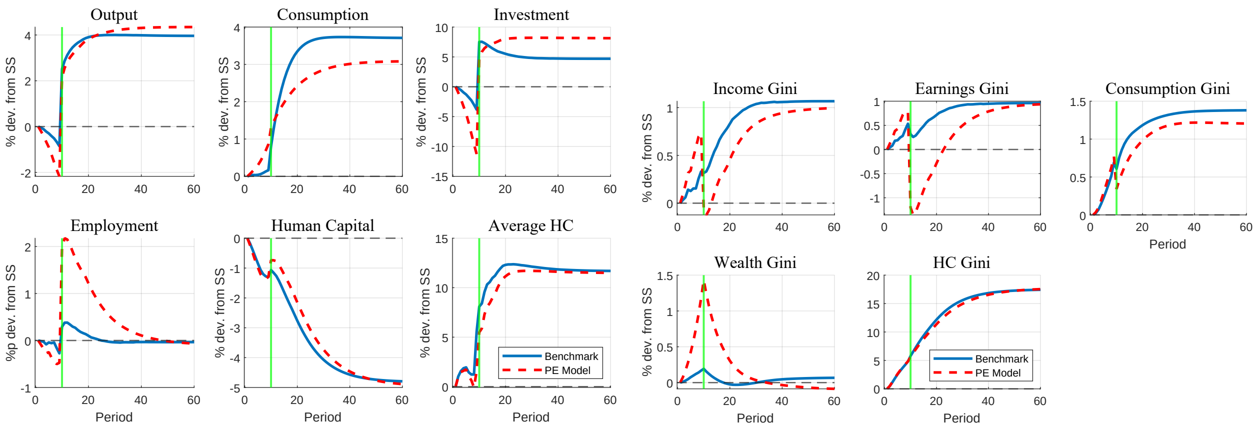

(a) Aggregate Variables

(b) Inequality Measures

Note: The left panel shows the transition paths of aggregate variables, and the right panel shows the transition paths of inequality measures, comparing the benchmark and PE models. The x-axis represents years, and the y-axis shows the percentage deviation from the initial steady state. AI introduction is assumed to occur in period 10. The PE model is an economy in which factor prices are held fixed at their initial steady-state values until the new steady state is reached.

Figure F.1: Transition Path with $\gamma = 0.2$: Benchmark vs. No HC Models

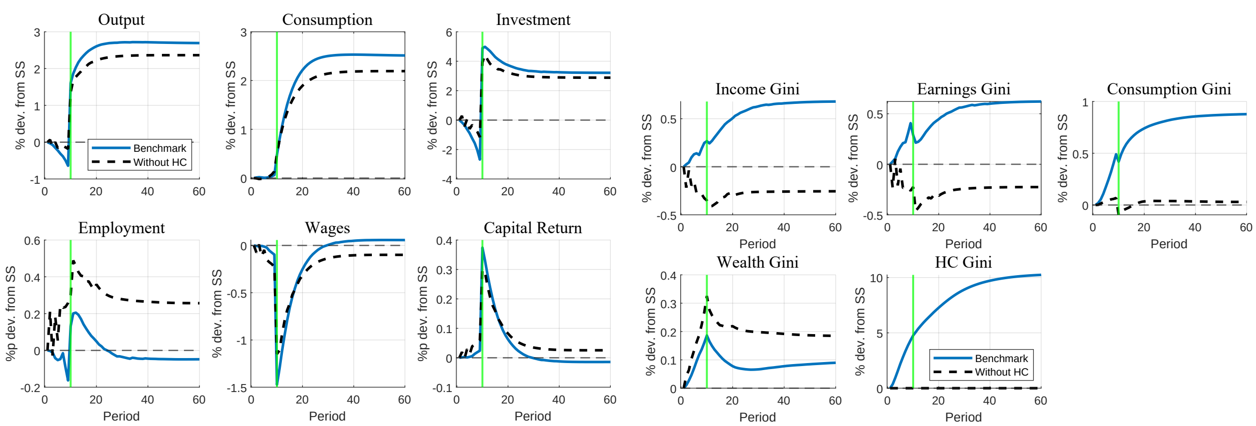

(a) Aggregate Variables

(b) Inequality Measures

Note: The left panel shows the transition paths of aggregate variables, and the right panel shows the transition paths of inequality measures, both for the benchmark and No HC models with $\gamma = 0.2$. The x-axis represents years, and the y-axis shows the percentage deviation from the initial steady state. AI introduction is assumed to occur in period 10. The No HC model is an economy in which workers maintain their initial steady-state level of human capital throughout the AI implementation until the new steady state is reached.

F Additional Figures and Tables

Figure F.2: Transition Path with $\gamma = 0.4$: Benchmark vs. No HC Models

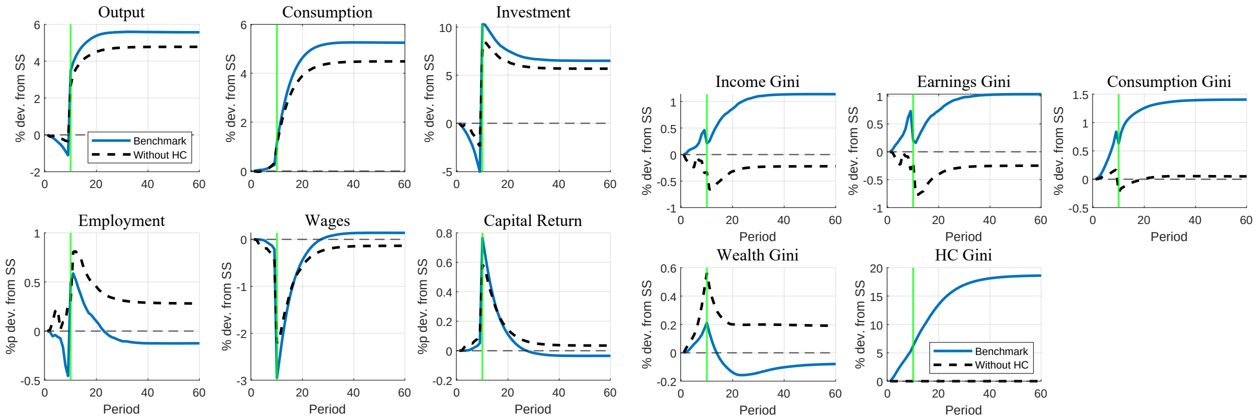

(a) Aggregate Variables

(b) Inequality Measures

Note: The left panel shows the transition paths of aggregate variables, and the right panel shows the transition paths of inequality measures, both for the benchmark and No HC models with $\gamma = 0.4$. The x-axis represents years, and the y-axis shows the percentage deviation from the initial steady state. AI introduction is assumed to occur in period 10. The No HC model is an economy in which workers maintain their initial steady-state level of human capital throughout the AI implementation until the new steady state is reached.

Figure F.3: Transition Path under Gradual AI Adoption: Benchmark vs. No HC Models

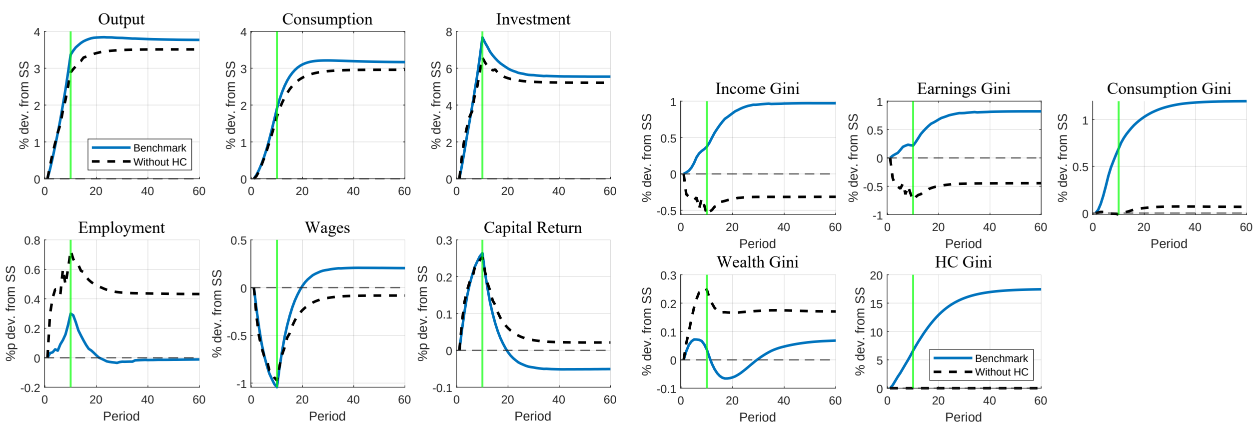

(a) Aggregate Variables

(b) Inequality Measures

Note: The left panel shows the transition paths of aggregate variables, and the right panel shows the transition paths of inequality measures, both for the benchmark and No HC models under gradual AI adoption. The x-axis represents years, and the y-axis shows the percentage deviation from the initial steady state. In this exercise, AI is introduced gradually from period 0 and fully implemented by period 10. The No HC model is an economy in which workers maintain their initial steady-state level of human capital throughout the AI implementation until the new steady state is reached.