

The Demographics of Technology Adoption and Productivity Growth*

Nils H. Lehr
International Monetary Fund
nlehr@imf.org

September 3, 2026

Young workers adopt new technologies at higher rates than old workers. I develop an overlapping-generations model linking this lifecycle gradient to technology adoption, wages, and innovation. Using U.S. commuting-zone variation and a historical-births instrument, I show that younger workforces expand computer-science employment and raise wages, especially for young workers. Exposure through trade to youth waves in foreign markets also increases R&D employment and patenting. Calibrated to these responses, the model attributes about one quarter of the late-1990s U.S. productivity acceleration to the baby boom's youth wave, partly recasting the IT boom as a demographic episode. Aging workforces therefore create a headwind to endogenous technology adoption and innovation.

JEL Codes: O33, J11

Keywords: Technology Adoption, Demographics, Productivity


* I benefited from advice from many including Stephen Terry, Pascual Restrepo, Tarek Hassan, David Lagakos, Daniele Paserman, Petia Topalova, Nan Li, Marti Mestieri, Huiyu Li, Robert King, and Susanto Basu as well as the useful comments and suggestions of participants in the Boston University Lunch Seminar, Green Line Macro Meeting, and Lecznar Memorial Lectures. Parts of this paper were previously circulated under the title "Aging, Technology Adoption, and Growth", which was first made publicly available in 2020, and "Innovation in an Aging Economy". The views expressed in this paper are those of the authors and do not necessarily represent those of the IMF, its Executive Board, or IMF management.

1. INTRODUCTION

Who adopts new technologies? Survey evidence on generative artificial intelligence shows a steep age gradient: workers in the youngest age group are almost three times as likely to use the technology as those in the oldest (Eurostat, 2026; Pew Research Center, 2025). The same gradient marked the arrival of the personal computer four decades earlier. Because productivity gains materialize only when technologies diffuse and are used, workforce age composition can shape productivity dynamics.

The timing of the U.S. demographic cycle and the IT boom is suggestive of a possible interplay between demographics and productivity growth. As the baby-boom cohorts moved through the age distribution, the share of workers aged 25–44 among those aged 25–64 rose from 53 percent in 1970 to nearly two-thirds around 1990, before falling to about one half by 2020 (Figure 1, panel a). TFP growth accelerated sharply in the following decade and subsequently slowed as the working-young share declined (panel b).

Figure 1: U.S. Productivity Growth Accelerated after a Youth Wave



Notes: The working-young share is the population aged 25–44 as a share of the population aged 25–64. Panel (b) reports average annualized TFP growth within each decade. Sources: U.S. Census/IPUMS, Penn World Table (Feenstra et al., 2015), and the utilization-adjusted TFP series of Fernald (2014), accessed via FRED.

This paper proposes a structural link between demographics and productivity dynamics through a technology-adoption channel and quantifies its contribution to the dot-com productivity boom. I develop an overlapping-generations model of technology adoption and semi-endogenous growth in which demographics operate at three levels. At the worker level, adopting a technology is a human-capital investment: workers pay a learning cost to supply labor on a technology, and young workers—with longer remain-

ing careers over which to recoup the cost—adopt at higher rates. At the regional level, technologies must also be made locally available through an upfront implementation investment undertaken by local firms and technology providers. Because the profitability of this investment rises with workers' adoption, it generates strategic complementarity that can amplify the direct effect of demographic composition. Finally, an aggregate innovation sector expands the technology frontier in response to adoption-driven profits. These channels distinguish two demographic forces. Population growth changes market size and therefore the long-run growth rate of the technology frontier, while age composition changes average technology adoption and accumulated experience. In a semi-endogenous growth environment, the composition effects shape the transition and productivity levels but not the long-run growth rate.

I provide evidence on these margins. First, CPS and Eurostat data document steep age gradients in workplace computer use and generative-AI adoption. Second, I study decadal changes across U.S. commuting zones from 1980 to 2010. Using historical local births as an instrument for the working-young share, I find that younger workforces experience faster growth in computer-science employment. A decomposition shows that this response occurs predominantly within age groups rather than through the mechanical shift toward age groups with higher baseline CS employment. The same instrumented variation raises wages for both age groups, with a larger response for young workers. Finally, commuting zones more exposed through trade to youth waves in foreign markets expand R&D employment and patenting, providing supporting evidence for the market-size channel.

I calibrate a quantitative extension of the model to these responses and to lifecycle profiles of wages and technology adoption. Feeding the observed U.S. demographic transition into the model raises annual TFP growth during 1995–2004 by 0.36 percentage points above its initial balanced-growth rate, accounting for about one quarter of the measured acceleration. The worker component contributes 0.26 pp and regional adoption contributes 0.09 pp. Accumulated experience shapes the timing: a youth wave initially lowers average experience, but this effect reverses as the large cohort ages.

Finally, I use the calibrated model to ask whether demographic aging as currently observed in the U.S. and other advanced economies (see Appendix F) changes the productivity response to an AI-like improvement in innovation and technology adoption. Across the demographic paths considered, the incremental TFP response is similar. Aging therefore lowers the underlying path of adoption-driven growth but has little effect on the gains from this particular technology shock.

This paper connects several strands of the literature. First, I contribute to the macroe-

economics of demographic change. Existing work has focused primarily on public finances, aggregate saving, and secular stagnation (Teulings and Baldwin, 2014; Eggertsson et al., 2019), while a smaller strand links age structure to aggregate productivity across countries (Feyrer, 2007; Aiyar et al., 2016), across U.S. states (Maestas et al., 2023), and in general-equilibrium models with age-specific research productivity or direction (Ang and Madsen, 2015; Aksoy et al., 2019; Irmen and Litina, 2022). A parallel literature argues that the slow-down in labor-force growth accompanying population aging has reduced firm entry and business dynamism (Hopenhayn et al., 2022; Engbom, 2023; Peters and Walsh, 2026; Karahan et al., 2024; Inokuma and Sánchez, 2025; Hoshi et al., 2025).¹ Closest to my empirical analysis, Derrien et al. (2023) find that younger workforces innovate more and attribute this relationship to comparative advantage. I instead provide instrumental-variable evidence on local technology adoption and complementary evidence on innovation through foreign-market exposure, consistent with a demand-side market-size channel. This interpretation also accords with evidence that firms reduce technology investment when young labor-market entrants become scarce (vom Baur, 2026) and that businesses in younger areas adopt new financial technologies faster (Crouzet et al., 2026). A related literature links aging to automation: labor scarcity induces firms to automate, potentially offsetting the output costs of workforce aging over the medium to long run (Acemoglu and Restrepo, 2022; Basso and Jimeno, 2021; Acemoglu et al., 2026). In my model, aging instead raises productivity only temporarily through accumulated experience and slows adoption-driven growth over longer horizons. The two mechanisms are complementary: an aging economy may automate more even as it adopts other learning-intensive technologies more slowly.

Second, the model connects endogenous growth, technology diffusion, and human-capital investment. Innovation expands the set of available technologies, as in expanding-variety and Schumpeterian models (Romer, 1990; Aghion and Howitt, 1992), while frontier growth follows the semi-endogenous structure of Jones (1995). The adoption side builds on work on technology diffusion (Comin and Hobijn, 2010; Comin and Mestieri, 2018; Argente et al., 2026) and, at the regional level, on Berkes et al. (2025), whose barriers to idea diffusion shape local growth. My regional block shares this local-diffusion structure but makes adoption demand responsive to demographic composition. At the household level, technology adoption is a human-capital investment with heterogeneous costs, in the spirit of Ben-Porath (1967) and Manuelli and Seshadri (2014). Closest to

¹Related work locates peak inventive and entrepreneurial productivity in middle age (Jones, 2010; Jones and Weinberg, 2011; Azoulay et al., 2020; Liang et al., 2018), potentially implying a non-monotonic relationship between age and innovation.

this mechanism, [Adão et al. \(2024\)](#) show that technology transitions slow when skills are technology-specific because the skill distribution adjusts as generations turn over, while [Hassan et al. \(2026\)](#) show that new technologies are initially skill-intensive and standardize as they age. Relatedly, [Xu \(2025\)](#) studies a cohort-specific complementarity in which workers entering after a technology wave are better matched to new technologies, emphasizing implications for new-firm growth, interest rates, and asset prices. I instead emphasize workers' remaining horizons as the source of generational differences in adoption and link those differences to regional diffusion and the demand for innovation. In this respect, the model also contributes to the market-size and directed-technical-change literatures, in which the size and composition of demand shape the rate and direction of innovation ([Acemoglu, 2002](#); [Acemoglu and Linn, 2004](#); [Acemoglu, 2010](#); [Gancia and Zilibotti, 2005](#)). Because younger workers adopt new technologies at higher rates, workforce age composition provides a specific and measurable source of variation in the market for innovation.

Finally, I speak to the literature on the IT productivity boom and its aftermath. The late-1990s resurgence in U.S. productivity growth is typically attributed to the production and diffusion of information technology ([Oliner et al., 2007](#); [Brynjolfsson et al., 2021](#)); I identify a demographic contribution to this episode. The subsequent slowdown has been linked to declining research productivity and business dynamism ([Syverson, 2017](#); [Bloom et al., 2020](#); [Akcigit and Ates, 2021](#); [Lehr, 2025b](#)) and, more broadly, to slowing population growth ([Jones, 2022](#)). I show that demographics contributed materially to the IT boom, while demographic aging need not substantially alter the productivity response to an AI-like technology shock.

2. A MODEL OF DEMOGRAPHICS AND TECHNOLOGY

I develop an overlapping-generations framework of technology adoption and endogenous innovation in which demographic shifts propagate through three levels. At the *household* level, young and old workers decide which of the available technologies to adopt; young workers, who reap the labor-income gains over a longer remaining horizon, adopt at a higher rate. At the *regional* level, I summarize implementation by local firms and technology providers with a surplus problem that determines how much of the frontier becomes locally available. As the returns to implementation depend on household adoption, this mechanism gives rise to strategic complementarity. At the *aggregate* level, an innovation sector expands the technology frontier in response to

aggregate adoption-driven profits. The model thus links demographics to productivity dynamics.

2.1. SETUP

Time is discrete and indexed by t . The economy consists of a unit mass of regions, each populated by a mass L_t of workers split across two cohorts: a mass L_{yt} of young workers and L_{ot} of old workers. Regions share an integrated market for the homogeneous final good and a common technology frontier, while technology availability and labor markets are local. I abstract from migration across regions. I first describe household and regional adoption in a representative region and then the aggregate innovation sector that determines the common frontier.

Final goods production. A representative firm hires technology-specific labor $l_t(m)$ and intermediate goods $x_t(m)$ to produce final output Y_t using the set M_t of locally adopted technologies, indexed by m :²

$$Y_t = \int_{M_t} (\tilde{\zeta}(m) l_t(m))^\alpha x_t(m)^{1-\alpha} dm. \quad (1)$$

The multiplier $\tilde{\zeta}(m)$ depends on when the region adopted technology m : $\tilde{\zeta}(m) = 1$ for technologies adopted in the current period ($m \in M_t^N$) and $\tilde{\zeta}(m) = \tilde{\zeta} \geq 1$ for technologies adopted in prior periods ($m \in M_{t-1}$). The vintage productivity premium $\tilde{\zeta} > 1$ for old technologies captures the rise in effective productivity over time as input prices decline and complementary intangibles accumulate (Nordhaus, 2007; Brynjolfsson et al., 2021). The firm takes wages $W_t(m)$ and intermediate good prices $P_t(m)$ as given and solves

$$\max_{\{l_t(m), x_t(m)\}_{m \in M_t}} Y_t - \int_{M_t} (P_t(m) x_t(m) + W_t(m) l_t(m)) dm. \quad (2)$$

Intermediate goods. Each adopted technology $m \in M_t$ is associated with a monopolist who produces the intermediate good at unit cost ψ in terms of the final good. The monopolist sets the price $P_t(m)$ taking the demand schedule implied by the final-goods

²The expanding-variety formulation provides a transparent benchmark for the analytical results. Section 4 allows innovation to improve and replace existing varieties. With an age-independent replacement hazard and a fresh adoption decision for an improved version, the lifecycle and regional comparative statics continue to operate, while replacement changes the vintage distribution, accumulated experience, and quantitative magnitude of the responses.

firm's problem (2) as given, and solves

$$\pi_t(m) = \max_{x_t(m)} P_t(x_t(m)) x_t(m) - \psi x_t(m). \quad (3)$$

Adoption costs and labor supply. For each household i and technology m , adoption costs are Cc_{im} , where $C > 0$ is a common scaling parameter and c_{im} is drawn from $F(c) = c^{1/\eta}$ on $[0, 1]$.³ Thus, $1/\eta$ is the elasticity of the adoption rate with respect to the cost cutoff. Draws are independent across households and technologies and remain fixed over a household's life: a technology not adopted when young carries the same cost in old age. Costs are paid in units of the final good and capture both direct learning expenses and the opportunity cost of time devoted to learning.⁴

The cost of adopting a set of technologies is the integral of the household's own draws. With a slight abuse of notation, I write A_{it} for both the set of technologies a household adopts and its measure. For a young household, who enters without prior adoption, the cost of adopting $A_{it} \subseteq M_t$ is

$$C_{it}(A_{it}) = C \int_{A_{it}} c_{im} dm. \quad (4)$$

An old household enters with a stock A_{it-1} of previously adopted technologies and adopts an additional set $A_{it}^N \subseteq M_t \setminus A_{it-1}$, drawn from the technologies it has not yet adopted—both newly available technologies in M_t^N and old technologies in M_{t-1} it did not adopt while young—at cost

$$C_{it}(A_{it}^N) = C \int_{A_{it}^N} c_{im} dm. \quad (5)$$

Each adopted technology yields ω_m efficiency units of labor at wage $W_t(m)$, with $\omega_m = 1$ in the adoption period and $\omega_m = \omega \geq 1$ thereafter. The parameter ω is an exogenous technology-specific experience-productivity multiplier; it is distinct from the equilibrium wage per efficiency unit. To economize on notation, I suppress time and household subscripts on these multipliers; their values depend on a technology's adoption vintage and, for ω_m , on the household's time since adoption.

³Households and workers coincide: each household supplies one worker's labor and makes that worker's adoption decisions.

⁴The distribution captures heterogeneity in the household's marginal cost of acquiring technology-specific skills, in the spirit of Ben-Porath (1967): some technologies are intuitive, while others require sustained learning. The setup is related to Manuelli and Seshadri (2014), who study human-capital investment across countries and model adoption costs partly in units of final output. In my framework, households invest in human capital by adopting technologies.

Young households. Young households live for two periods, young and old. They discount the future at rate β and face mortality risk δ_t at the end of their young period. To insure against mortality, they trade a state-contingent bond B_{it+1} paying $\tilde{R}_{t+1} \equiv R_{t+1}/(1 - \delta_t)$ units of the final good in case of survival and zero otherwise, such that the expected return equals the risk-free rate.⁵ In the young period, the household chooses a set of technologies $A_{it} \subseteq M_t$ to adopt, consumption C_{it} , and bond holdings B_{it+1} . In the old period, conditional on survival, the household chooses an additional set $A_{it+1}^N \subseteq M_{t+1} \setminus A_{it}$ and consumes the remaining income. The young household solves

$$\begin{aligned} & \max_{A_{it}, A_{it+1}^N, B_{it+1}} \ln C_{it} + \beta (1 - \delta_t) \ln C_{it+1}, \\ \text{s.t. } & C_{it} + B_{it+1} = \int_{A_{it}} W_t(m) dm - C_{it}(A_{it}), \\ & C_{it+1} = \tilde{R}_{t+1} B_{it+1} + \int_{A_{it} \cup A_{it+1}^N} \omega_m W_{t+1}(m) dm - C_{it+1}(A_{it+1}^N), \end{aligned} \quad (6)$$

where $C_{it}(\cdot)$ denotes the cost function from (4)–(5).

Old households. Old households live for one period. They enter with a stock A_{it-1} of adopted technologies and bond holdings B_{it} inherited from their young-period decisions. They choose an additional set $A_{it}^N \subseteq M_t \setminus A_{it-1}$ of technologies to adopt and consume the remaining income. The old household solves

$$\begin{aligned} & \max_{A_{it}^N} \ln C_{it}, \\ & C_{it} = \tilde{R}_t B_{it} + \int_{A_{it-1} \cup A_{it}^N} \omega_m W_t(m) dm - C_{it}(A_{it}^N), \end{aligned} \quad (7)$$

where $C_{it}(\cdot)$ is the adoption cost function (5).

Regional technology adoption. I represent decentralized implementation by local firms and technology providers with a regional adoption problem. At each point in time, the region inherits a stock M_{t-1} of previously adopted technologies and chooses the mass $M_t^N \in [0, M_t - M_{t-1}]$ of newly adopted technologies, where M_t is the aggregate technology frontier. Each newly adopted technology generates one period of monopoly profits $\pi_{N,t} \equiv \pi_t(m)$ for $m \in M_t^N$, with $\pi_t(m)$ given by (3) and requires an infrastructure investment paid once per technology at the time of adoption.⁶ In the following period,

⁵The state-contingent bond ties the return to mortality risk. Without this link, adoption would be insensitive to δ_t because households could perfectly insure against mortality via the bond. Similar dependencies arise under binding borrowing constraints or bequest motives.

⁶The regional stage introduces complementarities: each worker's adoption raises the value of making the technology available locally, and the region-level adoption decision internalizes these spillovers.

the technology enters the inherited stock and no longer earns rents. The one-period duration is a tractable version of the temporary monopoly rents in Schumpeterian growth models (Aghion and Howitt, 1992); it prevents current adoption from depending on an additional stream of monopoly rents in later periods. The marginal cost of adoption at level a is

$$\tilde{\kappa}_t(a) = L_t \cdot \kappa \cdot \left(\frac{a}{M_{t-1}} \right)^{1/\phi} \cdot \left(\frac{M_{t-1}}{M_{t-1}} \right)^{\theta/\phi}. \quad (8)$$

The elasticity ϕ governs the response of regional adoption to profitability, while θ governs its response to distance from the frontier. Scaling costs by L_t makes the adoption rate independent of regional size, in the spirit of Young (1998) and Howitt (1999). Increasing marginal costs in the adoption flow capture congestion, while the distance term makes the remaining technologies more costly to adopt near the frontier. The latter ensures a stable long-run distance between regional technology and the frontier. The region chooses M_t^N to maximize the surplus from new adoption,

$$\max_{M_t^N \in [0, M_t - M_{t-1}]} \pi_{N,t} \cdot M_t^N - \int_0^{M_t^N} \tilde{\kappa}_t(a) da, \quad (9)$$

yielding the marginal implementation condition

$$\pi_{N,t} = \tilde{\kappa}_t(M_t^N) \quad (10)$$

at the marginal adopted technology. Because $\tilde{\kappa}_t$ is increasing in a , inframarginal technologies are adopted at strictly lower cost than the marginal one, and the region retains a positive surplus that partly flows to the innovation sector.⁷

Innovation sector. The aggregate innovation sector expands the technology frontier M_t by producing $M_t^N \equiv M_t - M_{t-1}$ new technologies each period using final output and subject to the cost function

$$\mathcal{K}_t(M_t^N) = \chi \cdot (M_t^N)^{1+1/\varphi} \cdot M_{t-1}^{(\vartheta-1)/\varphi}, \quad (11)$$

where $\varphi > 0$ is the elasticity of frontier expansion to profitability, $\vartheta > 0$ governs how innovation costs vary with the inherited frontier, and $\chi > 0$ is a level shifter. Rising

Alvarez et al. (2023) document such strategic complementarities in technology adoption directly in the context of peer-to-peer digital payments.

⁷The regional problem is a reduced-form aggregate of decentralized local investments needed to make technologies usable, including installation, training, systems integration, and complementary infrastructure. The increasing cost schedule orders implementation projects by difficulty; local firms and providers undertake projects until the return to the marginal project equals its implementation cost.

marginal cost in M_t^N captures congestion in the innovation production process, and the dependence on M_{t-1} captures intertemporal fishing-out as the frontier expands as in Jones (1995). Each new technology is sold to regional adopters, and the innovation sector captures the inframarginal adoption surplus. Under the marginal-cost function (8) this surplus equals a fraction $1/(1 + \phi)$ of gross monopoly profits per adopted technology. Denote by $\mathbb{E}_t[\cdot]$ the cross-sectional average across regions, by $\pi_{N,t}$ the profits per regionally adopted technology, and by $a_{N,t}^R \equiv M_t^N / (M_t^N + (M_{t-1} - M_{t-1}))$ the average adoption rate of newly available technologies.⁸ Then the innovation sector's revenue per new technology is given by

$$\Pi_t^I = \frac{1}{1 + \phi} \mathbb{E}_t[a_{N,t}^R \pi_{N,t}]. \quad (12)$$

It chooses M_t^N to maximize total profit,

$$\max_{M_t^N} \Pi_t^I \cdot M_t^N - \mathcal{K}_t(M_t^N). \quad (13)$$

The growth rate of the frontier is $g_t^* \equiv M_t^N / M_{t-1}$.

Demographics. Each period, a mass $(1 + n_t)L_{yt-1}$ of young households is born, of which a fraction $(1 - \delta_t)$ survives into the old period. The share of young workers in the region's labor force is then

$$\ell_{yt} \equiv \frac{L_{yt}}{L_{yt} + L_{ot}} = \frac{1 + n_t}{2 + n_t - \delta_{t-1}}, \quad (14)$$

and the total labor force evolves according to

$$\frac{L_{t+1}}{L_t} = (1 + n_t) \cdot \frac{2 + n_{t+1} - \delta_t}{2 + n_t - \delta_{t-1}}. \quad (15)$$

2.2. COMPETITIVE EQUILIBRIUM

I focus on an equilibrium with a constant, exogenous gross risk-free rate R to isolate the mechanism of interest.

Definition 1 (Competitive equilibrium). *Given the gross risk-free rate R , demographic primitives $\{n_t, \delta_t\}_{t \geq 0}$, an initial regional technology stock M_{-1} , an initial global frontier M_{-1} , and initial population sizes $\{L_{y0}, L_{o0}\}$, a competitive equilibrium consists of*

⁸The denominator captures the assumption that regional adoption is random across non-adopted technologies.

1. household allocations $\{C_{yt}, C_{ot}, B_{yt+1}, A_{yt}, A_{ot}^N\}_{t \geq 0}$,
2. intermediate-good and labor allocations $\{l_t(m), x_t(m)\}_{m \in M_t, t \geq 0}$,
3. regional adoption $\{M_t^N\}_{t \geq 0}$ and frontier expansion $\{\mathcal{M}_t^N\}_{t \geq 0}$,
4. prices $\{W_t(m), P_t(m)\}_{m \in M_t, t \geq 0}$ and profits $\{\pi_t(m)\}_{m \in M_t, t \geq 0}$,

such that

1. the final-goods firm solves (2) taking $W_t(m)$ and $P_t(m)$ as given,
2. the intermediate-good monopolist for each $m \in M_t$ solves (3) taking final-good demand as given,
3. households solve (6) and (7) taking prices and profits as given,
4. the regional adoption decision satisfies (9)–(10),
5. the innovation sector solves (13) taking profits from new inventions Π_t^I as given,
6. the law of motion $M_t = M_{t-1} + M_t^N$ holds, labor markets clear within each adopted technology, and the frontier evolves according to $\mathcal{M}_t = \mathcal{M}_{t-1} + M_t^N$.

Definition 2 (Balanced growth path). A balanced growth path (BGP) is a competitive equilibrium in which demographic primitives are constant ($n_t = n$, $\delta_t = \delta$ for all t), and the regional technology stock M_t and global frontier $\mathcal{M}_t$ grow at constant common rate $g^* \equiv M_t/M_{t-1} - 1 = \mathcal{M}_t/\mathcal{M}_{t-1} - 1$.

Before stating the main results, I derive the key equilibrium relationships. In a symmetric equilibrium, perfect substitution equalizes returns across technologies, so household adoption follows a reservation-cost rule. Normalizing $\psi = 1 - \alpha$ and defining $\omega \equiv \alpha(1-\alpha)^{(1-\alpha)/\alpha}$ as the symmetric per-technology wage, the old- and young-household adoption rates absent vintage effects ($\zeta = 1$) are⁹

$$a_{ot} = \left(\frac{\omega}{C}\right)^{1/\eta} \quad \text{and} \quad a_{yt} = \left(\frac{\omega}{C}\right)^{1/\eta} \left(1 + \frac{1 - \delta_t}{R} \omega\right)^{1/\eta}. \quad (16)$$

Young households adopt at a higher rate because, beyond the current-period wage, they also collect income $\omega \omega$ on the technology in old age, adjusted for the survival-adjusted discount factor $(1 - \delta_t)/R$.

⁹Vintage effects enter asymmetrically across technology generations: for new technologies, they multiply the experience term ω ; for old technologies, they multiply $\omega \omega$. Throughout, I assume an interior optimum, with adoption rates below one.

At the regional level, the marginal implementation condition (10) delivers the regional technology growth rate

$$g_t^R \equiv \frac{M_t^N}{M_{t-1}} = \left(\frac{\pi_{N,t}}{L_t \kappa} \right)^\phi \left(\frac{M_{t-1}}{M_{t-1}} \right)^\theta. \quad (17)$$

Since $\pi_{N,t} = \alpha(1 - \alpha)^{1/\alpha} L_t \bar{a}_t^N$, where $\bar{a}_t^N$ is the average household-level adoption rate of new technologies, regional scale cancels: the adoption rate is driven by the average household adoption rate of newly available technologies, $\bar{a}_t^N = \ell_{yt} a_{yt} + (1 - \ell_{yt}) a_{ot}$, and by the distance to the frontier M_{t-1}/M_{t-1} , but not by the size of the region. I assume $\theta > 0$ throughout, so that a region further from the frontier adopts faster and therefore catches up.

At the aggregate level, the innovation sector's optimality condition yields the frontier growth rate

$$g_t^* = \left(\frac{\Pi_t^I}{\chi(1 + 1/\varphi)} \right)^\varphi M_{t-1}^{-\vartheta}, \quad (18)$$

where revenue Π_t^I rises in aggregate adoption and in market size through $\pi_{N,t} \propto L_t$. The fishing-out elasticity $\vartheta > 0$ anchors the level of the frontier relative to market size.

Along the BGP, the labor force grows at rate n , and the frontier growth rate is pinned down by the population growth rate and the elasticities of the innovation cost function as in Jones (1995):

$$(1 + g^*) = (1 + n)^{\varphi/\vartheta}. \quad (19)$$

2.3. MAIN RESULTS

I now characterize the equilibrium response of technology adoption, wages, productivity, and innovation to demographic shifts. All statements assume that elasticities $\{\eta, \phi, \theta, \varphi, \vartheta\}$ are strictly positive and that regions operate below the technology frontier.¹⁰ Proofs are provided in Appendix A.

The first result summarizes the lifecycle mechanism at the household level. Young workers adopt more because they recover learning costs over a longer horizon; mortality narrows this advantage.

Lemma 1 (Individual technology adoption). *Young workers adopt new technologies at a higher rate than old workers, $a_{yt} > a_{ot}$, reflecting the longer remaining horizon over*

¹⁰This is an interiority condition: some frontier technologies remain available for local implementation, so the upper bound on M_t^N does not bind. At the frontier, the cap on technologies available for adoption can truncate the local comparative statics.

which young workers reap the labor-income gains from adoption. The gap is larger when mortality δ_t is lower. Both cohorts adopt the same share of old technologies along the BGP.

The lifecycle gradient in individual adoption translates into a regional response through the profitability of new technologies. A larger working-young share raises household adoption of newly available technologies and therefore effective labor input and profits on new intermediate goods. The regional marginal implementation condition then implies faster adoption of the frontier and a higher regional technology level. Greater technology availability raises wages for both cohorts, but more for young workers because they adopt more new technologies and do not yet benefit from experience with old ones.

Proposition 1 (Regional adoption and wages). *Starting on a BGP, consider a region with a temporarily higher working-young share ℓ_{yt} driven by a one-time spike in population growth n_t at time t . Then, at that time, the region has a higher average household-level adoption rate, which raises the profitability of newly adopted technologies and induces faster regional adoption of the global frontier: g_t^R rises. The associated rise in local technology availability raises both cohorts' wages, and raises young wages by more.*

The aggregated regional responses then shape the development of new technologies. The frontier-growth response to a younger workforce induced by faster population growth has a permanent component and a transitory one. Permanently, higher population growth raises the market size for each technology and, through the innovation sector's optimality condition, the balanced-growth rate of frontier expansion, as in Jones (1995). On impact, a higher working-young share further raises household and regional adoption, lifting innovation-sector revenue and frontier growth by more than the permanent effect; this amplification decays as the wave passes.

Proposition 2 (Aggregate innovation). *A rise in population growth raises the growth rate of the technology frontier through two distinct effects. On the balanced growth path the effect is permanent and operates through market size, with elasticity $\partial \ln(1 + g^*) / \partial \ln(1 + n) = \varphi / \vartheta > 0$. On impact, a rise in the working-young share additionally raises frontier growth through higher household adoption $\bar{a}_t^N$ and the induced rise in regional adoption g_t^R , each lifting innovation-sector revenue.*

The implications for productivity growth reflect these channels jointly and can be nuanced. Because the model abstracts from physical capital and has a constant intermediate-input share, model TFP is proportional to value added per worker. Its

growth can be written as the sum of regional technology growth and growth in a human-capital-adjusted average adoption rate:

$$\Delta \ln TFP_t \approx g_t^R + \Delta \ln \tilde{a}_t \quad \text{with} \quad \tilde{a}_t = \ell_{yt} a_y + (1 - \ell_{yt}) \left(\frac{\omega}{1 + g_t^R} a_y + \frac{g_t^R}{1 + g_t^R} a_o \right).$$

The second term includes both technology adoption and productivity gains from technology-specific experience. Its mapping to measured TFP therefore depends on the labor-quality adjustment in the data.¹¹

On impact, higher household adoption of new technologies raises regional adoption. Frontier growth also increases through both higher adoption and the pure market-size effect of a larger population.¹² By contrast, the human-capital-adjusted average adoption rate initially declines whenever old workers earn more than young workers in cross section, as they do in the data: a younger workforce has less accumulated experience. This composition effect reverses as the larger cohort ages, generating a delayed but transitory increase in TFP growth. If technologies become more productive as they mature, $\zeta > 1$ provides an additional source of delayed growth.¹³

3. EMPIRICAL EVIDENCE

This section provides empirical evidence on four model predictions. First, using the 1989 CPS Computer Use Supplement and 2025 Eurostat data, I document lifecycle gradients in workplace computer adoption and generative AI adoption, consistent with Lemma 1. Second, consistent with Proposition 1, I show that U.S. commuting zones (CZs) with a higher working-young share (WYS) experienced faster growth in computer-science employment. Third, the same CZs experienced faster wage growth, particularly among young workers, as the proposition also predicts. Finally, I show that CZs more exposed through trade to youth waves in foreign markets expanded their innovation activity,

¹¹The Penn World Table human-capital index is based primarily on schooling and returns to education, so technology-specific experience remains in its residual. Fernald (2014)'s series removes observable labor-composition changes based on education and experience; part of the model's worker component may therefore be recorded as labor-quality growth rather than Fernald TFP, although conventional experience controls need not capture technology-specific learning fully.

¹²The relative strength of the regional and frontier responses determines whether M_t/M_t rises or falls and therefore shapes subsequent regional-adoption dynamics.

¹³These dynamics echo Solow's observation that the computer age was visible everywhere except in the productivity statistics. In the short run, young workers' advantage in technology adoption is offset by their lower accumulated experience. As experience recovers, high adoption can translate into delayed TFP growth.

consistent with Proposition 2.

3.1. DATA

My analysis links a range of data sources across time and space. I mostly focus on the United States and, for the cross-sectional exercises, use 1990 commuting zones (CZs) as the unit of geographic analysis. CZs, developed by [Tolbert and Sizer \(1996\)](#), are designed to capture local labor markets and have become the standard geography in the local-labor-markets literature ([Autor and Dorn, 2013a,b](#)). Unless otherwise noted, I map geographies to CZs using the crosswalk in [Autor and Dorn \(2013b\)](#).

Employment, occupations, and wages. I construct employment- and population-based measures from the 1980, 1990, and 2000 decadal Censuses and the 2010–2012 three-year American Community Survey from IPUMS ([Ruggles et al., 2020](#)). I focus on full-time full-year (FTFY) workers, defined as those reporting at least 40 weeks of work in the prior year and at least 35 weekly hours, following [Acemoglu and Autor \(2011\)](#). Occupations are mapped to the consistent codes developed in [Autor and Dorn \(2013a\)](#); I define the computer science occupation using their codes 64–65 and 229–233. For wage outcomes, I construct both unadjusted average log hourly wages and a composition-adjusted variant obtained by reweighting the CZ-year worker distribution to the national 1980 distribution over age, gender, race, and education ([DiNardo et al., 1996](#)). The composition-adjusted variant isolates changes in wages net of shifts in workforce composition. Wages by age group are constructed analogously.

Innovation. I measure local innovation along two margins. Innovation inputs are proxied by the share of FTFY workers in R&D occupations, defined as natural scientists (codes 68–83), social scientists (166–169), computer scientists (64–65, 229–233), and engineers (44–59). Labor accounts for approximately two-thirds of total R&D cost according to the NSF Business R&D and Innovation Survey, so this share is a reasonable proxy for local R&D investment intensity [Lehr \(2025a\)](#). Innovation outputs are measured using citation-weighted patents per 1,000 residents, as in [Terry et al. \(2026\)](#), based on the patent data of [Berkes \(2018\)](#). Patents are mapped to CZs via the inventor’s county of residence and recorded in the application year. Citation weights are constructed relative to patents in the same five-year window and CPC subsection.

Historical births. Historical births by county from 1900 onward come from historical Censuses, the NBER digitization of the U.S. Vital Statistics, and the Surveillance, Epidemiology, and End Results (SEER) program. I map them to CZs using the crosswalks in [Eckert et al. \(2020\)](#) and [Autor and Dorn \(2013a\)](#).

International trade and foreign demographics. For the international exercise in Section 3.6, I further rely on industry-level U.S. exports from the Census Foreign-Trade Statistics for 2000, local industry employment from the 2010–2012 County Business Patterns, and population by age and country from the United Nations World Population Prospects.

Technology adoption and calibration. I measure workplace computer adoption using the 1989 CPS Computer Use Supplement (Flood et al., 2025) and generative AI adoption using Eurostat (Eurostat, 2026). The calibration additionally uses TFP estimates from the Penn World Table and Fernald (2014), together with the OECD Trade in Value Added statistics (OECD, 2025).

3.2. EMPIRICAL STRATEGY

The CZ-level exercises in Sections 3.4 and 3.5 estimate the response of local outcomes to changes in the working-young share in a decadal first-difference specification:

$$\Delta Y_{gt} = \alpha_g + \gamma_t + \beta \Delta WYS_{gt} + \varepsilon_{gt}, \quad (20)$$

where g indexes CZs, t indexes decades, and Δ denotes the 10-year change. The first-difference structure eliminates time-invariant unobserved heterogeneity across CZs; the further inclusion of CZ fixed effects flexibly absorbs CZ-specific linear trends. I weight observations by the 1980 working-age population and cluster standard errors at the state level, following Autor and Dorn (2013b).

The working-young share is defined as the population age 25–44 as a fraction of the population age 25–64:

$$WYS_{gt} = \frac{\sum_{a=25}^{44} \text{Pop}_{g,a,t}}{\sum_{a=25}^{64} \text{Pop}_{g,a,t}} \times 100. \quad (21)$$

OLS estimates of β may reflect short-run mobility responses to local conditions: productive or innovative CZs may attract young workers, while rising local costs may repel them. I therefore construct a predicted WYS using long-dated variation in local births, following related demographic instruments in Shimer (2001); Engbom (2023); Acemoglu and Restrepo (2022); Derrien et al. (2023); Karahan et al. (2024). For each CZ g and year t , I define the predicted age- a population as

$$\widehat{\text{Pop}}_{g,a,t} = \frac{\text{Births}_{g,t-a}}{\sum_{g'} \text{Births}_{g',t-a}} \times \sum_{g'} \text{Pop}_{g',a,t}, \quad (22)$$

i.e., the share of national births occurring in CZ g approximately a years before t , scaled by the current national population of age a . The predicted WYS, $\widehat{WYS}_{gt}$, is constructed from (22) using the same age aggregation as in (21). Intuitively, $\widehat{WYS}_{gt}$ captures the WYS the CZ would have had if residents never moved and were subject to national-average mortality.

Identifying assumption. Conditional on year and CZ fixed effects, changes in $\widehat{WYS}_{gt}$ affect local outcomes only through changes in the local WYS.

The instrument directly addresses contemporaneous mobility responses but does not by itself rule out correlation between long-run fertility patterns and persistent CZ characteristics whose effects evolve differentially over time. CZ fixed effects in the main specification absorb time-invariant CZ-level confounders, but not such time-varying effects. To probe this concern, Appendix D.2 reports a complementary specification that drops CZ fixed effects and instead controls for year effects and 1980 pre-characteristics: the college, non-white, working-young, and female shares of the population, the metropolitan share, and the working-age population.

International exercise. The fourth exercise (Section 3.6) uses cross-country variation in the working-young share of U.S. trading partners as a shifter of foreign demand for U.S. innovation. The regression is then estimated by OLS rather than IV because the demographic composition of U.S. trading partners is plausibly exogenous to local U.S. innovation activity. I describe this design at the start of Section 3.6.


3.3. WORKER TECHNOLOGY ADOPTION

The model predicts that young workers adopt newly available technologies at higher rates than old workers. I document this gradient using workplace computer use in the 1989 CPS Computer Use Supplement and generative AI adoption in the 2025 Eurostat data. The CPS measures both whether a worker uses a computer and the number of tasks performed with it conditional on use. Their product forms a combined adoption index. For generative AI, I report overall and workplace adoption for the EU15.

Figure 2 plots both adoption rates. Computer adoption at work declines monotonically with age across the working life. The pattern is robust to residualizing for gender, education, and occupation: the lifecycle gradient is not driven by cohort differences in education or by sorting into more or less computer-intensive occupations.¹⁴ Appendix

¹⁴The same pattern arises in the later waves of the CPS-CIU with a flatter gradient, as predicted by the model: as the technology gets older, older workers adopted it when young.

Figure 2: Young Workers Are More Likely To Adopt the New Technologies



Notes: Left panel reports computer use at work across age for the U.S. from the 1989 wave of the CPS Computer Use Supplement. “Adjusted” residualizes adoption rates for gender, education, and occupation. The adoption rate is measured as the product of extensive and normalized intensive margin adoption. The latter is measured as the share of tasks for which a computer was employed and normalized to an average of 1. Right panel reports data on adoption of generative AI across age groups for EU15 countries from Eurostat for 2025.

Table D.1 documents related sectoral heterogeneity: workplace computer use is more prevalent in service-sector employment, and its young–old gradient is more positive there. This pattern suggests a potential interaction with the demand-driven shift of innovation toward services emphasized by Comin et al. (2025). Generative AI adoption, arguably a more nascent technology in 2025 than the computer was in 1989, features an even steeper age profile. Relative to the oldest group, the youngest age group is about 2.7 times as likely to use generative AI at all and at work.¹⁵

3.4. REGIONAL ADOPTION OF COMPUTING

Proposition 1 predicts faster adoption of new technologies in local labor markets with younger workforces. I proxy regional computer adoption with the change in computer-science employment as a share of CZ employment over 1980–2010. Computer-science occupations are a natural occupational complement to the individual workplace adoption

¹⁵Existing literature similarly documents a declining age profile in computer and AI adoption during the respective periods (Friedberg, 2003; Weinberg, 2004; Aubert et al., 2006; Bick et al., 2024; Humlum and Vestergaard, 2025; Pew Research Center, 2025).

documented above.¹⁶ This proxy does not measure computer use directly: it captures complementary technical labor associated with implementing and maintaining computing technology locally, and it may partly reflect the local production or relocation of computing services rather than local use.

Table 1: Workforce Aging and Computer Science Employment – Decomposition

	(1)	(2)	(3)	(4)
	Δ Computer science employment			
ΔWYS	0.199*** (0.040)	0.172*** (0.039)	0.032*** (0.006)	-0.005 (0.004)
Component	Total	Within	Across	Cross
Year FE	✓	✓	✓	✓
CZ FE	✓	✓	✓	✓
First-stage F	37.2	37.2	37.2	37.2
Observations	2,166	2,166	2,166	2,166

Note: All regressions are estimated by 2SLS, instrumenting ΔWYS with ΔWYS, the historical-fertility instrument described in the text. Reported first-stage statistics are Kleibergen–Paap robust *F*-statistics. CZ observations are weighted by 1980 working-age population. Standard errors are clustered at the state level. See the text for variable descriptions.

Standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Table 1 reports 2SLS estimates of (20). Column 1 reports the total effect: a one-percentage-point increase in the local WYS raises the CZ’s CS employment share by 0.20 percentage points, statistically significant at the 1% level.¹⁷ The Kleibergen–Paap robust first-stage *F*-statistic is 37.2.

To distinguish a behavioral adoption response from a mechanical composition effect, columns 2–4 decompose the total response. Writing the CS employment share as $CS_{gt} = \sum_a s_{g,a,t} r_{g,a,t}$, where $s_{g,a,t}$ is the workforce share in age bin a and $r_{g,a,t}$ is the age- a CS share, the decadal change satisfies

$$\Delta CS_{gt} = \underbrace{\sum_a s_{g,a,t-1} \Delta r_{g,a,t}}_{\text{Within}} + \underbrace{\sum_a \Delta s_{g,a,t} r_{g,a,t-1}}_{\text{Across}} + \underbrace{\sum_a \Delta s_{g,a,t} \Delta r_{g,a,t}}_{\text{Cross}}. \quad (23)$$

¹⁶Beaudry et al. (2010) document substantial geographic variation in PC adoption and relate it to the local supply of skilled labor; I instead examine the role of age composition. See also Kalyani et al. (2025) for complementary evidence on technology diffusion.

¹⁷Appendix Table D.2 reports specifications with year fixed effects alone and with controls for 1980 characteristics. The corresponding coefficients are 0.137 and 0.166 and remain statistically significant. OLS estimates are close to zero and insignificant.

The Within component captures changes in age-specific CS shares holding the CZ's age composition fixed; the Across component captures the effect of shifts in age composition holding age-specific CS shares fixed; the Cross term captures the residual interaction.

The Within component carries 86% of the total response (0.17 of 0.20 in columns 2 and 1, respectively). The Across component is small (0.03), and the Cross term is statistically indistinguishable from zero. The adoption response to a younger workforce is therefore not driven by a mechanical shift toward age groups with higher baseline CS shares: it operates predominantly within age cells, consistent with greater local computer implementation and use rather than a shift toward high-adoption age groups.

3.5. WAGES

Proposition 1 further predicts that a rise in the local WYS raises wages for all workers and disproportionately benefits young workers, since the gains from adopting new technologies accrue more strongly to workers with the longest remaining horizon in the labor force. Table 2 reports estimates of (20) with the change in composition-adjusted (log) wages as the dependent variable.¹⁸

Columns 1 and 2 report the response of the average wage. A one-percentage-point increase in ΔWYS raises wages by 0.11 log points by OLS and 0.19 log points by 2SLS, both statistically significant at the 1% level. The IV estimate exceeds the OLS estimate, consistent with contemporaneous mobility responses attenuating the OLS coefficient.

The remaining columns split the response by age group. Wages of young workers rise nearly twice as much as wages of old workers: 0.23 vs. 0.12 by IV, with the same ordering by OLS (0.12 vs. 0.09). The signs and the relative magnitude are consistent with Proposition 1: a rise in the WYS raises wages for all workers and disproportionately benefits young workers, with the cross-sectional experience premium declining.

Appendix Table D.4 shows that the estimates are stable to winsorizing wages, capping population weights, and reweighting on age and education alone. The population-weighted baseline targets the response for the average worker and the aggregate economy, whereas the unweighted specification targets the average CZ. The wage response is insignificant in the latter, suggesting that the effect is concentrated in larger CZs. Appendix Table D.5 shows that contemporaneous changes in demographic composition and working-age population do not absorb the response, although the estimates become somewhat smaller. Appendix Table D.6 obtains qualitatively similar results using

¹⁸The composition-adjusted variant holds the local distribution over age, gender, race, and education fixed at the national 1980 distribution.

Table 2: Workforce Aging and Wages

	(1)	(2)	(3)	(4)
A. OLS Estimates	Δ Wages			
Δ WYS	0.089*** (0.021)	0.108*** (0.024)	0.124*** (0.028)	0.086*** (0.022)
B. IV Estimates	Δ Wages			
Δ WYS	0.142*** (0.049)	0.191*** (0.050)	0.231*** (0.059)	0.117** (0.047)
Workers		All	Young	Old
Comp.-Adj.		✓	✓	✓
First-stage F	37.2	37.2	37.2	37.2
Observations	2,166	2,166	2,166	2,166

Note: Panel A reports OLS estimates; Panel B reports 2SLS estimates instrumenting Δ WYS with the historical-fertility instrument Δ WYS. Reported first-stage statistics are Kleibergen–Paap robust F -statistics. Wages in columns (2)–(4) are composition-adjusted (DFL-reweighted to 1980) and grouped by worker age (all; young, 25–44; old, 45–64). All regressions include CZ and year fixed effects. CZ observations are weighted by 1980 working-age population. Standard errors are clustered at the state level. See the text for variable descriptions.

Standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

alternative wage adjustments, including Mincer residualization.

3.6. INNOVATION VIA TRADE

The fourth and final exercise examines an empirical analogue of Proposition 2. In the model, a younger customer base adopts more technology, raising the profits from developing new technologies. The regression asks whether this market-size logic extends to foreign customers of U.S. producers; it is therefore supporting evidence rather than a literal regression counterpart to the baseline domestic model. The exercise differs from the previous three in two ways. First, the marginal value of an invention likely depends on the worldwide market for the resulting technology, not just the local CZ. The relevant demographic structure is therefore that of the market for innovation, which can extend far beyond the CZ in which inventive activity occurs.¹⁹ Second, because foreign demographics are exogenous to local U.S. innovation activity at decadal frequency, identification does not require instrumentation.

¹⁹Relatedly, Cai and Stoyanov (2016) and Gu and Stoyanov (2019) document that population age structure shapes countries' comparative advantage and trade patterns.

Table 3: Workforce Aging via Trade and Innovation

	(1)	(2)	(3)
	Δ Innovation outcome		
ΔWYS^{Trade}	0.100*** (0.034)	0.118*** (0.015)	0.027*** (0.006)
Outcome	R&D emp.	Cite-wt. patents	Raw patents
Observations	2,166	2,166	2,166

Note: OLS estimates with ΔWYS^{Trade} , the change in the working-young share of U.S. export markets, using 2000 country shares within U.S. industries and 2010–12 CZ industry-employment weights. Column (1) reports R&D employment as a share of the working-age population. Columns (2)–(3) report patenting per 1,000 working-age population; column (2) weights patents by 5-year forward citations within CPC subsection. All regressions control for CZ and year fixed effects. CZ observations weighted by 1980 working-age population. Standard errors clustered at state level. See text for variable descriptions.

Standard errors in parentheses. Significance levels: * 10% , ** 5%, *** 1%.

For each CZ g at decade t , I construct a trade-weighted WYS as

$$\Delta WYS_{gt}^{Trade} = \sum_k \omega_{gk} \Delta WYS_{k,t}, \quad (24)$$

where $\Delta WYS_{k,t}$ is the decadal change in the working-young share of country k , and ω_{gk} is the share of country k in CZ g 's industry-weighted exports as of 2010–12. The weights ω_{gk} combine the industry composition of CZ g 's 2010–12 employment with the country composition of U.S. industry-level exports in 2000. Intuitively, ΔWYS_{gt}^{Trade} is large in CZs that specialize in industries whose export destinations are experiencing a youth wave.

Table 3 reports OLS estimates. A one-percentage-point increase in the trade-weighted WYS raises R&D employment by 0.10 percentage points of the working-age population, citation-weighted patents by 0.12 per 1,000 working-age residents, and raw patents by 0.03. Thus, exposure to younger foreign markets predicts expansion in both innovation inputs and outputs, consistent with the model's market-size channel. Because the exposure measure combines foreign demographics with CZ industry composition, the design cannot fully rule out differential trends in industries serving younger markets; I therefore interpret the evidence as supporting rather than standalone causal evidence.

4. DEMOGRAPHICS AND PRODUCTIVITY DYNAMICS

This section quantifies the effect of demographic shifts on productivity. I first extend and calibrate the model and then feed the U.S. demographic transition surrounding the IT boom into it. Appendix C.1 provides the full quantitative model.

To match the decadal frequency of the data, I first extend the model to feature four working generations (age groups 25–34 through 55–64). Age groups 1–3 face a forward-looking problem as young households in the baseline model, but with varying horizons, while age group 4 operates as the old generation.

Second, I allow for imperfect substitution across technologies in final production via a CES production function with an elasticity of substitution σ :

$$Y_t = Q_t^{-\frac{1}{\sigma-1}} \left(\int_{M_t} q_{jt} (l_t(j)^\alpha x_t(j)^{1-\alpha})^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}}, \quad (25)$$

where $q_{jt} \geq 1$ counts the product lines of variety j —which grows through improvements, as described next— $Q_t = \int q_{jt} dj$ is the total measure of products, and labor and intermediates are hired per product line. The normalization by Q_t removes mechanical love-of-variety effects, and $\sigma \rightarrow \infty$ recovers the linear aggregator of Section 2. With imperfect substitution, the scarcity of labor supply in new technology vintages yields a wage premium—fewer workers have adopted them, and none yet carry experience premia on them—and household adoption decisions interact through equilibrium vintage prices. This endogenous vintage wage premium is distinct from the calibrated experience-productivity multipliers ω_s .

Third, technologies depreciate through replacement. Each innovation event either creates a new variety or improves an existing one. An improvement adds a product line but renders the previous version obsolete, requiring regions and workers to adopt the improved version before earning income from it. I denote the resulting replacement hazard by ρ . In the baseline, improvements generate fresh adoption-cost draws and reset experience premia. Replacement shifts the technology portfolio toward recent vintages and therefore raises the share of wages tied to newly adopted technologies, an important determinant of the wage response to demographic shocks.

4.1. CALIBRATION

I calibrate the model through a combination of externally chosen parameters and internal moment matching. Table 4 reports both sets, together with the moments and their

model counterparts. Section 5 discusses results for alternative calibration targets and parameters.

Externally set parameters. The gross decadal risk-free rate is $R = 1.5$, in line with the average real return on long-term U.S. government bonds over 1980–2010 (Jordà et al., 2019).²⁰ The labor share in final production is $\alpha = 0.2$, implying a markup in the intermediate sector of 25%. The survival rates $\tilde{\delta}_s$ are taken from period life tables, and the effective labor-supply transitions δ_s combine survival with the age profile of labor-force participation and hours from the Census (Appendix B.1). The elasticity of the regional adoption rate to distance from the frontier is set to $\theta = 1$. In the model, the implied per-decade rate of β -convergence in regional technology is $\lambda = \theta(g^* + \rho)/(1 + g^*) \approx 0.18$ per decade, or about 2% per year, matching the canonical historical estimate of Barro and Sala-i Martin (1992).²¹ The improvement share of innovation events is set to $s_I = 0.73$, matching the share of productivity growth attributable to improvements of existing products—own-product improvements plus creative destruction—in Klenow and Li (2021).²² On the balanced growth path this implies a technology replacement rate of $\rho = s_I g^*/(1 - s_I) \approx 14\%$ per decade. The vintage productivity premium is set to $\xi = 1.15$, so that mature technologies are 15% more productive than newly adopted ones.²³ The elasticity of substitution across technologies σ is calibrated so that the newest technology vintage commands a wage premium—the log wage premium over the second-newest vintage—of 10%, in line with the evidence on wage premia from early computer and AI adoption in Krueger (1993) and Alekseeva et al. (2021).²⁴ Finally, the

²⁰Holding this return fixed during the demographic transition excludes aging-induced changes in saving and equilibrium interest rates. A lower return would raise the discounted future payoff to young adopters and could partially offset the adoption headwind. The common return also abstracts from historical differences in local credit conditions, which could generate spatial variation in adoption incentives.

²¹The adoption rule (17) makes the adoption flow respond to distance with elasticity θ ; since the flow replaces only a fraction $(g^* + \rho)/(1 + g^*)$ of the technology stock per decade, the stock gap closes at rate $\lambda = \theta(g^* + \rho)/(1 + g^*)$ per decade. Section 5 reports robustness to $\theta = 0.5$ and $\theta = 2$.

²²Klenow and Li (2021) attribute about 60% of growth to incumbents improving their own products and 13% to creative destruction, while Garcia-Macia et al. (2019) attribute 87–95% of productivity growth to improvements of existing products.

²³This magnitude is consistent with within-vintage productivity gains that accrue only gradually after adoption: Brynjolfsson et al. (2021) estimate that unmeasured IT-related intangible investment raises measured productivity by about 16%, and mature technologies additionally face more competitive entry, compressing markups over the vintage’s lifecycle. The markup channel alone could rationalize substantially larger values: fully eliminating the monopoly markup of $1/(1 - \alpha) = 1.25$ would raise per-technology output by a factor of $(1 - \alpha)^{-(1-\alpha)/\alpha} \approx 2.44$.

²⁴DiNardo and Pischke (1997) suggest a lower premium for computers after accounting for unobservable worker skills. Kalyani et al. (2025) find that job postings associated with newly emerging technologies pay a large premium, which is partly accounted for by worker skills. Finally, Deming and Noray (2020) find that occupations with rapidly evolving skills pay a 20% premium, but attribute it partly to compensation for rapid skill obsolescence. Section 5 reports robustness to $\sigma = 13$, implying a premium of roughly 20%,

innovation-cost level χ is normalized to one.

Internally calibrated parameters. I organize the internally calibrated parameters into five blocks. First, the steady-state entry rate n matches the working-young share of 0.53 in 1970. Second, the experience premia ω_s , learning-cost elasticity η , and learning-cost level C match the lifecycle profiles of relative wages and AI adoption in 2025. Third, the regional adoption elasticity ϕ matches the young- and old-worker wage-response semi-elasticities in Table 2, after accounting for the difference between average wage and TFP growth. I obtain these model moments from the exact simulated wage-growth response to a one-time cohort shock in the multigenerational model, rather than equating the empirical coefficients mechanically to an analytical derivative. Fourth, κ normalizes the balanced-growth regional adoption rate to $M/\mathcal{M} = 0.5$.

Finally, φ matches the response of R&D employment to trade-weighted demographics. I convert the 0.10-percentage-point response in Table 3 into a semi-elasticity using the sample-average R&D employment share of 2.5% and account for the 10% foreign share of demand for U.S. innovation. The mapping first simulates a foreign market's general-equilibrium adoption-demand response to a cohort shock and then applies the innovation-sector optimality condition. Given φ , the balanced-growth condition $(1 + g^*) = (1 + n)^{\varphi/\vartheta}$ identifies ϑ from the long-run TFP growth target of 0.5% per year. Appendix C.3 derives the corresponding first-order formulas in the two-generation model.

Calibration fit. Table 4 and Figure 3 report the calibration targets and their model counterparts. The model matches the targeted moments closely. The calibrated regional adoption elasticity ϕ is large, indicating strong strategic complementarity between household and regional adoption. For a given empirical innovation response, a larger ϕ requires a smaller frontier-profit elasticity φ , because demographics already generate a larger regional-adoption response. The overall innovation response depends approximately on their product, although obsolescence and endogenous vintage wage premia modify this relationship in general equilibrium.


and to the linear limit $\sigma = \infty$ without any premium.

Table 4: Calibration and Targeted Moments

A. Parameters	Symbol	Value	Method
<i>A.1. External</i>			
Interest rate	R	1.50	External
Profit share	α	0.20	External
Survival (effective)	δ	{1.00, 1.04, 0.95, 0.67}	External
Survival (headcount)	$\tilde{\delta}$	{1.00, 0.98, 0.97, 0.93}	External
Regional convergence	θ	1.00	External
Improvement share	s_I	0.73	External
Vintage premium	ζ	1.15	External
R&D cost (normalized)	χ	1.00	Normalization
<i>A.2. Internal</i>			
Population Growth	n	0.03	Moment matching
Experience premium	ω	{1.40, 1.83, 2.16}	Moment matching
Adoption cost elasticity	η	1.10	Moment matching
Adoption cost shifter	ω/C	0.18	Moment matching
Substitution elasticity	σ	23.9	Moment matching
Regional revenue elasticity	ϕ	6.35	Moment matching
Regional adoption cost	κ	0.05	Moment matching
Contemp. R&D returns	φ	0.13	Moment matching
Fishing-out	ϑ	0.08	Moment matching
<i>B. Moments</i>			
	Target	Model	Source
TFP Growth Rate	0.05	0.05	Fernald (2014)
Working Young Share	0.53	0.53	Census
Wage Elasticity			Table D.5
Young	0.23	0.24	
Old	0.12	0.10	
Novel technology premium	0.10	0.10	Targeted
R&D semi-elasticity	0.04	0.04	Table 3
Regional adoption rate	0.50	0.50	Normalization

Notes: Table reports calibrated parameters and targeted moments. Parameters in Panel A.1 are externally chosen, while those in Panel A.2 target specific moments. The novel technology premium is the implied log wage premium of the newest technology vintage over the second-newest; it motivates the choice of σ and is not targeted. All values rounded to two digits.

Figure 3: Calibration Fit: Lifecycle Wages and Technology Adoption




Notes: Panel A compares relative wage income across age groups in the model and the data (potential-experience profile, census). Panel B compares technology adoption rates across age groups in the model with AI adoption rates from Eurostat.

4.2. DEMOGRAPHIC TRANSITIONS AND PRODUCTIVITY GROWTH

To understand the importance of demographics for productivity in the model, I explore the U.S. dynamics in the 1990s, as depicted in Figure 1 as a case study. I feed the observed demographic transition into the model: the entry-rate path n_t is calibrated so that the model's working-young share replicates the U.S. working-young share over 1970–2020 and stabilizes thereafter. Figure 4 reports the experiment. Panel B confirms that the model replicates the targeted transition, which, as per Panel C, induces a temporary rise in TFP growth as in the data, although the peak in 2000 is not as pronounced. The latter finding is unsurprising as technological forces themselves likely contributed to the productivity boom around the dot-com boom. In the model, the demographic shift raises technology adoption rates for new technologies, which induces regional adoption. The latter partly offsets the initial drag on TFP growth coming from the demographic composition effect. At the peak in the decade around the year 2000, workforce composition makes a direct contribution to TFP growth, as experience recovers, together with regional adoption. Afterwards, regional adoption declines strongly, driven by a smaller share of young workers, while experience composition effects maintain productivity growth overall. In the medium run, lower population growth dominates and TFP growth declines below the initial BGP level.

Table 5 quantifies the decomposition. In the baseline calibration, the demographic boom raises annual TFP growth by 0.36 percentage points above the initial balanced-

Figure 4: The Demographic Boom and the IT Productivity Boom



Notes: The entry-rate path is calibrated to replicate the U.S. working-young share (Panel B). TFP data in Panel C are from the Penn World Table (Feenstra et al., 2015) and scaled to have 0.5% productivity growth in the initial period. Panel D reports household-level technology rates of new and all technologies. The latter reflects jointly the age distribution of technologies and their generation-specific adoption rates. Panel E reports the model's TFP growth decomposition relative to the initial balanced growth path. Panel F reports the growth rate of the technology frontier.

growth rate during 1995–2004, with 0.09 pp operating through regional adoption and 0.26 pp through workforce composition. Measured TFP growth accelerated by 1.44 pp, implying that the model accounts for approximately one quarter of the acceleration. Averaged over 1975–2025, the model generates an annual acceleration of 0.07 pp, compared with 0.32–0.43 pp in the data, implying a demographic contribution of approximately 20% over the longer interval.

5. DISCUSSION

This section examines the mechanisms behind the quantitative results, their robustness to alternative calibrations, and the responses to demographic and technology shocks.

Mechanisms. Online Appendix Figure E.2 traces the boom through model variants that shut down individual channels—generational adoption differences, the regional adoption response, and market size effects from L_t —holding the baseline calibration fixed. The figure confirms technology adoption differences as a key driver of productivity dynamics.

Table 5: Excess Productivity Growth in Data and Model

Calibration	Workers	Regional	Total	Frontier
<i>A. Decadal Average (1975–2025)</i>				
<i>Data</i>				
(Scaled) PWT	—	—	0.43	—
Fernald (2014)	—	—	0.32	—
BLS labor productivity	—	—	—	—
<i>Model</i>	0.06	0.00	0.07	0.02
<i>B. Dot-Com Decade (1995–2004)</i>				
<i>Data</i>				
(Scaled) PWT	—	—	1.44	—
Fernald (2014)	—	—	1.41	—
BLS labor productivity	—	—	1.46	—
<i>Model</i>	0.26	0.09	0.36	0.03

TFP-growth contributions and technology-frontier growth above the initial steady state, together with TFP data from Penn World Tables (PWT) and [Fernald \(2014\)](#) and labor productivity from the BLS. PWT growth is scaled by 0.5/0.3 to yield an initial growth rate around 0.5% per year; the PWT and Fernald rows subtract the 0.5% initial BGP growth rate. The BLS row in Panel B reports the change in annual nonfarm-business labor-productivity growth from 1975–1985 to 1995–2005; it is not reported in Panel A.

In their absence, human capital dynamics would be significantly less pronounced and the regional adoption rate essentially flat.

Alternative calibrations and targets. Appendix Table [E.1](#) recalibrates the model using computer rather than AI adoption, a flatter wage profile, lower and higher R&D semi-elasticities, alternative wage-response estimates, no wage-growth rescaling, and a lower long-run TFP growth target. The implied contribution to dot-com-era TFP growth ranges from 0.20 pp under the Mincer-based wage estimates to 0.68 pp without wage rescaling, compared with 0.36 pp in the baseline. Most alternatives lie between 0.24 and 0.39 pp.

Appendix Table [E.2](#) varies the externally chosen assumptions: the vintage premium, regional catch-up, obsolescence, the foreign share of innovation demand, and substitution across technologies. The results are relatively insensitive to θ and ξ but more sensitive to ρ and σ . Lower obsolescence and greater substitutability reduce the wage share associated with new technologies, requiring a larger regional adoption elasticity ϕ to match the wage evidence. Consequently, $\rho = 0$ and $\sigma = \infty$ raise the dot-com contribution to 0.58 and 0.59 pp, respectively, whereas $\sigma = 13$ lowers it to 0.22 pp.

Demographics and the AI Shock. Many advanced economies face demographic headwinds, as documented in Figure F.1. Could these headwinds dampen the productivity gains from artificial intelligence and other broad technology waves? Appendix E.3 applies a permanent reduction in innovation and regional-adoption costs under alternative demographic paths. The resulting TFP and frontier-growth responses are similar, indicating limited interaction in the calibrated model.²⁵ Thus, within this calibration, demographics do not materially alter the transmission of a multiplicative cost-reduction shock to innovation and adoption.

Dynamics Around Transitory and Permanent Shocks. Appendix E.4 contrasts permanent and transitory demographic shocks. A permanently lower entry rate gradually lowers the working-young share, new-technology adoption, and balanced-growth productivity growth, although greater average experience partly offsets the transition. A transitory increase initially lowers average experience but soon raises new-technology adoption; as the larger cohort ages, recovering experience generates a delayed but temporary acceleration in TFP growth.

6. CONCLUSION

This paper argues that the age composition of the workforce shapes the adoption of new technologies and, through it, wages and innovation. Young workers adopt new technologies at markedly higher rates than old workers—a pattern visible for the personal computer in the 1980s and for generative artificial intelligence today. An overlapping-generations model translates this lifecycle gradient into aggregate outcomes: household adoption raises the returns to regional technology adoption, while adoption-driven profits strengthen the incentive to innovate. Consistent with the model, U.S. commuting zones with younger workforces experienced faster growth in computer-science employment and wages, particularly among young workers, while commuting zones more exposed through trade to youth waves in foreign markets expanded their innovation activity. Calibrated to these responses, the model attributes about one quarter of the late-1990s U.S. productivity acceleration to the baby boom’s youth wave.

The findings partly recast the IT boom as a demographic episode: the personal computer arrived as the largest cohorts in U.S. history entered their prime adoption years. As the working-young share declines across advanced economies, the same

²⁵Absent obsolescence, the direct log elasticities implied by (17) and (18) are $\frac{\partial \ln g_t^R}{\partial \ln \kappa} = -\phi$ and $\frac{\partial \ln g_t^*}{\partial \ln \chi} = -\varphi$, which do not depend directly on demographics.

demand-side mechanism becomes a headwind to endogenous technology adoption and innovation. Yet aging need not substantially attenuate the gains from the particular AI-like shock studied here: in the calibrated model, demographic and technology shocks are approximately separable because of the constant-elasticity adoption and innovation structure. Thus, demographics can shape productivity dynamics without necessarily amplifying or dampening other forces acting on productivity at the same time.

REFERENCES

- Acemoglu, Daron, “Directed Technical Change,” *Review of Economic Studies*, 10 2002, 69, 781–809.
- , “When Does Labor Scarcity Encourage Innovation?,” *Journal of Political Economy*, 2010, 118 (6), 1037–1078.
- and David Autor, “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, Vol. 4, Elsevier, 2011, pp. 1043–1171.
- and Joshua Linn, “Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry,” *Quarterly Journal of Economics*, 2004, 119 (3), 1049–1090.
- and Pascual Restrepo, “Demographics and Automation,” *The Review of Economic Studies*, 2022, 89, 1–44.
- , David Autor, Keelan Beirne, and Andrew Scott, “Baby Busts and Growth Booms: Demographic Change and the Macroeconomy,” *NBER Working Paper 35401*, 2026.
- Adão, Rodrigo, Martin Beraja, and Nitya Pandalai-Nayar, “Fast and Slow Technological Transitions,” *Journal of Political Economy Macroeconomics*, 2024, 2, 183–227.
- Aghion, Philippe and Peter Howitt, “A Model of Growth through Creative Destruction,” *Econometrica*, 1992, 60 (2), 323–351.
- Aiyar, Shekhar, Christian Ebeke, and Xiaobo Shao, “The Impact of Workforce Aging on European Productivity,” *IMF Working Paper 16/238*, International Monetary Fund 2016.
- Akcigit, Ufuk and Sina T. Ates, “Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory,” *American Economic Journal: Macroeconomics*, 2021, 13, 257–298.
- Aksoy, Yunus, Henrique S Basso, Ron P Smith, and Tobias Grasl, “Demographic structure and macroeconomic trends,” *American Economic Journal: Macroeconomics*, 2019, 11, 193–222.
- Alekseeva, Liudmila, José Azar, Mireia Giné, Sampsa Samila, and Bledi Taska, “The Demand for AI Skills in the Labor Market,” *Labour Economics*, 2021, 71, 102002.
- Alvarez, Fernando, David Argente, Francesco Lippi, Esteban Méndez, and Diana Van Patten, “Strategic Complementarities in a Dynamic Model of Technology Adoption: P2P Digital Payments,” *NBER Working Paper 31280*, 2023. *American Economic Review*, conditionally accepted.

- Ang, James B and Jakob B Madsen, "Imitation versus innovation in an aging society: international evidence since 1870," *Journal of Population Economics*, 2015, 28, 299–327.
- Argente, David, Esteban Méndez, and Diana Van Patten, "How New Products Enter Markets: Social Networks, Consumer Imports, and Retailer Learning," *Working Paper*, 2026. Quarterly Journal of Economics, accepted.
- Aubert, Patrick, Eve Caroli, and Muriel Roger, "New Technologies, Organisation and Age: Firm-Level Evidence," *The Economic Journal*, 2006, 116, F73–F93.
- Autor, David H. and David Dorn, "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review*, 2013, 103 (5), 1553–1597.
- and —, "The growth of low-skill service jobs and the polarization of the US Labor Market," *American Economic Review*, 2013, 103, 1553–1597.
- Azoulay, Pierre, Benjamin F. Jones, J. Daniel Kim, and Javier Miranda, "Age and High-Growth Entrepreneurship," *American Economic Journal: Insights*, 2020, 2, 65–82.
- Barro, Robert J. and Xavier Sala i Martin, "Convergence," *Journal of Political Economy*, 1992, 100 (2), 223–251.
- Basso, Henrique S. and Juan F. Jimeno, "From secular stagnation to robocalypse? Implications of demographic and technological changes," *Journal of Monetary Economics*, 2021, 117, 833–847.
- Beaudry, Paul, Mark Doms, and Ethan Lewis, "Should the personal computer be considered a technological revolution? Evidence from U.S. metropolitan areas," *Journal of Political Economy*, 2010, 118, 988–1036.
- Ben-Porath, Yoram, "The Production of Human Capital and the Life Cycle of Earnings," *Journal of Political Economy*, 1967, 75 (4), 352–365.
- Berkes, Enrico, "Comprehensive Universe of U.S. Patents (CUSP): Data and Facts," *Working Paper*, 2018, pp. 1–27.
- , Ruben Gaetani, and Martí Mestieri, "Technological Waves, Knowledge Diffusion, and Local Growth," *Journal of Political Economy Macroeconomics*, 2025, 3 (1), 75–121.
- Bick, Alexander, Adam Blandin, and David J. Deming, "The Rapid Adoption of Generative AI," NBER Working Paper 32966, National Bureau of Economic Research, Cambridge, MA 2024. Forthcoming, *Management Science*.
- Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb, "Are Ideas Getting Harder to Find?," *American Economic Review*, 2020, 110, 1104–1144.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, "The Productivity J-Curve: How

- Intangibles Complement General Purpose Technologies,” *American Economic Journal: Macroeconomics*, 2021, 13 (1), 333–372.
- Cai, Jie and Andrey Stoyanov, “Population Aging and Comparative Advantage,” *Journal of International Economics*, 2016, 102, 1–21.
- Comin, Diego A., Danial Lashkari, and Martí Mestieri, “The Structural Transformation of Innovation,” NBER Working Paper 33855, National Bureau of Economic Research 2025.
- Comin, Diego and Bart Hobijn, “An Exploration of Technology Diffusion,” *American Economic Review*, 2010, 100 (5), 2031–2059.
- and Martí Mestieri, “If Technology Has Arrived Everywhere, Why Has Income Diverged?” *American Economic Journal: Macroeconomics*, 2018, 10 (3), 137–178.
- Crouzet, Nicolas, Pulak Ghosh, Apoorv Gupta, and Filippo Mezzanotti, “Demographics and Technology Diffusion: Evidence from Mobile Payments,” *Working Paper*, 2026.
- Deming, David J. and Kadeem Noray, “Earnings Dynamics, Changing Job Skills, and STEM Careers,” *Quarterly Journal of Economics*, 2020, 135 (4), 1965–2005.
- Derrien, Francois, Ambrus Kecskes, and Phuong-Anh Nguyen, “Labor Force Demographics and Corporate Innovation,” *The Review of Financial Studies*, 2023, 36, 2797–2838.
- DiNardo, John E. and Jörn-Steffen Pischke, “The Returns to Computer Use Revisited: Have Pencils Changed the Wage Structure Too?” *Quarterly Journal of Economics*, 1997, 112 (1), 291–303.
- DiNardo, John, Nicole M. Fortin, and Thomas Lemieux, “Labor Market Institutions and the Distribution of Wages, 1973–1992: A Semiparametric Approach,” *Econometrica*, 1996, 64 (5), 1001–1044.
- Eckert, Fabian, Andrés Gvartz, Jack Liang, and Michael Peters, “A Method to Construct Geographical Crosswalks with an Application to US Counties since 1790,” NBER Working Paper 26770, National Bureau of Economic Research 2020.
- Eggertsson, Gauti B., Neil R. Mehrotra, and Jacob A. Robbins, “A model of secular stagnation: Theory and quantitative evaluation,” *American Economic Journal: Macroeconomics*, 2019, 11, 1–48.
- Engbom, Niklas, “Misallocative Growth,” *Working Paper*, 2023.
- Eurostat, “Individuals — Use of Generative AI Tools [isoc_ai_iaiu],” Dataset, European Commission, Luxembourg 2026. Accessed [DATE].
- Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer, “The Next Generation of the Penn World Table,” *American Economic Review*, 2015, 105 (10), 3150–3182.

- Fernald, John G., "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," Working Paper 2012-19, Federal Reserve Bank of San Francisco 2014.
- Feyrer, James, "Demographics and Productivity," *The Review of Economics and Statistics*, 2007, 89, 100–109.
- Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, Daniel Backman, Etienne Breton, Grace Cooper, Julia A. Rivera Drew, Stephanie Richards, and David Van Riper, "Integrated Public Use Microdata Series, Current Population Survey: Version 13.0," [dataset]. Minneapolis, MN: IPUMS 2025.
- Friedberg, Leora, "The Impact of Technological Change on Older Workers: Evidence from Data on Computers," *Industrial and Labor Relations Review*, 2003, 56, 511–529.
- Gancia, Gino and Fabrizio Zilibotti, "Horizontal Innovation in the Theory of Growth and Development," *Handbook of Economic Growth*, 2005, 1, 111–170.
- Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J. Klenow, "How Destructive Is Innovation?," *Econometrica*, 2019, 87 (5), 1507–1541.
- Gu, Ke and Andrey Stoyanov, "Skills, Population Aging, and the Pattern of International Trade," *Review of International Economics*, 2019, 27, 499–519.
- Hassan, Tarek Alexander, Aakash Kalyani, and Pascual Restrepo, "The Skill Premium in Times of Rapid Technological Change," *NBER Working Paper 34939*, 2026.
- Hopenhayn, Hugo A., Julian Neira, and Rish Singhania, "From Population Growth To Firm Demographics: Implications for Concentration, Entrepreneurship and the Labor Share," *Econometrica*, 2022, 90, 1879–1914.
- Hoshi, Takeo, Naomi Kodama, and Huiyu Li, "The Effect of Aging on Entrepreneurship and Aggregate Productivity," *The Journal of the Economics of Ageing*, 2025, 32, 100605.
- Howitt, Peter, "Steady Endogenous Growth with Population and R&D Inputs Growing," *Journal of Political Economy*, 1999, 107 (4), 715–730.
- Humlum, Anders and Emilie Vestergaard, "The Unequal Adoption of ChatGPT Exacerbates Existing Inequalities among Workers," *Proceedings of the National Academy of Sciences*, 2025, 122 (1), e2414972121.
- Inokuma, Hiroshi and Juan M. Sánchez, "From Population Growth to TFP Growth," *American Economic Journal: Macroeconomics*, 2025. Forthcoming.
- Irmen, Andreas and Anastasia Litina, "Population Aging and Inventive Activity," *Macroeconomic Dynamics*, 2022, 26, 1127–1161.
- Jones, Benjamin F., "Age and great invention," *Review of Economics and Statistics*, 2010, 92, 1–14.

- and Bruce A. Weinberg, “Age dynamics in scientific creativity,” *PNAS*, 2011, 108, 18910–18914.
- Jones, Charles I., “R&D-Based Models of Economic Growth,” *Journal of Political Economy*, 1995, 103 (4), 759–784.
- , “The End of Economic Growth? Unintended Consequences of a Declining Population,” *American Economic Review*, 11 2022, 112 (11), 3489–3527.
- Jordà, Òscar, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M. Taylor, “The Rate of Return on Everything, 1870–2015,” *Quarterly Journal of Economics*, 2019, 134 (3), 1225–1298.
- Kalyani, Aakash, Nicholas Bloom, Marcela Carvalho, Tarek A. Hassan, Josh Lerner, and Ahmed Tahoun, “The Diffusion of New Technologies,” *Quarterly Journal of Economics*, 2025, 140 (2), 1299–1365.
- Karahan, Fatih, Benjamin Pugsley, and Ayşegül Şahin, “Demographic Origins of the Startup Deficit,” *American Economic Review*, 2024, 114 (7), 1986–2023.
- Klenow, Peter J. and Huiyu Li, “Innovative Growth Accounting,” *NBER Macroeconomics Annual*, 2021, 35, 245–295.
- Krueger, Alan B., “How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984–1989,” *Quarterly Journal of Economics*, 1993, 108 (1), 33–60.
- Lagakos, David, Benjamin Moll, Tommaso Porzio, Nancy Qian, and Todd Schoellman, “Life cycle wage growth across countries,” *Journal of Political Economy*, 2018, 126, 797–849.
- Lehr, Nils H, “Does Monopsony Matter for Innovation?,” *Working Paper*, 2025.
- Lehr, Nils Haakon, “Did R&D Misallocation Contribute to Slower Growth?,” 2025.
- Liang, James, Hui Wang, and Edward P. Lazear, “Demographics and Entrepreneurship,” *Journal of Political Economy*, 10 2018, 126, S140–S196.
- Maestas, Nicole, Kathleen J. Mullen, and David Powell, “The Effect of Population Aging on Economic Growth, the Labor Force, and Productivity,” *American Economic Journal: Macroeconomics*, 2023, 15, 306–332.
- Manuelli, Rodolfo E. and Ananth Seshadri, “Human Capital and the Wealth of Nations,” *American Economic Review*, 2014, 104 (9), 2736–2762.
- Nordhaus, William D., “Two Centuries of Productivity Growth in Computing,” *Journal of Economic History*, 2007, 67 (1), 128–159.
- OECD, “Trade in Value Added (TiVA): 2025 Edition,” OECD Inter-Country Input-Output Database, OECD, Paris 2025. Accessed [DATE] via <https://www.oecd.org/en/topics/sub-issues/trade-in-value-added.html>.

- Oliner, Stephen D., Daniel E. Sichel, and Kevin J. Stiroh, "Explaining a Productive Decade," *Brookings Papers on Economic Activity*, 2007, 2007 (1), 81–137.
- Peters, Michael and Conor Walsh, "Population Growth and Firm Dynamics," *Journal of Political Economy Macroeconomics*, 2026, 4, 256–296.
- Pew Research Center, "34% of U.S. Adults Have Used ChatGPT, About Double the Share in 2023," Short read, June 25 June 2025.
- Romer, Paul M, "Endogenous Technological Change," *Journal of Political Economy*, 1990, 98.
- Ruggles, Steven, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek, "IPUMS USA: Version 10.0 [dataset]," 2020.
- Shimer, R., "The Impact of Young Workers on the Aggregate Labor Market," *The Quarterly Journal of Economics*, 8 2001, 116, 969–1007.
- Syverson, Chad, "Challenges to Mismeasurement Explanations for the US Productivity Slowdown," *Journal of Economic Perspectives*, 5 2017, 31, 165–186.
- Terry, Stephen J., Thomas Chaney, Konrad B. Burchardi, Lisa Tarquinio, and Tarek A. Hassan, "Immigration, Innovation, and Growth," *American Economic Review*, 2026, 116 (3), 828–861.
- Teulings, Coen and Richard Baldwin, *Secular Stagnation: Facts, Causes and Cures*, CEPR Press, 2014.
- Tolbert, Charles M. and Molly Sizer, "US Commuting Zones and Labor Market Areas - A 1990 Update," 1996.
- vom Baur, Cäcilia, "No Teens, No Tech: How Shortages of New Entrants Hinder Firm Technology Investments," *Working Paper*, 2026.
- Weinberg, Bruce A., "Experience and Technology Adoption," IZA Discussion Paper 1051, IZA 2004.
- Xu, Zeshu, "Demographics and Technology Adoption," 2025. Working paper, BI Norwegian Business School.
- Young, Alwyn, "Growth without Scale Effects," *Journal of Political Economy*, 1998, 106 (1), 41–63.

APPENDIX

A. DERIVATIONS AND PROOFS

A.1. DERIVATIONS

Production. The final-good first-order conditions are $P_t(m) = (1 - \alpha) (l_t(m)/x_t(m))^\alpha$ and $W_t(m) = \alpha (l_t(m)/x_t(m))^{\alpha-1}$. The intermediate monopolist faces a constant demand elasticity $1/\alpha$ and sets $P_t(m) = \psi/(1 - \alpha)$; normalizing $\psi = 1 - \alpha$ gives $P_t(m) = 1$ and $x_t(m) = (1 - \alpha)^{1/\alpha} l_t(m)$. Substituting back, the wage per efficiency unit is common across technologies,

$$W_t(m) = \alpha (1 - \alpha)^{(1-\alpha)/\alpha} \equiv \omega, \quad (26)$$

and monopoly profits per technology are

$$\pi_t(m) = \alpha (1 - \alpha)^{1/\alpha} L_t a_t(m), \quad a_t(m) \equiv \frac{l_t(m)}{L_t}, \quad (27)$$

where $l_t(m)$ is the efficiency-weighted labor adopting technology m and $a_t(m)$ is the corresponding per-capita adoption rate. A fresh adopter supplies one efficiency unit and earns ω ; a household that retains a technology into old age supplies ω efficiency units and earns $\omega\omega$.

Household adoption. Substituting B_{it+1} out of the budget constraints, the adoption sets enter the young household's problem only through lifetime wealth; adoption is therefore wealth-maximizing and independent of the consumption-smoothing margin. A household adopts technology m whenever its marginal contribution to wealth weakly exceeds the cost $C c_{im}$, with $c_{im} \sim F(c) = c^{1/\eta}$ on $[0, 1]$ drawn once and permanent. The marginal new technology returns ω to an old household, and $\omega + \frac{1-\delta_t}{R_{t+1}} \omega\omega$ to a young household (the current wage plus the expected discounted old-age return on the retained technology). Because a fresh adopter earns ω on any technology regardless of its calendar vintage, each household applies a single reservation cost across M_t , and the adoption rates are

$$a_{ot} = \left(\frac{\omega}{C} \right)^{1/\eta}, \quad a_{yt} = \left(\frac{\omega}{C} \right)^{1/\eta} \left(1 + \frac{1 - \delta_t}{R_{t+1}} \omega \right)^{1/\eta}. \quad (28)$$

Neither rate depends on the population growth rate n_t .

Aggregation. Because $\bar{c}_{yt} > \bar{c}_{ot}$ and draws are permanent, every technology a surviving household would adopt while old was already adopted while young; old-age adoption is confined to technologies that first became available in the current period (Lemma 1). Hence each technology is adopted at its cohort's young-period rate. The per-capita efficiency-weighted adoption rates of newly and previously available technologies are

$$\bar{a}_t^N = \ell_{yt} a_{yt} + (1 - \ell_{yt}) a_{ot}, \quad \bar{a}_t^O = \ell_{yt} a_{yt} + (1 - \ell_{yt}) \omega a_{yt-1}, \quad (29)$$

where the old cohort supplies ω efficiency units on technologies it adopted while young. With $g_t \equiv M_t^N / M_{t-1}$, output per worker is $Y_t / L_t = (1 - \alpha)^{(1-\alpha)/\alpha} M_t \tilde{a}_t$, and total factor productivity is $TFP_t \equiv M_t \tilde{a}_t$ with

$$\tilde{a}_t = \ell_{yt} a_{yt} + (1 - \ell_{yt}) \left(\frac{\omega}{1 + g_t} a_{yt-1} + \frac{g_t}{1 + g_t} a_{ot} \right). \quad (30)$$

Regional adoption and frontier. The marginal implementation condition (10) with (27) gives the regional adoption rate (17), in which the scale L_t cancels,

$$a_t^R \equiv \frac{M_t^N}{M_{t-1}} = \left(\frac{\alpha(1 - \alpha)^{1/\alpha}}{\kappa} \bar{a}_t^N \right)^\phi \left(\frac{M_{t-1}}{M_{t-1}} \right)^\theta, \quad (31)$$

so the regional *rate* responds to demographics only through the household-adoption composite $\bar{a}_t^N$ and the distance to the frontier M_t / M_{t-1} . Scale instead operates on the *level* of the frontier through innovation: with $\Pi_t^I = \frac{1}{1+\phi} \bar{a}_{N,t}^R \bar{\pi}_{N,t}$ and $\bar{\pi}_{N,t} = \alpha(1 - \alpha)^{1/\alpha} L_t \bar{a}_t^N$, frontier growth (18) is $g_t^* = (\Pi_t^I / \chi(1 + 1/\phi))^\varphi M_{t-1}^{-\vartheta}$. Along the BGP Π_t^I grows at rate n while $M_{t-1}^{-\vartheta}$ falls at rate ϑg^* ; constancy of g^* requires $\varphi n = \vartheta g^*$, i.e. (19).

A.2. PROOFS

Proof of Lemma 1. (i)–(ii). From (28), $a_{yt}/a_{ot} = (1 + \frac{1-\delta_t}{R_{t+1}} \omega)^{1/\eta} > 1$, so $a_{yt} > a_{ot}$. The bracket is strictly decreasing in δ_t , so the gap is larger the lower is mortality.

(iii). Since $\bar{c}_{yt} > \bar{c}_{ot}$ and the draws c_{im} are permanent, any technology a surviving household would adopt while old (draw at most $\bar{c}_{ot}$) lies below $\bar{c}_{yt-1}$ and was already adopted while young; old-age adoption is thus confined to technologies that first became available in the current period. Each cohort therefore adopts a given technology at its young-period rate. Along the BGP $a_{yt} = a_y$ is constant, so the share of the old-technology stock M_{t-1} held by the young cohort at t (rate a_y) equals the share held by the old cohort, which adopted those technologies while young at $t - 1$ (also rate a_y). $\square$

Proof of Proposition 1. Population-growth channel. From (14), $\partial \ell_{yt} / \partial n_t = (1 - \delta_{t-1}) / (2 + n_t - \delta_{t-1})^2 > 0$, and by (28) the rates a_{yt}, a_{ot} are invariant to n_t . Hence $\partial \bar{a}_t^N / \partial \ell_{yt} = a_{yt} - a_{ot} > 0$ by Lemma 1, so a higher ℓ_{yt} driven by n_t raises $\bar{a}_t^N$ purely through composition. Since a_t^R is increasing in $\bar{a}_t^N$ with elasticity $\phi > 0$ in (31), the regional adoption rate rises.

Wages. Per-capita labor incomes are $W_{yt} = \omega a_{yt} M_{t-1} (1 + g_t)$ and $W_{ot} = \omega M_{t-1} (\omega a_{yt-1} + g_t a_{ot})$. Both rise with $g_t = a_t^R$, and

$$\frac{\partial \ln W_{yt}}{\partial \ln g_t} = \frac{g_t}{1 + g_t}, \quad \frac{\partial \ln W_{ot}}{\partial \ln g_t} = \frac{g_t}{\omega (a_{yt-1}/a_{ot}) + g_t}.$$

Since $\omega \geq 1$ and $a_{yt-1}/a_{ot} > 1$, the old-worker elasticity is strictly smaller, so wages rise disproportionately for the young and the experience premium W_{ot}/W_{yt} falls.

Mortality channel. A sustained increase in mortality raises the young share through δ_{t-1} in (14) but lowers the young adoption rate a_{yt} through δ_t in (28). The two enter $\bar{a}_t^N = \ell_{yt} a_{yt} + (1 - \ell_{yt}) a_{ot}$ with opposite sign: a larger young share raises $\bar{a}_t^N$ (composition), while a lower a_{yt} lowers it (within-cohort). The net effect on $\bar{a}_t^N$, and hence on a_t^R , is therefore ambiguous. $\square$

Proof of Proposition 2. Permanent channel. Along the BGP, (19) gives $\partial \ln(1 + g^*) / \partial \ln(1 + n) = \phi / \vartheta > 0$: a higher population growth rate raises frontier growth permanently through the market-size scale effect on $\Pi_t^I \propto L_t$.

Impact channels. Off the BGP, linearizing (18) and (12),

$$\hat{g}_t^* = \phi \hat{\Pi}_t^I - \vartheta \hat{M}_{t-1}, \quad \hat{\Pi}_t^I = \hat{L}_t + \hat{a}_t^N + \hat{a}_{N,t}^R.$$

A rise in the working-young share raises $\hat{a}_t^N$ through composition (Lemma 1) and, through Proposition 1, raises $\hat{a}_{N,t}^R$; both lift innovation-sector revenue and hence g_t^* on impact with elasticity ϕ . These two terms move with the youth share and revert as the demographic wave passes, so they amplify frontier growth only transiently; the scale term $\hat{L}_t$ is what survives into the balanced-growth elasticity above. $\square$

B. EMPIRICAL APPENDIX

This appendix collects supporting tables for the empirical analysis in Section 3.

B.1. DATA

Census and ACS samples. Employment, population, and wage measures are constructed from the 1980, 1990, and 2000 decadal Censuses and the 2010–2012 ACS from IPUMS (Ruggles et al., 2020), restricted to full-time full-year workers as described in Section 3.1. Geographies are mapped to 1990 commuting zones using the crosswalks of Autor and Dorn (2013a) and Eckert et al. (2020).

Lifecycle wage profiles. The relative-wage targets are estimated on white full-time full-year men, regressing log hourly wages on age- or potential-experience-group indicators with education, state, and year fixed effects. Potential experience is defined as age minus years of schooling minus six for workers with at least a high-school degree, and age minus 18 otherwise, following Lagakos et al. (2018); the experience profile aligns workers at labor-market entry and is the baseline calibration target, with the age profile used as a robustness variant.

Survival and effective labor supply. The survival rates $\tilde{\delta}_s$ are ten-year conditional survival probabilities between the age bands 25–34, 35–44, 45–54, and 55–64, computed from U.S. period life tables averaged over 1980–2010. The effective labor-supply transitions δ_s multiply survival by the growth of effective hours per capita across bands, where effective hours combine the extensive margin (employment) and the intensive margin (weeks times usual weekly hours), with non-workers contributing zero hours.

Computer use. The adoption targets combine the extensive margin of computer use at work with the intensity of use (the number of distinct tasks performed, normalized by its mean) from the 1989 CPS Computer Use Supplement, with non-users contributing zero. Rates are composition-adjusted by residualizing on industry, occupation, education, and year, and profiles are measured by the cohort’s age in 1989 so that they capture adoption of a newly arriving technology.

Productivity series. TFP growth is measured using the Penn World Table (Feenstra et al., 2015), the utilization-adjusted TFP series of Fernald (2014), and the BLS multifactor productivity series for the private business sector; labor productivity is output per hour in the nonfarm business sector (BLS). All series are accessed via FRED (mnemonics RTFPNAUSA632NRUG, MFPPBS, OPHNFB) and aggregated to decade averages.

C. QUANTITATIVE APPENDIX

This appendix collects supporting material for the quantitative analysis in Section 4.

C.1. A MULTI-GENERATIONAL EXTENSION

This section describes the multi-generational extension of the main model, which underlies the quantitative analysis in Section 4.

Technology replacement and effective units. As introduced in Section 4, technologies depreciate through improvement-driven replacement. Let E_t denote the flow of innovation events and M_t^c the count of frontier varieties. Each event either creates a new variety with counter $q = 1$ (probability $1 - s_I$) or improves an existing variety drawn uniformly at random (probability s_I), raising its counter by one and obsoleting the previous version. The replacement hazard faced by any version in use is therefore

$$\rho_t = s_I e_t, \quad e_t \equiv \frac{E_t}{M_{t-1}^c}, \quad (32)$$

where e_t is the event intensity. I measure all technology stocks in *effective units* (product lines): a variety with counter q accounts for q units. Denoting by $\bar{q}_t \equiv M_t / M_t^c$ the average counter, the variety count, the effective frontier stock, and the effective regional stock evolve as

$$M_t^c = (1 + (1 - s_I) e_t) M_{t-1}^c, \quad M_t = (1 - \rho_t) M_{t-1} + M_t^N, \quad M_t = (1 - \rho_t) M_{t-1} + M_{N,t}, \quad (33)$$

where the gross frontier flow $M_t^N = E_t (1 + s_I \bar{q}_{t-1})$ counts one line per new variety plus, for each improvement, the $q+1$ lines of the improved version—the q lines of the obsoleted version exit through the term $\rho_t M_{t-1}$ —and $M_{N,t}$ is gross regional adoption in effective units. Along the balanced growth path these laws of motion imply $\bar{q} = 1/(1 - s_I)$ and the event intensity $e = g^* + \rho$, so the hazard satisfies $\rho = s_I (g^* + \rho)$, i.e., $\rho = s_I g^* / (1 - s_I)$, and the gross flows are $M_t^N = (g^* + \rho) M_{t-1}$ and $M_{N,t} = (g^* + \rho) M_{t-1}$. Off the balanced growth path, ρ_t moves with innovation through (32).

In the baseline, re-adoption of an improved version carries a fresh idiosyncratic cost draw and restarts experience premia. Replacement appears in four places below: the laws of motion (33); the vintage distribution (46), which tilts toward recent versions; the household adoption values (40), which discount future income by the version's survival

probability; and the regional and frontier optimality conditions (48) and (49), which are stated in terms of the gross adoption flow and the event flow, respectively.

Demographic notation. Let $s \in \{1, \dots, S\}$ denote the “stage of life” of a generation. I denote by $\delta_{ss'}$ the probability of surviving from stage s to stage $s' \geq s$, with $\delta_{ss} = 1$; note that $\delta_{ss'}$ is a *survival* probability, whereas δ_t in Section 2 denotes mortality. To account for reduced labor supply in later stages, population sizes evolve with the effective rates $\tilde{\delta}_{ss'}$, which combine survival with labor-force participation and hours. A mass $(1 + n_t) L_{1t-1}$ of households enters at stage 1 each period, so that

$$L_{1t} = (1 + n_t) L_{1t-1} \quad \text{and} \quad L_{st} = \tilde{\delta}_{s-1s} L_{s-1t-1} \quad \text{for } 1 < s \leq S. \quad (34)$$

Experience premia ω_e accrue with periods of use: a technology adopted $e - 1$ periods ago yields ω_e efficiency units of labor, with $\omega_1 = 1$ and $\omega_{e+1} \geq \omega_e$. Vintage age is indexed by $g \geq 0$, the number of periods since the technology entered use; the vintage productivity multiplier is $\xi_0 = 1$ for technologies adopted in the current period and $\xi_g = \xi$ for $g \geq 1$, as in the baseline model.

Production. Final production uses the CES aggregator (25) over product lines. I first restate it at the level of individual lines—which removes the improvement counters from the system entirely—and then record the mapping back to the variety-level statement (25). Because stocks are measured in effective units, the regional stock M_t is also the measure of product lines available for production. Indexing lines by $u \in [0, M_t]$ and denoting by $g(u)$ the vintage age of line u , a representative firm hires line-specific labor $l_t(u)$ and intermediate goods $x_t(u)$ to produce final output

$$Y_t = M_t^{-\frac{1}{\sigma-1}} \left(\int_0^{M_t} y_t(u)^{\frac{\sigma-1}{\sigma}} du \right)^{\frac{\sigma}{\sigma-1}}, \quad y_t(u) = (\xi_{g(u)} l_t(u))^\alpha x_t(u)^{1-\alpha}, \quad (35)$$

taking wages $W_t(u)$ and intermediate goods prices $P_t(u)$ as given:

$$\max_{\{l_t(u), x_t(u)\}_{u \in [0, M_t]}} Y_t - \int_0^{M_t} (P_t(u) x_t(u) + W_t(u) l_t(u)) du. \quad (36)$$

The normalization by M_t removes mechanical love-of-variety effects, and $\sigma \rightarrow \infty$ recovers the linear aggregator of Section 2. I normalize the price of the final good to one and

denote by $\mathcal{P}_t(u) \equiv \partial Y_t / \partial y_t(u)$ the shadow price of line u 's output,

$$\mathcal{P}_t(u) = M_t^{-\frac{1}{\sigma-1}} \left(\int_0^{M_t} y_t(v)^{\frac{\sigma-1}{\sigma}} dv \right)^{\frac{1}{\sigma-1}} y_t(u)^{-\frac{1}{\sigma}}. \quad (37)$$

Mapping between the line-level and variety-level aggregators. Equation (35) is the main-text aggregator (25) written in different units. A variety j with counter q_{jt} consists of q_{jt} product lines: each improvement adds a line. Workers' technology-specific know-how is non-rival within a variety—an adopter supplies labor to every line of the variety—while intermediates are line-specific. In a symmetric equilibrium, all lines of a variety receive identical inputs, so grouping the line integral by variety gives

$$\int_0^{M_t} y_t(u)^{\frac{\sigma-1}{\sigma}} du = \int q_{jt} y_t(j)^{\frac{\sigma-1}{\sigma}} dj \quad \text{and} \quad M_t = \int q_{jt} dj = Q_t,$$

which is exactly (25): the counter enters as a linear line count, not as a quality shifter inside the CES kernel. The two formulations divide equilibrium objects cleanly. Per-line objects—the shadow price $\mathcal{P}_t(u)$, the wage per efficiency unit, and per-line intermediate quantities—are identical across the lines of a variety and independent of its counter. Variety-level objects—income per adopter, profits, and household and regional adoption costs—equal the per-line objects times q_{jt} . Because the counter multiplies both the benefits and the costs of adoption, adoption decisions are independent of q , and the system below can be stated entirely per line.²⁶

Intermediate goods. Each line is served by a monopolist who produces the intermediate good at constant unit cost ψ in terms of the final good and solves

$$\pi_t(u) = \max_{x_t(u)} P_t(x_t(u)) x_t(u) - \psi x_t(u), \quad (38)$$

where $P_t(x_t(u))$ is the inverse factor demand implied by (36). As in the baseline model, the monopolist takes the line's shadow price $\mathcal{P}_t(u)$ as given, so it faces the Cobb–Douglas factor demand $x_t(u) \propto P_t(u)^{-1/\alpha}$ and charges the constant markup $1/(1 - \alpha)$ over marginal cost, independently of σ . I maintain the normalization $\psi = 1 - \alpha$, so intermediate prices satisfy $P_t(u) = 1$.

²⁶Equivalently, a variety with counter q can be read as a single product whose q improvements make each physical unit q times more valuable at q times the marginal production cost; the two readings imply identical values for all equilibrium objects and differ only in the units in which prices and quantities are quoted.

Household problem. The household problem follows the baseline setup with S stages of life and stage-specific mortality risk. A household i at stage s solves

$$\begin{aligned} \max_{\{A_{it+k}^N, B_{it+k+1}\}} \quad & \sum_{k=0}^{S-s} \beta^k \delta_{ss+k} \ln C_{it+k} \quad \text{s.t.} \quad A_{it+k} = A_{it+k-1} + A_{it+k}^N \\ \text{and} \quad & B_{it+k+1} + C_{it+k} = \tilde{R}_{s+k} B_{it+k} + \int_{A_{it+k}} \omega_{e_i(m)} \mathcal{W}_{g(m)t+k} dm - C_{it+k}(A_{it+k}^N), \end{aligned} \quad (39)$$

where $e_i(m)$ counts the periods of use since household i adopted technology m (with $e_i(m) = 1$ in the adoption period), $g(m)$ is the vintage age of m , and adoption costs follow the baseline cost structure (4)–(5), scaled by the counter as described above. Households enter with zero assets and no technology know-how, $B_{it} = 0$ and $A_{it-1} = \emptyset$ for $s(i) = 1$. Mortality risk is insured through annuities: survivors earn the gross return $\tilde{R}_{s,t+1} = R_{t+1}/\delta_{ss+1}$, so that the expected return equals the risk-free rate.

In equilibrium, prices—and hence wages $\mathcal{W}_{gt}$ —are homogeneous within a vintage age g . Adoption follows a reservation-cost rule, and the adoption rate of stage- s households on vintage- g technologies satisfies

$$a_{sgt} = \left(\frac{\mathcal{W}_{gt}}{C} \right)^{\frac{1}{\eta}} \left(\sum_{k=0}^{S-s} \frac{\mathcal{W}_{g+k,t+k}}{\mathcal{W}_{gt}} \omega_{k+1} \frac{\delta_{ss+k}}{R^k} \prod_{i=1}^k (1 - \rho_{t+i}) \right)^{\frac{1}{\eta}}. \quad (40)$$

The $k = 0$ term is the current wage; future terms carry the experience premium ω_{k+1} and are discounted by the worker's survival δ_{ss+k} , the version's survival $\prod_i (1 - \rho_{t+i})$, and interest. Because the vintage multiplier is embedded in the wage ($\mathcal{W}_{gt} \propto \xi_g$, see (42)), equation (40) covers both relevant cases: when adopting the newest vintage ($g = 0$), the multiplier ξ applies only from the second period of use onward through the wage ratio, whereas when adopting a mature vintage ($g \geq 1$), it applies throughout.

Labor supply by vintage. The effective labor input per line of a vintage- g technology aggregates adopters across stages of life, weighted by experience premia and the vintage multiplier:

$$L_{gt} = \xi_g \sum_{s=1}^S L_{st} a_{sgt} \omega_{\min(s, g+1)}, \quad (41)$$

where the experience index $\min(s, g+1)$ reflects that a worker at stage s can have used a technology of vintage age g for at most $\min(s, g+1)$ periods. No additional version-survival factor appears in the stock: a line of vintage age g has, by definition, survived to age g , and replaced versions exit the vintage distribution through (46).

Prices, wages, and output. Substituting $P_t(u) = 1$ into the firm's optimality conditions, the wage per efficiency unit of technology-specific labor and the per-line profit on a vintage- g line are

$$\omega_{gt} = \xi_g \alpha (1 - \alpha)^{\frac{1-\alpha}{\alpha}} \mathcal{P}_{gt}^{\frac{1}{\alpha}} \quad (42)$$

and

$$\pi_{gt} = \alpha (1 - \alpha)^{\frac{1}{\alpha}} \mathcal{P}_{gt}^{\frac{1}{\alpha}} L_{gt}, \quad (43)$$

where $\mathcal{P}_{gt}$ denotes the common shadow price of vintage- g lines and L_{gt} is the effective labor input from (41). Inverting the factor-demand system, output is a function of the vintage labor supplies,

$$Y_t = (1 - \alpha)^{\frac{1-\alpha}{\alpha}} M_t L_t \left(\sum_{g=0}^{\infty} m_{gt} \ell_{gt}^{\frac{\sigma-1}{\sigma+\frac{1-\alpha}{\alpha}}} \right)^{\frac{\sigma+\frac{1-\alpha}{\alpha}}{\sigma-1}} \quad \text{with} \quad m_{gt} = \frac{M_{gt}}{M_t} \quad \text{and} \quad \ell_{gt} = \frac{L_{gt}}{L_t}, \quad (44)$$

and the vintage price consistent with these supplies is

$$\mathcal{P}_{gt} = \left(\frac{\ell_{gt}^{\frac{\sigma}{\sigma+\frac{1-\alpha}{\alpha}}}}{\sum_{g'=0}^{\infty} m_{g't} \ell_{g't}^{\frac{\sigma-1}{\sigma+\frac{1-\alpha}{\alpha}}}} \right)^{-\frac{1}{\sigma}}. \quad (45)$$

The vintage distribution follows from the law of motion of the regional stock in (33):

$$m_{0t} = \frac{M_{N,t}}{M_t}, \quad m_{gt} = (1 - \rho_t) \frac{M_{t-1}}{M_t} m_{g-1,t-1} \quad \text{for } g \geq 1, \quad (46)$$

so that replacement tilts the distribution toward recent versions. Finally, expenditure on intermediate inputs satisfies

$$\psi \sum_{g=0}^{\infty} M_{gt} x_{gt} = (1 - \alpha)^2 Y_t, \quad (47)$$

where x_{gt} is the per-line intermediate quantity.

Regional technology adoption. The set of locally available lines is determined by *regional* technology adoption, as in the baseline model. In the symmetric equilibrium studied here, all regions are identical, so the regional stock coincides with the aggregate stock entering final production; an individual region nonetheless takes the vintage prices $\mathcal{P}_{gt}$ as given, since regions are atomistic in the aggregate market for products. Adoption

operates at the level of varieties; because both profits and sunk costs scale with the counter q , the marginal implementation condition can be stated per line. With the marginal-cost structure of (8), gross regional adoption $M_{N,t}$ satisfies

$$\pi_{0t} \geq L_t \kappa \left(\frac{M_{N,t}}{M_{t-1}} \right)^{1/\phi} \left(\frac{M_{t-1}}{M_{t-1}} \right)^{\theta/\phi} \quad \text{s.t. } M_t \leq M_t, \quad (48)$$

with equality whenever adoption is interior. By (43), profits on the newest vintage are $\pi_{0t} = \alpha(1 - \alpha)^{1/\alpha} \mathcal{P}_{0t}^{1/\alpha} L_{0t}$: relative to the baseline model, they carry the price factor $\mathcal{P}_{0t}^{1/\alpha}$, so scarcity of labor on new vintages raises regional adoption incentives.

Innovation. The innovation sector produces innovation events. Its optimality condition takes the same form as in the baseline model (18) but applies to the event intensity e_t from (32):

$$e_t = \left(\frac{\Pi_t^I}{\chi(1 + 1/\phi)} \right)^\varphi (M_{t-1}^c)^{-\vartheta}, \quad (49)$$

where the fishing-out term operates on the variety count and Π_t^I is the expected revenue per event. An event creates $1 + s_I \bar{q}_{t-1}$ product lines in expectation, each of which earns the innovation sector the per-line adoption surplus as in (12):

$$\Pi_t^I = (1 + s_I \bar{q}_{t-1}) \frac{1}{1 + \phi} \mathbb{E}_t[a_{N,t}^R \pi_{0t}], \quad a_{N,t}^R = \frac{M_{N,t}}{M_t - M_{t-1}}, \quad (50)$$

where $a_{N,t}^R$ is the regional adoption rate of not-yet-adopted lines. The laws of motion (33) then imply the frontier growth rate and the evolution of the average counter,

$$g_t^* = \frac{e_t}{\bar{q}_{t-1}} \quad \text{and} \quad \bar{q}_t = \bar{q}_{t-1} \frac{1 + g_t^*}{1 + (1 - s_I) e_t}.$$

In the implementation, the variety count is normalized to equal the effective stock along the balanced growth path and the expected-lines factor is expressed relative to its balanced-growth value $1 + s_I \bar{q} = \bar{q}$, absorbing both constants into χ ; with $s_I = 0$, the system collapses to the baseline condition (18).

Equilibrium system. Given demographic paths $\{n_t\}$ and parameters, the equilibrium is a fixed point in the vintage-price profile $\{\mathcal{P}_{gt}\}$ and the hazard path $\{\rho_t\}$: (i) given prices and hazards, wages follow from (42) and household adoption rates from (40); (ii) adoption rates aggregate to vintage labor supplies via (41); (iii) the regional marginal implementation condition (48) determines gross adoption $M_{N,t}$, and the innovation block (49)–(50) determines the event intensity e_t , hence the hazard $\rho_t = s_I e_t$ and the

technology stocks via (33); (iv) the vintage distribution (46) and labor supplies return updated prices through (45), closing the loop. Output and intermediate expenditure follow from (44) and (47).

C.2. MOMENTS

In this section, I further discuss the construction of moments used for the calibration.

Experience Wage profiles. Following Lagakos et al. (2018), I construct wage profiles according to potential experience using the census data. I first restrict the data to FTFY workers and then calculate potential experiences as age minus years of schooling, where I infer the latter based on the worker’s highest degree. I then construct the wage profile based on the decades of potential experience, i.e., 0–9, 10–19, 20–29, and 30–39.

Adoption profiles. To capture the intensive and extensive margin of technology adoption, I calculate the adoption rate as the product of extensive adoption times the intensive adoption rate, where I normalize the latter to an average of 1. Using these categories, I then construct average adoption rates by age group.

C.3. IDENTIFICATION OF THE RESPONSE ELASTICITIES

This section derives the mapping from the two response moments—the wage-growth responses to the local working-young share and the R&D response to trade-weighted demographics—to the elasticities ϕ and φ in the two-generation model of Section 2, augmented with technology replacement at rate ρ as in the quantitative implementation. Setting $\rho = 0$ recovers the model of Section 2. The derivations are first-order impact responses around the balanced growth path. The quantitative calibration matches the wage moments to exact simulated counterparts in the multigenerational model and uses a simulated foreign-market demand response in mapping the R&D moment; the formulas below clarify which steady-state objects drive those mappings.

Wage responses identify ϕ . With replacement, the adopted technology stock evolves as $M_t = (1 - \rho)M_{t-1} + M_t^N$, so that along the balanced growth path the flow of newly adopted technologies is $M_t^N = (g^* + \rho) M_{t-1}$ and recent vintages account for a share $(g^* + \rho)/(1 + g^*)$. Consider an unexpected one-time increase in the size of the entering cohort at t , raising the working-young share by $d\ell_{yt}$. By Lemma 1, the individual adoption rates a_{yt}, a_{ot} are unchanged, so the shock moves wages only through the supply of newly

adopted technologies. First, the adoption composite responds through composition alone:

$$d \ln \bar{a}_t^N = \varepsilon_{\bar{a}} d\ell_{yt}, \quad \varepsilon_{\bar{a}} \equiv \frac{a_{yt} - a_{ot}}{\bar{a}_t^N}. \quad (51)$$

With the technology stock M_{t-1} and the distance term predetermined, the regional adoption rule (17) implies $d \ln M_t^N = \phi d \ln \bar{a}_t^N$. The frontier path is taken as given for a single region. Finally, newly adopted technologies enter group incomes in proportion to each group's income share on the newest vintage. In the two-generation model, young workers earn $W_{yt} = \omega a_{yt} M_t$, and old workers earn the experience-augmented wage on the surviving part of the stock they adopted when young plus the fresh wage on newly adopted technologies, $W_{ot} = \omega (\omega a_{yt-1} (1 - \rho) M_{t-1} + a_{ot} M_t^N)$. Therefore

$$\frac{d \ln W_{jt}}{d \ell_{yt}} = s_j^N \phi \varepsilon_{\bar{a}}, \quad s_y^N = \frac{g^* + \rho}{1 + g^*}, \quad s_o^N = \frac{a_{ot} (g^* + \rho)}{\omega a_{yt} (1 - \rho) + a_{ot} (g^* + \rho)}, \quad (52)$$

where s_j^N is group j 's steady-state income share on the newest vintage. Because lagged wages are predetermined, (52) is also the impact response of wage *growth*, the object estimated in the decadal first-difference regressions. Thus, conditional on $\{a_y, a_o, \omega, g^*\}$, each wage moment therefore identifies ϕ . Their ratio, i.e., relative response if given by

$$\frac{d \ln W_{yt}}{d \ln W_{ot}} = \frac{s_y^N}{s_o^N} = \frac{\omega a_{yt} (1 - \rho) + a_{ot} (g^* + \rho)}{a_{ot} (1 + g^*)},$$

and exceeds one whenever $\omega a_{yt} > a_{ot}$, which holds since $\omega \geq 1$ and $a_{yt} > a_{ot}$. Young wages respond more because young incomes load more heavily on new technologies. Replacement raises both newest-vintage income shares, so that a given empirical wage response is consistent with a smaller regional elasticity ϕ .

The innovation response identifies ϕ . Consider a shift $d\ell_y$ in the working-young share of markets accounting for a share μ of the demand for newly invented technologies. Innovation revenue per new technology is $\Pi_t^I = \frac{1}{1+\phi} a_{N,t}^R \bar{\pi}_{N,t}$. The shock raises profits per adopted technology directly ($\bar{\pi}_{N,t} \propto \bar{a}_t^N$, elasticity one) and the adoption hazard through free entry (elasticity ϕ), so, holding frontier growth fixed,

$$d \ln \Pi_t^I|_{g^*} = \mu (1 + \phi) \varepsilon_{\bar{a}} d\ell_y.$$

With replacement, the frontier evolves as $M_t = (1 - \rho) M_{t-1} + M_t^N$, so the gross flow of new frontier technologies along the balanced growth path is $M_t^N = (g^* + \rho) M_{t-1}$,

and with M_{t-1} predetermined, $d \ln M_t^N = d \ln(g^* + \rho)$. The innovation optimality condition (18) then implies $d \ln M_t^N = \varphi d \ln \Pi_t^I$. Faster frontier growth, however, enlarges the pool of technologies awaiting adoption and dilutes the adoption hazard, $a_{N,t}^R = M_t^N / (M_t - M_{t-1})$, which feeds back into revenue with semi-elasticity $-1/(1 + g^* - \hat{m})$. Solving the fixed point,

$$d \ln M_t^N = \frac{\varphi \mu (1 + \phi) \varepsilon_{\bar{a}}}{1 + \varphi r} d\ell_y, \quad r = \frac{g^* + \rho}{1 + g^* - \hat{m}}, \quad (53)$$

setting $\rho = 0$ recovers the feedback term of the model without replacement. Given ϕ , $\varepsilon_{\bar{a}}$, $\hat{m}$, g^* , ρ , and μ , the left-hand side is strictly increasing in φ , so the moment identifies φ uniquely.

SUPPLEMENTAL APPENDIX

D. ADDITIONAL EMPIRICAL RESULTS

D.1. SECTORAL HETEROGENEITY IN COMPUTER ADOPTION

Table D.1 compares workplace computer use in goods and service sectors. In the raw specifications, computer use among old workers is 10.8 percentage points higher in services in 1989 and 15.6 points higher in 2003. The young–old gradient is also 3.8 and 5.1 points more positive in services, respectively. Controlling for gender, education, and detailed occupation absorbs the service-sector level difference and reduces the differential age gradient to 1.2 points in 1989, which is statistically insignificant, and 2.9 points in 2003, which remains statistically significant. Thus, the concentration of computer adoption in service employment partly reflects its occupational composition, while the age profile remains relatively more favorable to young workers in services. These comparisons are descriptive and do not identify a causal sectoral effect.

Table D.1: Workplace Computer Adoption by Sector and Age

	(1)	(2)	(3)	(4)
	Computer use at work			
Services	0.108** (0.050)	-0.013 (0.027)	0.156*** (0.043)	-0.014 (0.014)
Young (25–44)	0.030* (0.017)	0.060*** (0.011)	-0.046* (0.026)	-0.013 (0.008)
Services × Young (25–44)	0.038* (0.021)	0.012 (0.013)	0.051* (0.027)	0.029*** (0.009)
Year	1989	1989	2003	2003
Controls	No	Yes	No	Yes
Observations	41,164	41,164	43,327	43,327

Note: The dependent variable equals one if a worker uses a computer at work. Young workers are ages 25–44; old workers are ages 45–64. Services comprise utilities and transportation, wholesale and retail trade, finance, insurance, and real estate, low-skill and professional services, and public administration; goods comprise agriculture, mining, construction, and manufacturing. The services coefficient is the service–goods difference among old workers, the young coefficient is the young–old difference in goods, and their interaction is the additional young–old difference in services. Adjusted specifications control for gender, education, and detailed occupation. Regressions use CPS Computer Use Supplement weights. Standard errors are clustered by detailed industry.

Standard Errors in Parentheses. Significance levels: * 10% , ** 5%, *** 1%.

D.2. SPECIFICATION LADDER FOR REGIONAL ADOPTION

Table D.2 reports the computer-science adoption response across increasingly saturated specifications: year fixed effects only, year fixed effects with 1980 pre-characteristics (the college, non-white, female, working-young, and metropolitan shares and the working-age population), and the baseline with year and CZ fixed effects.

Table D.2: Workforce Aging and Computer Science Employment

	(1)	(2)	(3)	(4)
A. First Stage	Δ Working Young Share			
$\Delta \widehat{WYS}$	0.135*** (0.031)	0.135*** (0.031)	0.161*** (0.034)	0.213*** (0.035)
B. Second Stage	Δ Computer science employment			
ΔWYS	0.010 (0.022)	0.137** (0.068)	0.166*** (0.042)	0.199*** (0.040)
1980 controls			✓	
Year FE		✓	✓	✓
CZ FE				✓
First-stage F		19.4	22.9	37.2
Observations	2,166	2,166	2,166	2,166

Note: Panel A reports first-stage results and Panel B reports second-stage results. All regressions are estimated by 2SLS, instrumenting ΔWYS with $\Delta \widehat{WYS}$, the historical-fertility instrument described in the text. Reported first-stage statistics are Kleibergen–Paap robust F -statistics. The 1980 controls comprise the bachelor, non-white, female, working-young, and metropolitan shares as well as the working-age population. CZ observations are weighted by 1980 working-age population. Standard errors are clustered at the state level. See the text for variable descriptions.

Standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Table D.3 adds Census-region-by-year and state-by-year fixed effects to absorb geographic shocks common to local labor markets in each decade. The computer-science employment and wage responses remain positive and statistically significant in both specifications.

Table D.3: Workforce Aging: Geographic-Time Fixed Effects

	(1)	(2)	(3)	(4)
A. Year FE	Δ Outcome			
Δ WYS	0.199*** (0.040)	0.191*** (0.050)	0.231*** (0.059)	0.117** (0.047)
B. Census region $\times$ year FE	Δ Outcome			
Δ WYS	0.289*** (0.086)	0.289*** (0.073)	0.361*** (0.085)	0.148** (0.055)
C. State $\times$ year FE	Δ Outcome			
Δ WYS	0.367*** (0.084)	0.322*** (0.083)	0.401*** (0.093)	0.176** (0.072)
Outcome	CS em- ployment	Wages, all	Wages, young	Wages, old

Note: Each panel reports 2SLS estimates of the response to Δ WYS, instrumented by the historical-fertility instrument Δ WYS. Column (1) reports computer-science employment as a share of CZ employment; columns (2)–(4) report composition-adjusted log wages for all, young (25–44), and old (45–64) workers. All specifications include CZ fixed effects. Panel A includes year fixed effects, Panel B Census-region-by-year fixed effects, and Panel C state-by-year fixed effects. The Kleibergen–Paap first-stage F-statistics are 37.2, 14.6, and 33.1, and the numbers of observations are 2,166, 2,166, and 2,157 in Panels A, B, and C, respectively. CZ observations are weighted by 1980 working-age population. Standard errors are clustered at the state level.

Standard errors in parentheses. Significance levels: * 10% , ** 5%, *** 1%.

D.3. ROBUSTNESS FOR WAGE REGRESSIONS

Table D.5 reports the bad-control exercise: the outcome responses to the predicted working-young share are stable to progressively adding contemporaneous demographic changes (working-age population, college share, non-white and female shares) as controls. Table D.4 probes the wage responses along the weighting and outlier dimensions.

Table D.7 reports the reduced-form counterparts of the instrumental-variable exercises: the computer-science employment and wage outcomes regressed directly on the change in the predicted working-young share constructed from historical local births.

Table D.8 probes the innovation-via-trade exercise of Table 3. Each column is an innovation outcome, including R&D employment excluding social scientists in the final column. Panel A repeats the baseline specification. Panel B adds the contemporaneous

Table D.4: Wages and Workforce Aging: Weighting and Outlier Robustness

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. Young workers (25–44)	$\Delta \ln \text{Wage (Young)}$						
ΔWYS	0.231*** (0.059)	0.225*** (0.062)	0.230*** (0.059)	0.218*** (0.057)	0.178*** (0.048)	0.169*** (0.047)	0.001 (0.082)
B. Old workers (45–64)	$\Delta \ln \text{Wage (Old)}$						
ΔWYS	0.117** (0.047)	0.052 (0.047)	0.117** (0.047)	0.113** (0.043)	0.070** (0.033)	0.055* (0.033)	-0.100 (0.070)
Outcome winsor.	–	–	1/99	5/95	–	–	–
Weights	Pop.	Pop.	Pop.	Pop.	Cap p95	Cap p90	None
Reweighting	4-way	Age×Edu	4-way	4-way	4-way	4-way	4-way
First-stage F	37.2	37.2	37.2	37.2	56.1	47.9	21.5
Observations	2,166	2,166	2,166	2,166	2,166	2,166	2,166

Note: Each column reports a 2SLS estimate of the wage response to ΔWYS , instrumented by the historical-fertility instrument ΔWYS , with CZ and year fixed effects and standard errors clustered at the state level. *Outcome winsor.* gives the percentiles at which the wage outcome is winsorized (– = none); *Weights* gives the observation weighting (Pop. = 1980 working-age population; Cap p_x = population weights winsorized at the x th percentile; None = unweighted); *Reweighting* gives the composition adjustment (4-way = age×sex×race×education; Age×Edu = age×college only). Capping population weights reduces the coefficients; in the unweighted specification, the estimates are 0.001 for young workers and -0.100 for old workers.

Standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

change in the *local* working-young share as a control: the responses to foreign workforce aging are essentially unchanged, and the local share enters insignificantly, indicating that the trade-weighted measure does not proxy for local demographic change.

E. ADDITIONAL QUANTITATIVE RESULTS

E.1. ALTERNATIVE CALIBRATIONS

Tables E.1 and E.2 report the alternative calibrations behind Tables E.3 and E.4.

Table D.5: Workforce Aging: Bad-Control

	(1)	(2)	(3)	(4)
A. Computer science employment		Δ CS Employment		
ΔWYS (pred.)	0.199*** (0.034)	0.196*** (0.031)	0.170*** (0.032)	0.141*** (0.027)
B. Young wages (25–44)		Δ ln Wage (Young)		
ΔWYS (pred.)	0.231*** (0.071)	0.242*** (0.058)	0.205*** (0.054)	0.158*** (0.044)
C. Old wages (45–64)		Δ ln Wage (Old)		
ΔWYS (pred.)	0.117** (0.055)	0.128** (0.049)	0.099** (0.047)	0.070* (0.040)
Δ WA pop.		✓	✓	✓
Δ College			✓	✓
Δ Non-white/Female				✓
Year FE	✓	✓	✓	✓
CZ FE	✓	✓	✓	✓
Observations	2,166	2,166	2,166	2,166

Note: Each column reports an OLS regression of the panel outcome on the predicted working-young share, obtained from the first-stage regression of ΔWYS on the historical-fertility instrument. Columns progressively add decadal changes in the working-age population, college share, and non-white and female shares as controls. Wages are composition-adjusted (DFL-reweighted to 1980). All regressions include year and CZ fixed effects. CZ observations are weighted by 1980 working-age population. Standard errors are clustered at the state level.

Standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Table D.6: Wages and Workforce Aging: Alternative Composition Adjustments

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	$\Delta \ln \text{Wages}$						
ΔWYS	0.191*** (0.050)	0.068*** (0.021)	0.142*** (0.049)	0.231*** (0.059)	0.109*** (0.027)	0.117** (0.047)	0.021 (0.026)
Workers	All	All	All	Young	Young	Old	Old
Adjustment	DFL	Mincer	Within-age	DFL	Mincer	DFL	Mincer
First-stage F	37.2	37.2	37.2	37.2	37.2	37.2	37.2
Observations	2,166	2,166	2,166	2,166	2,166	2,166	2,166

Note: Each column reports the 2SLS estimate of the wage response to ΔWYS , instrumented by the historical-fertility instrument ΔWYS , with CZ and year fixed effects and standard errors clustered at the state level. CZ observations weighted by 1980 working-age population. The *Adjustment* row gives the composition adjustment of the wage outcome: DFL reweighting to the national 1980 distribution over age, gender, race, and education (baseline); a Mincer residualization of log wages; and the within-age component of the CZ-level wage change, which is available for the aggregate only. See text for variable description.

Standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Table D.7: Workforce Aging: Reduced-Form Estimates

	(1)	(2)	(3)	(4)
	$\Delta \text{Outcome}$			
ΔWYS	0.042*** (0.007)	0.041*** (0.013)	0.049*** (0.015)	0.025** (0.012)
Outcome	CS emp.	Wages, all	Wages, young	Wages, old
Observations	2,166	2,166	2,166	2,166

Note: Reduced-form OLS estimates of the outcomes in Tables 1 and 2 on the predicted working-young share change ΔWYS constructed from historical local births. Column (1) reports computer-science employment as a share of CZ employment; columns (2)–(4) report composition-adjusted (DFL-reweighted) log wages for all, young (25–44), and old (45–64) workers. All regressions control for CZ and year fixed effects. CZ observations weighted by 1980 working-age population. Standard errors clustered at state level.

Standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Table D.8: Workforce Aging via Trade and Innovation: Local Demographics Horserace

	(1)	(2)	(3)	(4)
A. Baseline	Δ Innovation outcome			
ΔWYS^{Trade}	0.100*** (0.034)	0.118*** (0.015)	0.027*** (0.006)	0.080** (0.031)
B. + Local ΔWYS	Δ Innovation outcome			
ΔWYS^{Trade}	0.091** (0.036)	0.110*** (0.015)	0.026*** (0.006)	0.073** (0.032)
ΔWYS	0.018 (0.022)	0.017 (0.014)	0.002 (0.006)	0.014 (0.021)
Outcome	R&D emp.	Cite-wt. patents	Raw patents	R&D emp., narrow
Observations	2,166	2,166	2,166	2,166

Note: Each column reports the response of an innovation outcome to ΔWYS^{Trade} , the change in the working-young share of U.S. export markets. Columns (1) and (4) report R&D employment as a share of the working-age population, where column (4) excludes social scientists from the R&D occupation set; columns (2)–(3) report patenting per 1,000 working-age population. Panel A repeats the baseline specification of Table 3. Panel B adds the contemporaneous change in the local working-young share as a control. All regressions control for CZ and year fixed effects. CZ observations weighted by 1980 working-age population. Standard errors clustered at state level.

Standard errors in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Table E.1: Alternative Calibration Targets

A. Parameters	Baseline	Bad Control	OLS	Mincer	Low R&D	High R&D	Wages	0.3% TFP Gr.	Comp. Adop.	Wage Scaling
Substitution elast. σ	23.9	23.9	23.9	23.9	23.9	23.9	23.9	23.9	23.9	23.9
Improvement share s_I	0.73	0.73	0.73	0.73	0.73	0.73	0.73	0.73	0.73	0.73
Experience premium ω	{1.40, 1.83, 2.16}	{1.40, 1.83, 2.16}	{1.40, 1.83, 2.16}	{1.40, 1.83, 2.16}	{1.40, 1.83, 2.16}	{1.40, 1.83, 2.16}	{1.38, 1.67, 1.86}	{1.40, 1.83, 2.16}	{1.38, 1.76, 2.00}	{1.40, 1.83, 2.16}
Vintage premium ξ	1.15	1.15	1.15	1.15	1.15	1.15	1.15	1.15	1.15	1.15
Adop. cost elast. η	1.10	1.10	1.10	1.10	1.10	1.10	1.05	1.10	2.12	1.10
Adop. cost shifter ω/C	0.18	0.18	0.18	0.18	0.18	0.18	0.19	0.18	0.08	0.18
Reg. revenue elast. ϕ	6.35	4.30	3.59	2.68	6.35	6.35	6.38	6.35	10.9	17.5
Reg. adop. cost κ	0.05	0.05	0.06	0.08	0.05	0.05	0.04	0.05	0.04	0.04
Contemp. R&D returns φ	0.13	0.17	0.19	0.24	0.10	0.27	0.12	0.13	0.16	0.07
Fishing-out ϑ	0.08	0.11	0.13	0.16	0.06	0.18	0.08	0.08	0.11	0.04
B. Moments										
Relative Wage										
Age-Group 2	1.29	1.29	1.29	1.29	1.29	1.29	1.28	1.29	1.29	1.29
Age-Group 3	1.50	1.50	1.50	1.50	1.50	1.50	1.40	1.50	1.50	1.50
Age-Group 4	1.56	1.56	1.56	1.56	1.56	1.56	1.38	1.56	1.56	1.56
Wage Elasticity										
Young	0.24	0.16	0.13	0.10	0.24	0.24	0.24	0.24	0.23	0.24
Old	0.10	0.07	0.06	0.04	0.10	0.10	0.11	0.10	0.13	0.10
Novel Technology Premium	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.08	0.10
R&D semi-elasticity	0.04	0.04	0.04	0.04	0.03	0.08	0.04	0.04	0.04	0.04


Notes: Columns re-calibrate the full model under alternative targets. The baseline targets the unadjusted IV estimates of the wage regression; Bad Control uses the more conservative bad-control estimates, OLS uses the OLS estimates, and Mincer uses the Mincer-adjusted estimates. Low R&D targets a conservative R&D semi-elasticity of 0.03 instead of the baseline 0.04; High R&D doubles the baseline target to 0.08. Wages targets an alternative, flatter wage profile. Comp. Adop. targets the computer adoption rates in 1989. Wage Scaling removes the rescaling of the regression coefficients to match TFP growth.

Table E.2: Alternative External Parameter Choices

A. Parameters	Baseline	$\xi = 1$	$\theta = .5$	$\theta = 2$	$\rho = 0$	20% Foreign	$\sigma = 13$	$\sigma = \infty$
Substitution elast. σ	23.9	23.9	23.9	23.9	23.9	23.9	13.0	∞
Improvement share s_I	0.73	0.73	0.73	0.73	0.00	0.73	0.73	0.73
Experience premium ω	{1.40, 1.83, 2.16}	{1.40, 1.84, 2.18}	{1.40, 1.83, 2.16}	{1.40, 1.83, 2.16}	{1.32, 1.57, 1.68}	{1.40, 1.83, 2.16}	{1.40, 1.84, 2.18}	{1.39, 1.81, 2.13}
Vintage premium ξ	1.15	1.00	1.15	1.15	1.15	1.15	1.15	1.15
Adop. cost elast. η	1.10	1.03	1.10	1.10	1.20	1.10	1.05	1.18
Adop. cost shifter ω/C	0.18	0.21	0.18	0.18	0.16	0.18	0.18	0.18
Reg. revenue elast. ϕ	6.35	5.60	6.35	6.35	23.1	6.35	5.76	7.49
Reg. adop. cost κ	0.05	0.05	0.04	0.05	0.04	0.05	0.05	0.04
Contemp. R&D returns φ	0.13	0.14	0.13	0.13	0.04	0.06	0.24	0.05
Fishing-out ϑ	0.08	0.09	0.09	0.08	0.02	0.04	0.16	0.04
B. Moments								
Relative Wage								
Age-Group 2	1.29	1.29	1.29	1.29	1.29	1.29	1.29	1.29
Age-Group 3	1.50	1.50	1.50	1.50	1.50	1.50	1.50	1.50
Age-Group 4	1.56	1.56	1.56	1.56	1.56	1.56	1.56	1.56
Wage Elasticity								
Young	0.24	0.24	0.24	0.24	0.24	0.24	0.24	0.24
Old	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10
Novel Technology Premium	0.10	0.07	0.10	0.10	0.09	0.10	0.16	0.00
R&D semi-elasticity	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04

Notes: Columns re-calibrate the full model under alternative external parameter choices. θ is the elasticity of regional adoption to distance from the frontier (baseline 1), ξ the vintage premium, ρ the obsolescence hazard, and σ the elasticity of substitution across products (calibrated to a novel-technology wage premium of about 10% in the baseline; $\sigma = 13$ implies about 20%, and $\sigma = \infty$ is the linear model). 20% Foreign assumes a 20% foreign share of demand for U.S. innovation instead of the baseline 10%.

Figure E.1: The Demographic Boom and Imperfect Substitution across Products



Notes: Dynamic responses to the demographic transition for alternative values of σ .

Table E.3: Excess Productivity Growth — Alternative Calibration Targets

Calibration	Workers	Regional	Total	Frontier
<i>A. Decadal Average (1975–2025)</i>				
<i>Data</i>				
(Scaled) PWT	—	—	0.43	—
Fernald (2014)	—	—	0.32	—
BLS labor productivity	—	—	—	—
<i>Model</i>				
Baseline	0.06	0.00	0.07	0.02
Bad Control	0.05	-0.02	0.03	0.03
OLS	0.04	-0.03	0.01	0.03
Mincer	0.03	-0.05	-0.01	0.03
Low R&D	0.07	0.02	0.08	0.02
High R&D	0.05	-0.05	0.00	0.04
Wage Profile	0.06	0.01	0.06	0.02
0.3% TFP Gr.	0.06	0.00	0.07	0.02
Comp. Adop.	0.06	0.00	0.07	0.03
Wage Rescale	0.11	0.11	0.22	0.02
<i>B. Dot-Com Decade (1995–2004)</i>				
<i>Data</i>				
(Scaled) PWT	—	—	1.44	—
Fernald (2014)	—	—	1.41	—
BLS labor productivity	—	—	1.46	—
<i>Model</i>				
Baseline	0.26	0.09	0.36	0.03
Bad Control	0.24	0.03	0.28	0.04
OLS	0.23	0.01	0.24	0.04
Mincer	0.22	-0.02	0.20	0.04
Low R&D	0.27	0.11	0.38	0.02
High R&D	0.24	0.03	0.26	0.06
Wage Profile	0.24	0.11	0.35	0.03
0.3% TFP Gr.	0.26	0.09	0.36	0.03
Comp. Adop.	0.22	0.17	0.39	0.04
Wage Rescale	0.34	0.33	0.68	0.02

TFP-growth contributions and technology-frontier growth above the initial steady state, together with TFP data from Penn World Tables (PWT) and [Fernald \(2014\)](#) and labor productivity from the BLS. PWT growth is scaled by 0.5/0.3 to yield an initial growth rate around 0.5% per year; the PWT and Fernald rows subtract the 0.5% initial BGP growth rate. The BLS row in Panel B reports the change in annual nonfarm-business labor-productivity growth from 1975–1985 to 1995–2005; it is not reported in Panel A. Each model row re-calibrates the full model to the alternative target indicated in Appendix Table E.1.


Table E.4: Excess Productivity Growth — Alternative Parameter Choices

Calibration	Workers	Regional	Total	Frontier
<i>A. Decadal Average (1975–2025)</i>				
<i>Data</i>				
(Scaled) PWT	—	—	0.43	—
Fernald (2014)	—	—	0.32	—
BLS labor productivity	—	—	—	—
<i>Model</i>				
Baseline	0.06	0.00	0.07	0.02
$\xi = 1$	0.06	-0.01	0.05	0.02
$\theta = .5$	0.06	0.01	0.07	0.02
$\theta = 2$	0.07	-0.02	0.05	0.02
$\rho = 0$	0.06	0.16	0.22	0.01
20% Foreign	0.07	0.03	0.10	0.01
$\sigma = 13$	0.04	-0.06	-0.01	0.03
$\sigma = \infty$	0.11	0.10	0.21	0.01
<i>B. Dot-Com Decade (1995–2004)</i>				
<i>Data</i>				
(Scaled) PWT	—	—	1.44	—
Fernald (2014)	—	—	1.41	—
BLS labor productivity	—	—	1.46	—
<i>Model</i>				
Baseline	0.26	0.09	0.36	0.03
$\xi = 1$	0.25	0.07	0.32	0.03
$\theta = .5$	0.25	0.12	0.37	0.03
$\theta = 2$	0.28	0.05	0.34	0.03
$\rho = 0$	0.29	0.28	0.58	0.02
20% Foreign	0.27	0.13	0.41	0.02
$\sigma = 13$	0.21	0.01	0.22	0.05
$\sigma = \infty$	0.40	0.19	0.59	0.02

TFP-growth contributions and technology-frontier growth above the initial steady state, together with TFP data from Penn World Tables (PWT) and [Fernald \(2014\)](#) and labor productivity from the BLS. PWT growth is scaled by 0.5/0.3 to yield an initial growth rate around 0.5% per year; the PWT and Fernald rows subtract the 0.5% initial BGP growth rate. The BLS row in Panel B reports the change in annual nonfarm-business labor-productivity growth from 1975–1985 to 1995–2005; it is not reported in Panel A. Each model row re-calibrates the full model under the external parameter choice indicated in Appendix Table [E.2](#).

E.2. MECHANISMS

Figure E.2: Alternative Models: Demographic Boom



Notes: Each line feeds the calibrated demographic boom through a model variant that shuts down one channel: adoption differences across generations, the regional adoption response, or frontier movement.


E.3. DEMOGRAPHIC HEADWINDS TO ARTIFICIAL INTELLIGENCE?

Figure E.3 compares the response to an unanticipated shock to innovation and regional-adoption productivity under three demographic paths. Pre-shock decisions coincide with those along the corresponding no-shock path, while agents reoptimize from the date of the shock onward.

I consider a high-entry-rate balanced growth path, an early demographic transition completed immediately before the shock, and a late transition beginning when the shock arrives. Each shock path is compared with a no-shock economy following the

same demographic path. The frontier-growth responses are nearly identical on impact and remain close thereafter. Demographics generate somewhat larger differences in subsequent household and regional adoption, producing modest variation in the timing and peak of the TFP response. Overall, demographics alter the economy's underlying growth path but neither substantially amplify nor attenuate the productivity shock in this calibration. This limited interaction reflects the constant direct elasticities of regional adoption and frontier growth to their respective cost shifters; the remaining differences arise through the endogenous transition of adoption, profits, and technology stocks.

Figure E.3: Demographic Transitions and an AI Productivity Shock




Notes: The figure reports responses to an unanticipated, permanent reduction in the innovation-cost shifter χ and the regional-adoption cost shifter κ . "High n " remains on the high-entry-rate balanced growth path; "Early transition" completes the decline in the entry rate immediately before the shock; and "Late transition" begins the decline when the shock occurs. Panels A–D show paths under the shock, while panels E–F report differences between the shock and no-shock paths within each demographic scenario. The horizontal axis reports decades since the shock.

E.4. DYNAMICS AROUND TRANSITORY AND PERMANENT SHOCKS

Figure E.4 reports the response to a permanent halving of the labor-force entry rate. As smaller cohorts enter, the working-young share and the adoption of new technologies decline, while average adoption across all technologies rises as experienced workers become more prevalent. These forces generate a nonmonotonic TFP-growth response: workforce composition initially raises growth, partly offset by weaker regional adoption;

as this composition effect fades, growth temporarily falls below trend before recovering. Meanwhile, the smaller market size steadily reduces frontier growth, and the economy ultimately converges to a lower balanced-growth path.


Figure E.4: Permanent Decline in the Labor-Force Entry Rate



Notes: Response to a permanent halving of the entry rate n , in deviations from the initial balanced growth path. The horizontal axis reports model periods (decades).

Figure E.5 reports the response to a one-period, five-fold increase in the entry rate. The resulting youth wave raises new-technology adoption after a short lag but initially reduces average experience, causing TFP growth to fall on impact. Regional adoption then strengthens and, as the larger cohort gains experience, the worker contribution turns positive, producing a delayed acceleration in TFP growth. As these effects unwind, growth temporarily undershoots its initial path and gradually returns toward trend, while the frontier-growth response dissipates more slowly.


Figure E.5: One-Period Increase in the Labor-Force Entry Rate




Notes: Response to a one-period, five-fold increase in the entry rate n , in deviations from the initial balanced growth path. The horizontal axis reports model periods (decades).

F. GLOBAL DEMOGRAPHIC TRENDS


Figure F.1: Global Demographic Trends: Working-Young Share and Labor-Force Growth




(a) Working-Young Share: Countries




(b) Labor-Force Growth: Countries




(c) Working-Young Share: Income Groups



(d) Labor-Force Growth: Income Groups



(e) Working-Young Share: High-Income Regions



(f) Labor-Force Growth: High-Income Regions

Notes: The working-young share is the population aged 25–44 as a share of the population aged 25–64. Labor-force growth is the annualized growth rate of the population aged 25–64, computed over a ten-year window centered on the reported year. Income groups (high, middle, low) and the aggregates in the bottom panels (G7, OECD, EU15, high-income) follow the UN classifications. The dotted vertical line marks 2024, the start of the projection period. Source: UN World Population Prospects 2024, medium projection variant.