

Demographic Origins of Premature Deindustrialization

Marta Domínguez-Jiménez*

CEMFI

Job Market Paper

October 25, 2026

[Latest version](#)

Abstract

Many developing economies today experience limited industrialization and rapid growth in low-skill services. I show how workforce aging gives rise to this pattern of premature deindustrialization and slower economic growth. First, I document a new fact: within birth cohorts, the probability that a worker is employed in manufacturing declines sharply with age. I trace this pattern to life-cycle changes in productive abilities. Manufacturing relies more on early-peaking abilities such as physical strength and reasoning. Services rely on accumulated knowledge and interpersonal skills that improve throughout the working life. I then develop a quantitative open-economy model of structural transformation, where workers with heterogeneous abilities sort across sectors over the life cycle. I calibrate the model using five decades of Indonesian data. Younger demographics alone make the calibrated economy industrialize further and grow faster. Manufacturing's employment share peaks 4.8 percentage points higher under the younger path of an early industrializer than under the older path of a rapidly aging developing economy. This greater industrialization, supported by export opportunities, helps generate 20 percentage points more cumulative growth in real output per worker.

*Marta Domínguez-Jiménez: marta.dominguez@cemfi.edu.es. I thank Sebastian Fanelli, Nezih Guner, Josep Pijoan-Mas and Tommaso Porzio for their guidance and comments, as well as the participants of the CEMFI Macroeconomics Workshop, the Columbia Business School Macro Lunch, the RCEA International Conference 2026, the XXIX Workshop on Dynamic Macroeconomics in Vigo 2026, the Royal Economic Society Annual Conference 2026 and the EEA-ESEM Congress 2026. This project received the support of a fellowship from the "la Caixa" Foundation (ID 100010434), fellowship code LCF/BQ/DR23/12000010.

1 Introduction

Industrialization supports rising incomes by drawing more workers into manufacturing, a sector with rapid productivity growth. Yet many low- and middle-income economies are experiencing premature deindustrialization: manufacturing’s share of employment begins to decline at lower levels and lower incomes than in earlier industrializers (Rodrik, 2016a). At the same time, the world is aging rapidly. At comparable stages of development, many of these economies have older workforces than earlier industrializers did. This paper quantifies the effects of workforce aging on industrialization and growth in output per worker.

Existing explanations of premature deindustrialization emphasize sectoral technological change and its interaction with demand and trade. This paper shows that workforce aging also plays a central role. By changing the mix of productive abilities, aging erodes an economy’s comparative advantage in manufacturing, limiting its capacity to industrialize. Manufacturing relies more on physical and fluid-cognitive abilities such as strength and reasoning, which peak early in life. Services rely more on crystallized (i.e. accumulated) knowledge, such as vocabulary, which continues to develop throughout the working life. These differences shift workers’ comparative advantage toward services as they age, linking workforce age-composition to manufacturing’s labor supply. Export opportunities amplify the manufacturing advantage of a younger workforce by allowing economies to specialize and expand production beyond what domestic spending alone would support. To quantify the impact of workforce aging, I develop an open-economy model of structural transformation and calibrate it using Indonesian data for 1976–2025. In the model, workers’ changing abilities interact with sector-specific requirements to shape where they choose to work and how productive they are in that sector.

I find that an economy industrializes further and grows faster with a younger workforce, even under the same sectoral productivity paths and trade conditions. Holding all non-demographic features fixed, I solve the same Indonesian economy under the younger demographic path of the historical United States and the older path of Thailand, a rapidly aging economy experiencing premature deindustrialization. I align the paths at an initial stage of structural transformation and compare fifty years of subsequent industrialization. Under the younger U.S. path, manufacturing’s employment share peaks 4.8 percentage points higher than under the older Thai path—almost half the observed U.S.–Thailand gap.¹ This helps raise cumulative growth in real output per worker by 19.5 percentage points.

The mechanism behind these quantitative results is supported by worker-level evidence from repeated cross-sections covering 88 countries. I document a new fact: within birth cohorts, the probability of working in manufacturing falls sharply with age, while that of working in services rises. Manufacturing’s share of non-agricultural employment is 21 percentage points higher among workers aged 20 than among those aged 64; the corresponding gap in its share of total employment is 12 per-

¹Using the highest manufacturing employment shares reported by Felipe, Mehta, and Rhee (2019, Table 8).

centage points. Indonesian panel data follow the same workers for over two decades and record frequent individual transitions from manufacturing to services, consistent with this age pattern. Differences in the occupations held by younger and older workers account for almost all of the manufacturing decline, suggesting that what matters is the types of work done within a sector. These worker-level patterns matter for the aggregate composition of employment: larger young-adult shares predict higher industrial and manufacturing employment shares at the country level.

To assess whether changes in abilities can explain the worker-level sorting, I construct age–ability profiles from published studies of physical and cognitive aging.² I weight these profiles by each sector’s ability requirements, measured using the Occupational Information Network (O*NET). The resulting relative-productivity curves closely match the pattern of sectoral employment by age: they fall sharply in manufacturing and rise in low- and high-skill services. I use the same age–ability profiles and sectoral requirements directly in the quantitative model.

The model embeds life-cycle worker sorting in an open economy with four sectors: agriculture, manufacturing, low-skill services, and high-skill services. Workers differ in permanent ability endowments as well as schooling. Each ability then follows a distinct age trajectory, shared across workers and constructed from the studies of aging cited above. These characteristics together with measured sectoral requirements determine workers’ productivity in each sector and shape their sectoral choices. Permanent differences in workers’ relative strengths generate persistence in sectoral choice, while ability aging and changing relative wages can induce switching without idiosyncratic taste shocks. Workers move freely among manufacturing and services, but a cost of leaving agriculture captures the mobility barriers documented in the literature (Porzio, Rossi, and Santangelo, 2022; Adamopoulos, Brandt, Chen, Restuccia, and Wei, 2024).

Worker sorting operates alongside the standard forces of structural transformation: sector-specific productivity paths, international trade, and non-homothetic CES preferences. Age-specific taste weights also allow aging to affect the sectoral composition of household demand. In an extension, I incorporate manufactured durable goods into non-homothetic CES preferences. Young cohorts build up stocks, while accumulated holdings reduce older cohorts’ need for additional durables. Purchases depend on these holdings and respond to the same income and relative-price forces that drive structural transformation.

The baseline model is calibrated to Indonesia over 1976–2025. I take the two central inputs to worker sorting as given: the age–ability profiles constructed from the literature on aging and O*NET sectoral ability requirements. With these external inputs, the model closely reproduces the sectoral employment patterns by age observed in the Indonesian data. The same economy is then solved under

²See Medeiros, de Araújo, and de Araújo (2013); Dodds, Syddall, Cooper et al. (2014); Salthouse (2009); Hartshorne and Germine (2015); Verhaeghen (2003).

the U.S., Indonesian, and Thai population-by-age paths, with all non-demographic inputs unchanged. The demographic paths are aligned by the stage of structural transformation: each path begins with the last year in which agriculture accounted for at least 60 percent of employment. The quantitative analysis measures the effects of workforce aging, explains what drives them, and examines how they depend on export opportunities and whether schooling can offset them.

Demographic differences alone generate markedly different paths of industrialization and growth, reproducing the two defining features of premature deindustrialization. In the baseline Indonesian economy, manufacturing employment peaks at 21.0 percent, compared to 18.4 percent under the older Thai path and 23.3 percent under the younger U.S. path. Real income per capita at the peak is 36.3 percent lower under Thai demographics relative to US demographics. The U.S.–Thailand gap in manufacturing’s employment share is initially small but widens over time. Over 1976–2025, manufacturing accounts for 17.1 percent of total employment on average under the U.S. path, compared with 14.2 percent under the Thai path. Most of the manufacturing employment-share gap is mirrored by a larger share of employment in low-skill services under the Thai path. These demographic differences have important consequences for long-run growth: under the Thai path, cumulative growth in real output per worker over the fifty year period is almost 20 percentage points lower.

Life-cycle worker sorting is the main driver of these differences in industrialization and growth. Ability aging changes both workers’ sectoral choices and how productive they are in their chosen sector. This mechanism accounts for around 70 percent of the peak manufacturing gap (3.4 percentage points) and more than three-quarters of the average gap (2.2 points). It also explains about three-fifths of the difference in cumulative growth in real output per worker. Age-dependent demand contributes only 0.2 percentage points to the peak manufacturing gap and 0.1 points to the average manufacturing gap. Explicitly modeling the accumulation of manufactured durables yields similar manufacturing results and preserves the dominant role of worker sorting.

Demographic comparative advantage drives much of the manufacturing gap, and relies on international trade. Export markets allow a younger workforce’s advantage in manufacturing to translate into greater production and employment than domestic spending alone would support. By contrast, trade weakens the effect of age-dependent demand on domestic manufacturing employment: exports can absorb output no longer purchased at home, while imports can meet higher domestic spending. The baseline treats Indonesia as a small open economy in agriculture and manufacturing, given its small share of world goods exports. In an alternative specification, these sectors face downward-sloping foreign demand: expanding exports requires lower prices relative to foreign goods. Under this alternative economy, the U.S.–Thailand gap in manufacturing’s peak employment is 3.1 percentage points, two thirds of its baseline value. Worker sorting still accounts for around 70 percent of this gap, while age-dependent demand’s contribution rises to 20 percent. Consistent with this mechanism, cross-

country estimates show a stronger relationship between the young-adult share and manufacturing’s employment share in more open economies.

Can greater schooling offset these disadvantages of an older workforce? Only a very substantial schooling expansion offsets the older economy’s 2025 output-per-capita disadvantage: 3.2 additional years on average for every cohort entering the workforce after 1976. Schooling closes this income gap by raising workers’ productivity and shifting employment toward high-skill services, while manufacturing’s employment share rises only modestly. Even a five-year schooling increase for the same cohorts leaves three-quarters of the manufacturing-peak gap intact.

Related literature: This paper contributes to the theoretical and quantitative literature on structural transformation (Kongsamut, Rebelo, and Xie, 2001; Caselli and Coleman, 2001; Ngai and Pissarides, 2007; Acemoglu and Guerrieri, 2008; Foellmi and Zweimüller, 2008; Duarte and Restuccia, 2010; Buera and Kaboski, 2012; Herrendorf, Rogerson, and Valentinyi, 2013; Boppart, 2014; Ngai and Petrongolo, 2017; Comin, Lashkari, and Mestieri, 2021; García-Santana, Pijoan-Mas, and Villacorta, 2021; Fan, Peters, and Zilibotti, 2023). Explanations of premature deindustrialization emphasize technological change and globalization (Rodrik, 2016b), differences in sectoral productivity growth (Huneus and Rogerson, 2024), and technology-adoption lags (Fujiwara and Matsuyama, 2024). I quantify the importance of a new labor-supply driver: workforce aging changes the supply of workers with a comparative advantage in manufacturing, affecting both industrialization and growth in output per worker. The analysis also builds on open-economy models linking trade and international specialization to structural transformation (Matsuyama, 2009; Teignier, 2018; Cravino and Sotelo, 2019; Sposi, Yi, and Zhang, 2026). I quantify how export opportunities support the manufacturing advantage of a younger workforce, allowing for specialization and production beyond domestic demand.

The paper also contributes to research on how changes in the workforce drive structural transformation (Caselli and Coleman, 2001). Porzio, Rossi, and Santangelo (2022) quantify how human-capital differences across cohorts shift labor supply out of agriculture. I focus on a different force: changes in workers’ relative abilities within their careers alter their comparative advantage between manufacturing and services as they age. The empirical analysis adds to evidence on sectoral and occupational reallocation (Bustos, Caprettini, and Ponticelli, 2016; Duernecker and Herrendorf, 2022; Donovan, Lu, and Schoellman, 2023) by documenting a systematic age pattern. Within birth cohorts, the probability of working in manufacturing declines sharply with age. I connect this pattern to occupational ability requirements and published evidence on ability aging. Related work examines task requirements and occupational mobility (Gathmann and Schönberg, 2010; Deming and Noray, 2020), age differences in adjustment to changing occupational opportunities (Autor and Dorn, 2009; Dauth, Findeisen, Suedekum, and Woessner, 2021), and cross-country differences in life-cycle wage growth

and human-capital accumulation (Lagakos et al., 2018b,a).

Research on demographic change and growth emphasizes how population age structure affects income per capita through dependency ratios and aggregate labor supply (Bloom and Williamson, 1998; Bloom, Canning, and Sevilla, 2003). I extend this perspective to how age composition within the workforce shapes its allocation across sectors. Related work links age composition to aggregate productivity and growth (Feyrer, 2007; Maestas, Mullen, and Powell, 2023), examines endogenous technological responses to labor scarcity (Acemoglu and Restrepo, 2022; Acemoglu, Autor, Beirne, and Scott, 2026), and links population growth to firm dynamics (Peters and Walsh, 2026); I focus on how workforce aging reshapes manufacturing through changing comparative advantage and sectoral sorting. Cai and Stoyanov (2016); Kopecky (2026) link age-dependent skills and industry skill requirements to international specialization; I study how changing comparative advantage reallocates workers between manufacturing and services over their careers. On the demand side, Cravino, Levchenko, and Rojas (2022) show how aging shifts expenditure toward services. I incorporate age-dependent demand alongside worker sorting and quantify their respective contributions to differences in manufacturing employment shares across demographic paths. An extension incorporates manufactured durable goods into the non-homothetic CES demand framework, drawing on the life-cycle accumulation mechanism studied by Fernández-Villaverde and Krueger (2011). Purchases depend on existing stocks and respond to the same income and relative-price forces that drive structural transformation.

Finally, I develop a theory of life-cycle worker sorting to quantify how workforce aging changes sectoral employment and output. Building on models of sectoral sorting (Ohnsorge and Trefler, 2007; Lagakos and Waugh, 2013; Bárány and Siegel, 2018) and complementing work on multidimensional worker-job matching (Lindenlaub, 2017) and skill accumulation (Lise and Postel-Vinay, 2020; Guvenen, Kuruscu, Tanaka, and Wiczer, 2020), the framework incorporates life-cycle changes in abilities and differences in schooling distributions across cohorts. Permanent differences in workers' abilities and schooling generate persistence in sectoral choice, while ability aging and changing equilibrium wages can induce sector switching. This mechanism operates without the idiosyncratic preference shocks included in the benchmark dynamic labor-reallocation models of Artuç, Chaudhuri, and McLaren (2010) and Caliendo, Dvorkin, and Parro (2019). Sectoral choice depends on workers' relative strengths, while aggregation retains differences in their overall productivity. I derive exact aggregation formulas that remain valid with forward-looking choices and a cost of leaving agriculture. These results make it possible to quantify how workforce aging changes both sectoral employment and the productive contributions of the workers who select into each sector.

Section 2 documents the empirical facts, Sections 3 and 4 develop and calibrate the model, Section 5 reports the demographic counterfactuals, and Section 6 concludes.

2 Age Structure as a Force in Structural Transformation

I first document workers' reallocation from manufacturing toward services as they grow older. This sorting reflects changes in comparative advantage over the life cycle: manufacturing relies relatively more on abilities that peak early, while services place greater weight on abilities that develop for longer. I then move from micro evidence to aggregate outcomes: larger young-adult cohorts predict higher industry employment shares, while larger pre-peak young-adult shares are associated with higher industrial-employment peaks at higher incomes.

2.1 Data and measurement

I combine cross-country census microdata, Indonesian panel data, and cross-country aggregate data. Appendix A describes the samples, harmonization, and variable construction.

Cross-country census data. To document how workers sort across sectors as they age, I use harmonized IPUMS International censuses covering 88 countries from 1960 to 2020 (Minnesota Population Center, 2025). For each country, census year, and worker age, I calculate the share employed in agriculture, manufacturing, low-skill services, and high-skill services. A sub-sample of this dataset containing information on workers' occupations is used to measure the occupation composition of each sector. Additionally, Indonesian IPUMS censuses provide complementary sector-by-age evidence and schooling information. IPUMS USA Full Count censuses and the Human Mortality Database provide the historical U.S. age and sector distribution.

Indonesian Panel data. The census samples are repeated cross-sections and therefore do not follow the same workers over time. I consequently also use data from the Indonesia Family Life Survey (IFLS), which follows individuals across five survey waves from 1993 to 2014 and also contains retrospective employment histories (Strauss, Witoelar, and Sikoki, 2016). The labor, wage, and schooling records allow me to study sectoral transitions and wage-age profiles within individual careers and provide the moments used to discipline the model's worker block.

Cross-country aggregate data. The aggregate analysis combines sectoral employment shares from the World Bank's *World Development Indicators* (WDI) with population age shares from the United Nations' *World Population Prospects* (WPP). The principal within-country panel covers 147 countries. Because WDI's industry category combines manufacturing with mining, utilities, and construction, I use the Groningen Growth and Development Centre's Economic Transformation Database and 10-Sector Database to isolate manufacturing where their narrower coverage permits (de Vries et al., 2021, 2015). The birth-rate instrument is constructed from the historical crude birth rates in Delventhal, Fernández-Villaverde, and Guner (2024). Income and the remaining macroeconomic covariates come from the Maddison Project Database, WDI, Penn World Table, the United Nations Development

Programme’s Human Development Reports, and the Wittgenstein Centre. Finally, the model’s trade, sectoral-accounting, and world-price inputs are constructed from CEPII’s BACI merchandise-trade database, the NBER–United Nations World Trade Flows, OECD input–output tables, UNIDO industrial accounts, WTO and OECD–WTO services-trade data, the World Bank agricultural commodity-price index, and GGDC sectoral deflators.

2.2 Why is manufacturing young?

Individual abilities follow different productivity profiles over the life cycle: some peak early in adulthood while others continue improving for workers until they are past retirement. Since different sectors reward different abilities, workers sort across sectors as they age. Concretely, comparative advantage shifts from manufacturing toward services along the life cycle because manufacturing relies more heavily on abilities that peak early in life. This interpretation connects two relationships taken from external sources. First, research on cognitive and physical aging shows that abilities follow distinct life-cycle profiles. Physical abilities peak early in the working life (Dodds, Syddall, Cooper et al., 2014; Medeiros, de Araújo, and de Araújo, 2013), as do fluid-cognitive abilities like processing speed and reasoning (Salthouse, 2009; Hartshorne and Germine, 2015). On the other hand, crystallized-and-interpersonal abilities, for example written and oral comprehension or vocabulary, develop for longer (Verhaeghen, 2003; Salthouse, 2009; Hartshorne and Germine, 2015). Second, I use Occupational Information Network (O*NET) data to measure ability requirements by sector. I find manufacturing is relatively intensive in abilities that peak early, while services place greater weight on crystallized knowledge and interpersonal skills that improve throughout most of the working life.

Together, these relationships yield four empirical predictions: 1) within birth cohorts, manufacturing’s employment share should decline with age; 2) the corresponding reallocation from manufacturing to services should appear within individual careers; 3) occupations requiring early-peaking abilities should employ disproportionately younger workers, and much of the sectoral age gradient should be explained by the detailed occupations that make up the sector; and 4) sectoral employment shares over the life-cycle should follow the age patterns predicted by ability aging and sectoral requirements. The evidence supports all four predictions.

Prediction 1: Manufacturing’s employment share declines sharply with age within birth cohorts. Let $s_{c,t,a}^j$ be sector j ’s share of non-agricultural employment among workers of age a in country c and census year t . I estimate

$$s_{c,t,a}^j = \sum_{a' \neq 50} \alpha_{a'} \mathbf{1}[a = a'] + X'_{c,t,a} \beta + \gamma_1 \ln \text{GDP}_{ct} + \gamma_2 (\ln \text{GDP}_{ct})^2 + \mu_{c,b(t,a)} + \varepsilon_{c,t,a}, \quad (1)$$

where the age a fixed effects for every individual year are the main coefficients of interest. To separate life-cycle sorting from fixed differences across generations, I use country-five-year birth-cohort fixed

effects $\mu_{c,b(t,a)}$.³ A quadratic in log GDP per capita accounts for changes over time associated with economic development. $X_{c,t,a}$ includes the shares of male workers, urban residents, and workers previously living in another region or country, together with educational-attainment shares. Standard errors are clustered by country. Identification combines cohort groups observed at different ages across censuses with comparisons across adjacent single ages within each cohort bin. The identifying assumption is that, conditional on the cohort fixed effects and the GDP quadratic, residual period shocks do not vary systematically by age.

Figure 1 plots the coefficients on each of the age fixed effects, $\alpha_a \forall a$, for each of the three sectors.⁴ The figure documents a new empirical regularity: manufacturing’s share of employment declines sharply with age within birth cohorts. Within non-agricultural employment, the decline between ages 20 and 64 is 20.9 percentage points (s.e. 2.7), from 0.45 to 0.24. Using the same specification with total employment as the denominator, the manufacturing share instead falls by 12.1 percentage points (s.e. 2.1; Appendix Figure 26), representing a decline of over 50 percent.




Figure 1: Within-cohort sectoral age profiles. *Notes:* Non-agricultural employment shares from equation (1), recentered to their employment-weighted sample means. Shading shows manufacturing’s pointwise 95 percent confidence interval; standard errors are clustered by country.

Low-skill services employment rises from 0.32 to 0.42 between the ages of 20 and 64, an increase of 10.1 points (s.e. 3.0). High-skill services rise from 0.21 to 0.31, an increase of 9.6 points (s.e. 3.6). The shape of the curves survives alternative approaches to the age–period–cohort problem, alternative employment-share denominators, and restrictions by country income level (Appendix B). Manufacturing is more concentrated among young workers for both men and women and in both

³Where $b(t, a) = 5 \lfloor (t - a)/5 \rfloor$.

⁴The α_a coefficients are recentered such that each line’s employment-weighted mean equals the corresponding estimation-sample mean.

urban and rural areas (Appendix B.2).

A separate concern is that expanding sectors may employ younger workers simply because they hire more new entrants to the labor market. This explanation implies that the faster-growing sector should have the younger workforce. In these censuses, however, services employment grew faster than manufacturing, yet service workers were older. Sectoral growth is therefore, if anything, a mitigating factor. The decline with age also remains pronounced whether manufacturing’s aggregate employment share is rising or falling over time: between ages 20 and 64, manufacturing’s share of non-agricultural employment falls by 21.4 percentage points in the expanding group and 25.8 points in the contracting group (Appendices B.4.1 and B.4.2).

The census profiles reveal a pronounced shift from manufacturing toward services as cohorts age. I now turn to Indonesian panel data to establish whether the same reallocation occurs within individual careers.

Prediction 2: Many workers transition from manufacturing to services over their careers.

The census evidence presented above observes birth-cohort groups at different ages. I turn to the Indonesia Family Life Survey, which follows workers across successive interviews from 1993 to 2014, to verify that individual worker transitions underlie the sorting patterns observed. For each pair of interviews, I observe a worker’s sector at the beginning and end of the interval. Among workers employed outside agriculture at both dates, I estimate separately the probability that a manufacturing worker moves to services and that a service worker moves to manufacturing.

Let i index worker–wave-pair observations and let a_i denote the worker’s age band at the beginning of the interval: 20–24, 25–34, 35–44, 45–54, or 55–64. For each directional transition $d \in \{M \rightarrow S, S \rightarrow M\}$, I estimate

$$T_i^d = f^d(a_i) + X_i' \gamma^d + \lambda_{p(i), q(i)}^d + \varepsilon_i^d, \quad (2)$$

where T_i^d indicates the stated transition and $f^d(a_i)$ allows a separate effect for each age band. The indices $p(i)$ and $q(i)$ denote the worker’s province of residence at origin and the wave pair, respectively. The vector X_i contains sex, education, and urban-residence controls measured at origin. Province-by-wave-pair fixed effects compare workers of different ages within the same origin province and interview interval. They account for shifts in transition probabilities common to all ages, including those associated with local structural transformation. For each transition, I calculate what each worker’s predicted probability would be in each age group, keeping their other characteristics, province, and interview interval unchanged. I then average these predictions using the survey weights. Regressions use the IFLS person weight from the origin wave, and standard errors are clustered by worker.

The within-cohort profile in Prediction 1 shows manufacturing’s share of nonagricultural employment declining from the beginning of the working life; the IFLS reveals a pattern of individual transi-

Table 1: Predicted sectoral transition probabilities by age, Indonesia.

Age at origin	$M \rightarrow S$	$S \rightarrow M$
20–24	28.7 [25.5, 31.9]	12.2 [10.5, 14.0]
25–34	32.3 [30.2, 34.5]	9.2 [8.3, 10.0]
35–44	36.0 [33.2, 38.8]	6.4 [5.7, 7.2]
45–54	38.9 [34.9, 42.9]	4.6 [3.7, 5.6]
55–64	34.4 [27.6, 41.1]	4.4 [2.8, 6.1]
Transition observations	4,826	16,459
Distinct workers	3,197	9,485

Notes: Predicted transition probabilities (percent) among workers employed outside agriculture at both interviews. S pools both service sectors. IFLS 1993–2014; intervals span 3–7 years. Brackets give pointwise 95 percent confidence intervals, with standard errors clustered by worker. Appendix C.1 gives the calculation.

tions consistent with this decline. The predicted probability that a manufacturing worker observed at ages 20–24 will be in services by the next wave is 28.7 percent. This probability rises to 38.9 percent for workers observed in manufacturing at ages 45–54: movement into services becomes more likely through middle age. The reverse transition is much less likely even among the youngest workers: the predicted probability of moving into manufacturing falls from 12.2 percent at ages 20–24 to under 5 percent for workers observed when they are 45 or older.

Appendix C documents the probability calculations and supporting specifications, distinguishes low- and high-skill service destinations, and examines sectoral persistence. It also explores how real hourly earnings vary with age, using repeated observations of the same workers in the same sector and controlling for earnings changes shared within each sector and survey wave. The estimated age profiles rise through middle age in services, while remaining approximately flat early in working life and declining thereafter in manufacturing (Appendix C.4). Finally, the appendix compares earnings growth between employees who change sectors and those observed in their original sector at both interviews (Appendix C.5), finding no significant pattern.

Prediction 3: Occupations that require early-peaking abilities employ younger workers.

For each occupation, O*NET’s *Abilities* and *Generalized Work Activities* modules report importance and required level of each ability or activity. I classify items from both modules into three requirement bundles. Physical (P) requirements include strength, dexterity, motor coordination, and manual or equipment-intensive activities. Fluid-cognitive (F) requirements include reasoning, attention, perceptual speed, working memory, information processing, and problem solving. Crystallized-and-interpersonal (C) requirements combine learned verbal and knowledge demands with communication, teaching, care, persuasion, negotiation, coordination of others, and organizational activities. The three

scores add up to one within occupation. Appendix A documents their construction. The analysis below uses two subsets of IPUMS censuses with ISCO-08 and ISCO-88 occupation data respectively.

I ask whether occupations requiring early-peaking abilities are more common among younger workers. Let $\theta_o^{PF} = \theta_o^P + \theta_o^F = 1 - \theta_o^C$ denote the combined physical and fluid-cognitive share of occupation o 's requirements. I rank three-digit occupations by this measure and divide them into five approximately employment-balanced groups, keeping occupations intact and group membership fixed across ages. For each group, I estimate the age profile of its share of non-agricultural employment using equation (1), restricting the sample to workers with valid occupation codes. Figure 2 shows the estimated profiles as percentage-point differences from age 50.




Figure 2: Occupations requiring early-peaking abilities have younger employment profiles. *Notes:* Occupations are grouped by combined physical and fluid-cognitive requirements. Profiles show percentage-point differences from age 50 in non-agricultural employment shares among workers with valid occupation codes. Shading: pointwise 95 percent confidence intervals for the highest- and lowest-requirement groups, clustered by country. Appendix D.1.

The estimates support the prediction. In the ISCO-08 sub-sample, the employment share falls by 23.8 percentage points between ages 20 and 64 in the quintile with the highest combined physical and fluid-cognitive requirements, including garment trade workers and textile-products machine operators. The lowest-requirement group, including general office clerks and primary-school teachers, gains 12.8 points. ISCO-88 shows the same contrast: the highest-requirement group loses 13.4 points; the lowest gains 9.9 points. Appendix D.1 quantifies this relationship across individual occupations.

A complementary exercise asks whether the manufacturing age gradient in Prediction 1 reflects sorting across occupations and tasks or some other driving force. Sectors employ different mixes of

occupations, so age-related sorting across these jobs can produce the sectoral age profiles documented above. I re-estimate equation (1) in the smaller set of occupation-coded census samples, retaining only workers with a valid three-digit occupation. I then add three-digit occupation fixed effects, leaving the sample, controls, country–birth-cohort fixed effects, and weighting unchanged.

In the ISCO-08 sub-sample, the estimated decline in manufacturing’s share of non-agricultural employment between ages 20 and 64 is 31.2 percentage points. It falls to 4.7 percentage points (s.e. 0.9) after including three-digit occupation controls, an 85 percent decline. In ISCO-88, the corresponding gap falls from 14.3 to -0.5 percentage points, disappearing entirely. Differences in detailed occupations thus account for much of the estimated manufacturing age gradient, linking it to the type of work required in the sector (Appendix D.2, Table 22). Appendix D.3 extends the evidence to changes in occupational shares within manufacturing; Appendix D.4 compares young-worker shares across manufacturing divisions.

Prediction 4: Employment shares by sector and age track life-cycle changes in comparative advantage. We have seen that manufacturing’s share of non-agricultural employment falls with age, while both service sectors’ shares rise. Can changes in abilities over the life cycle explain these contrasting patterns? I combine evidence from studies of physical and cognitive aging with O*NET measures of how intensively each sector requires each ability. Using these two inputs, I calculate how productivity changes with age in each sector relative to the others. I compare the resulting curves with the employment-share profiles estimated in Prediction 1 (Figure 1), asking whether the direction and shape agree. Neither the ability curves nor the sectoral requirement weights are chosen to match those employment profiles.

Table 2 reports sectoral requirement weights constructed from O*NET, normalized to sum to one for each sector. Manufacturing places greater combined weight on physical and fluid-cognitive requirements, while both service sectors place greater weight on crystallized-and-interpersonal requirements.

Table 2: Ability requirements by sector, $\theta_{j,k}$ (O*NET)

Sector	Physical	Fluid-Cog.	Cryst. & interp.
Agriculture (A)	0.409	0.217	0.375
Manufacturing (M)	0.555	0.315	0.130
Low-skill services (L)	0.252	0.302	0.446
High-skill services (H)	0.044	0.455	0.501

I then construct age profiles for the three abilities from studies on physical and cognitive aging. The physical profile draws on flexibility and grip strength (Medeiros, de Araújo, and de Araújo, 2013; Dodds, Syddall, Cooper et al., 2014); the fluid-cognitive profile uses evidence on processing

speed and reasoning (Salthouse, 2009); and the crystallized-and-interpersonal profile uses vocabulary (Verhaeghen, 2003; Hartshorne and Germine, 2015; Salthouse, 2014).

I use the reported timing and magnitude of changes with age from these sources to construct annual ability curves over ages 20–64, first expressed in standard deviations across individuals and then converted into productivity indices, and normalize each index to have an arithmetic mean of one over ages 20–64. Figure 5 in Section 4.2 plots these profiles. Appendix E explains how each curve is constructed from the sources and documents the conversion into productivity indices.

Let $\rho_k^{\text{pub}}(a)$ denote ability k 's productivity profile and θ_{jk} its fixed O*NET weight in sector j . I calculate ability-implied sectoral log productivity $\ell_j(a)$ and subtract its mean across the three non-agricultural sectors at each age to obtain relative log productivity $r_j(a)$:

$$\ell_j(a) = \sum_{k \in \{P, F, C\}} \theta_{jk} \log \rho_k^{\text{pub}}(a), \quad r_j(a) = \ell_j(a) - \frac{1}{3} \sum_{h \in \{M, L, H\}} \ell_h(a).$$

If $r_j(a)$ falls with age for one sector and rises for another, ability aging shifts comparative advantage toward the latter.

In Figure 3, the lines show the ability-implied relative-productivity profiles $r_j(a)$, while the dots show the census employment profiles by age estimated in Figure 1. The comparison concerns direction and shape rather than amplitude, because relative efficiency and employment shares have different units. Both profiles are expressed relative to ages 50–54. Employment coefficients are shown in percentage points on the left axis, with relative log productivity on the right.




Figure 3: External evidence on ability aging predicts sectoral sorting. *Notes:* Dots: census employment coefficients (left axis, percentage points). Lines: ability-implied relative productivity (right axis, log units). Both are relative to ages 50–54. Appendix E gives the construction.

The ability-implied profiles get all three directions right: from the early twenties to the early sixties, relative efficiency declines in manufacturing and rises in both service sectors, similarly to the estimated employment profiles. The correspondence extends across age groups: across the nine age bands, the correlation between the predicted and estimated profiles is +0.997 for manufacturing, +0.905 for low-skill services, and +0.981 for high-skill services. These ability curves by sector are constructed without any worker employment data, and yet they mirror the sorting choices that workers make.

The result is robust to alternative cognitive and physical profiles. I consider 81 combinations varying the fluid-cognitive peak, the crystallized-and-interpersonal plateau, and their amplitudes, and separately nine combinations varying the rate of physical decline and its conversion into standardized units. In every case, relative productivity falls between ages 20 and 64 in manufacturing and rises in both service sectors, with correlations with the employment profiles remaining close to their baseline values (Appendix E).

Together, Predictions 3 and 4 link workers' sector of choice to the occupation composition of the sector and the distinct abilities required by those occupations. Because each ability evolves in a different way over the life-cycle and requirements vary sector by sector, a worker's comparative advantage evolves as they age. It shifts from manufacturing to services.

2.3 Age structure and manufacturing trajectories

The age patterns in sectoral employment raise an aggregate question: do countries with different age structures follow different industrialization trajectories? I examine whether larger young-adult cohorts predict higher industry employment shares within countries, and whether larger young-adult shares before the industrial-employment peak are associated with higher peaks at higher incomes.

2.3.1 Young-adult cohorts and industry employment

Let s_{it}^j denote sector j 's share of total or non-agricultural employment in country i and year t . I estimate

$$s_{it}^j = \beta_j \text{Young}_{it} + \lambda_1 \ln(\text{GDPpc}_{it}) + \lambda_2 [\ln(\text{GDPpc}_{it})]^2 + \mathbf{X}_{it}'\theta + \alpha_i + \gamma_t + \varepsilon_{it}, \quad (3)$$

where Young_{it} is the share of the total population aged 25–34, and α_i and γ_t are country and year fixed effects.

The contemporaneous young-adult share may respond to sectoral conditions, for example through migration. To isolate variation in cohort size determined before current employment outcomes, I instrument Young_{it} with the mean crude birth rate over years $t - 34, \dots, t - 25$, when the cohorts currently aged 25–34 were born. As an alternative, I age the population aged 5–14 in $t - 20$ forward using country-specific survival probabilities and scale the surviving cohort by the predicted population aged 20 and above. The two instruments measure the history of the same cohorts.

For identification, I rely on the assumption that past birth rates and childhood cohort sizes affect the sectoral composition of employment only through the current proportion of young adults in the population. A quadratic in current log GDP per capita controls for the relationship between development and sectoral employment. The vector $\mathbf{X}_{it}$ additionally contains female labor-force participation, mean years of schooling among individuals aged 25 and above, imports relative to GDP, the log capital–output ratio, the investment share, and log TFP. The capital and productivity controls are motivated by evidence that a declining supply of young workers can induce investment in labor-saving technologies (Acemoglu, Autor, Beirne, and Scott, 2026). Finally, country fixed effects control for permanent differences across countries while year fixed effects absorb common shocks.

Table 3: Young-adult cohort size and sectoral employment shares

	OLS	Birth-rate 2SLS	Child-stock 2SLS
<i>Panel A. Share of total employment</i>			
Agriculture	−0.307** (0.138)	−0.258 (0.323)	−0.260 (0.207)
Industry	0.863*** (0.184)	0.746** (0.313)	0.730*** (0.180)
Services	−0.556*** (0.211)	−0.487 (0.336)	−0.470** (0.184)
<i>Panel B. Share of non-agricultural employment</i>			
Industry	0.840*** (0.210)	1.027*** (0.350)	0.696*** (0.226)
Full controls	Yes	Yes	Yes
Country and year FE	Yes	Yes	Yes
Sample years	1991–2016	1991–2016	1991–2016
Observations	3,650	3,650	3,650
Countries	147	147	147
Kleibergen–Paap rk Wald F statistic	—	87.1	186.4

Notes: Coefficients give employment-share percentage-point changes per one-percentage-point increase in the population share aged 25–34. Industry includes manufacturing, mining, utilities, and construction. Country-clustered standard errors are in parentheses. Specifications: equation (3) and Appendix G.1. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 3 exhibits the coefficients on the young adult share under the different specifications. The birth-rate 2SLS estimate implies that a one-percentage-point larger young-adult share is associated with a 0.75-percentage-point larger industry share of total employment and a 1.03-point larger share within non-agricultural employment. The child-stock estimates are 0.73 and 0.70 points, respectively. Appendix Table 27 reports all non-fixed-effect coefficients underlying Table 3.

Appendix G provides supporting evidence for this specification and its identifying assumptions. Both cohort instruments strongly predict the young-adult share and give similar coefficients; the esti-

mates also remain significant under weak-instrument-robust inference. When seven lagged birth-rate windows enter the first stage jointly, positive coefficients are concentrated in the windows corresponding to the current 25–34 population. Because some contemporaneous controls may themselves respond to demographic change, I also estimate the relationship with country and year fixed effects alone.⁵ With either cohort instrument, the estimate for industry’s share of total employment remains positive and close to one (0.97 and 1.00); Appendix G.3 reports these estimates and the corresponding reduced forms. The employment relationship is stable when conditioning on mortality and income when the cohorts were born, allowing annual shocks to differ across World Bank regions, and removing countries one at a time. Controlling for birth-window income and its growth gives a coefficient of 0.68 for industry’s share of total employment, compared with 0.66 on the same sample; adding region–year effects gives 0.82, compared with the 0.75 baseline.

The relationship is also present under narrower definitions of manufacturing. Using a smaller panel of countries from the GGDC, the birth-rate 2SLS coefficients are 0.926 for manufacturing’s share of total employment and 1.266 for its share of non-agricultural employment (Appendix G.5). I also explore the relationship at the sub-national level: across 8,316 region–years in 1,463 regions and 72 countries, a one-percentage-point larger young-adult share is associated with a 0.381-point larger manufacturing share after region and country–year fixed effects (Appendix G.6).

Finally, trade can give a younger workforce greater scope to specialize in manufacturing by allowing production to expand beyond domestic demand. In order to explore this relationship, I extend the baseline instrumental-variable specification by interacting the young-adult share with trade openness, measured as exports plus imports of goods and services relative to GDP. With the birth-rate instrument, a twenty-percentage-point higher trade/GDP ratio is associated with a 0.48-percentage-point larger response of industry’s share of total employment to a one-percentage-point increase in the young-adult share. The interaction remains positive and significant at the one-percent level when using the alternative cohort instrument or running the regression in the GGDC panel for manufacturing specifically (Appendix G.7). The positive interaction also remains when the additional economic controls are omitted; Appendix G.7.1 reports these estimates and the corresponding reduced forms.

2.3.2 Age structure before the industrial-employment peak

I next examine whether the young adult share in the years leading up to the industrial employment peak is associated with the two outcomes that define premature deindustrialization: the peak share of industrial employment and income per capita at that peak (Rodrik, 2016b). For each country i , let t_i^{peak} denote the year in which the industry employment share reaches its maximum and let s_i^{peak}

⁵See Angrist and Pischke (2009, Section 3.2.3) on how conditioning on variables that may themselves be affected by the treatment can introduce selection bias.

denote the corresponding share. Keeping only countries that peak within the sample, I estimate

$$s_i^{\text{peak}} = \beta_L^s \overline{\text{Young}_i}^{(-L)} + \mathbf{Z}_{i,s}^{(-L)'} \lambda^s + \eta^s t_i^{\text{peak}} + \delta_{d(i)}^s + \varepsilon_i^s. \quad (4)$$

Here, $\overline{\text{Young}_i}^{(-L)}$ is the population share aged 25–34 averaged over the $L \in \{5, 10, 15\}$ years strictly preceding the peak; the peak year is excluded. The vector $\mathbf{Z}_{i,s}^{(-L)}$ contains pre-peak GDP per capita growth, average log GDP per capita and its square, and average imports relative to GDP. The term $\delta_{d(i)}^s$ denotes peak-decade fixed effects.

For income at the industrial-employment peak, I estimate

$$\ln(\text{GDPpc})_i^{\text{peak}} = \beta_L^y \overline{\text{Young}_i}^{(-L)} + \mathbf{Z}_{i,y}^{(-L)'} \lambda^y + \eta^y t_i^{\text{peak}} + \delta_{d(i)}^y + \varepsilon_i^y, \quad (5)$$

where $\mathbf{Z}_{i,y}^{(-L)}$ contains pre-peak GDP per capita growth and average imports relative to GDP. Both equations give countries equal weight, include a linear peak-year control and peak-decade fixed effects, and are estimated without a constant. I report heteroskedasticity-robust standard errors.

Table 4: Pre-peak age structure and industrial-employment peaks

	5-year window	10-year window	15-year window
<i>Panel A. Peak industry employment share</i>			
One-percentage-point larger young-adult share	1.106*** (0.288)	1.438*** (0.410)	1.650*** (0.485)
<i>Panel B. Log GDP per capita at the peak</i>			
One-percentage-point larger young-adult share	0.177*** (0.025)	0.220*** (0.029)	0.261*** (0.041)
Peak-height controls	Yes	Yes	Yes
Peak-income controls	Yes	Yes	Yes
Peak-year control and peak-decade FE	Yes	Yes	Yes
Observations	70	68	62

Notes: Coefficients are per one-percentage-point larger population share aged 25–34, averaged over the indicated pre-peak window: percentage points in Panel A and log units in Panel B. The peak year is excluded. Robust standard errors are in parentheses. Appendix H.1. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

A one-percentage-point larger average pre-peak young-adult share in the 10 years preceding the peak is associated with a 1.44-percentage-point higher industrial-employment peak and 0.220 log units higher GDP per capita at the peak, equivalent to approximately 25 percent. Across the five- to fifteen-year windows, the peak-share estimates range from 1.11 to 1.65 percentage points, while the peak-income estimates range from 0.18 to 0.26 log units, or approximately 19 to 30 percent. Appendix Table 37 reports all non-fixed-effect coefficients underlying Table 4.

Appendix H.2 replicates the analysis after reconstructing the peak using a centered five-year moving average rather than the raw annual maximum. Appendix H.3 uses the young-adult share averaged within a certain GDP-per-capita band, rather than over the lead-up to the peak, as the main indepen-

dent variable of interest. Appendix H.4 conditions on pre-peak capital deepening and TFP growth. The results are robust under all these exercises.

Appendix H.5 replicates the analysis for manufacturing only using GGDC data. A one-percentage-point larger young-adult share is associated with a 2.6–3.6-percentage-point higher manufacturing peak. I also consider an intermediate aggregate comprising manufacturing, construction, and utilities. In the baseline ten-year specification, manufacturing accounts for 91.5 percent of its peak-height coefficient. Finally, Appendix H.6 reports results using non-agricultural employment as the denominator, HC3 standard errors, and exact leave-one-country-out estimates. The coefficient remains positive and statistically precise throughout.

Recap. The empirical evidence traces a consistent pattern from life-cycle behavior to aggregate industrialization. Manufacturing’s share of employment within cohorts declines with age, with the same reallocation visible within individual careers. The task and ability evidence supports an age-varying comparative-advantage mechanism driving sectoral sorting: manufacturing relies more on physical and fluid-cognitive abilities, which peak early in the working life, while services rely more on crystallized-and-interpersonal skills, which improve until retirement. At the aggregate level, larger young-adult cohorts predict higher industry employment shares within countries, while larger pre-peak young-adult shares are associated with higher industrial-employment peaks at higher incomes. The model quantifies how age structure affects manufacturing and output in general equilibrium.

3 Model

I develop an open-economy overlapping-generations model to quantify how aging during industrialization changes manufacturing employment and output in general equilibrium. The model embeds the worker-sorting mechanism documented above and retains the usual forces of structural transformation: non-homothetic CES preferences, sector-specific productivity growth, and international trade.

Workers differ in permanent ability endowments and schooling. Each ability follows a distinct life-cycle productivity profile, common across workers. Sectors use these abilities with different intensities: permanent differences in workers’ relative strengths generate persistence in sectoral choice, while age-related changes in comparative advantage, together with general equilibrium forces, can induce switching without taste shocks. The model’s main analytical contribution is to make this life-cycle sorting tractable in general equilibrium: I derive exact expressions for how each cohort divides across sectors and how much its workers contribute to production. I then extend household demand to include manufactured durable goods. Cohorts enter without any and accumulate holdings through purchases over the life cycle. Durable holdings enter the same non-homothetic CES demand system as non-durable consumption, linking durable accumulation over the life cycle to the income and relative-price

changes that drive structural transformation.

The environment is a small open economy populated by overlapping generations, with four sectors indexed by $\mathcal{J} \equiv \{A, M, L, H\}$, where A denotes agriculture, M manufacturing, L low-skill services, and H high-skill services. All individuals belong to a single representative dynasty that pools their resources.

3.1 Demographics

Time is discrete and a period is one year. Individuals enter the model at age $a_0 = 20$, live at most until age $a_{\max} = 89$, and are indexed by $\mathcal{A} \equiv \{20, \dots, a_{\max}\}$. Working ages are $\mathcal{A}^W \equiv \{20, \dots, \bar{a}\}$, where $\bar{a} = 64$. The model's age distribution $\{\pi_t(a)\}_{t,a}$ is exogenous and deterministic, with $\sum_{a \in \mathcal{A}} \pi_t(a) = 1$. For $a = 20, \dots, a_{\max} - 1$, whenever $\pi_t(a) > 0$, I define

$$s_t(a) \equiv \frac{\pi_{t+1}(a+1)}{\pi_t(a)}. \quad (6)$$

The term $s_t(a)$ is a cohort-accounting object: it absorbs mortality and net migration.

Workers in $\mathcal{A}^W$ supply one unit of time inelastically and choose a sector. Individuals older than $\bar{a}$ are retired; they remain members of the dynasty and consume but do not supply labor. Let $L_{j,t}$ denote the mass of workers employed in sector j . The model employment-to-population ratio is

$$L_t = \sum_{j \in \mathcal{J}} L_{j,t} = \sum_{a \in \mathcal{A}^W} \pi_t(a), \quad (7)$$

since every working-age individual is employed in one sector. Equation (7) is the headcount margin: a younger age structure raises the working-age mass relative to total population.⁶

3.2 Endowments, abilities, and schooling

Each worker carries a bundle of *physical* (P), *fluid-cognitive* (F), and *crystallized-and-interpersonal* (C) abilities. Three components determine the level of each of the three abilities for a given worker at a given age. The first is a common age profile for each ability $\rho_k(a) > 0$, $k \in \{P, F, C\}$: physical ability fades with age, fluid-cognitive ability peaks early, and crystallized-and-interpersonal ability rises throughout the working life. The other two are permanent, fixed at labor-market entry, and idiosyncratic to the individual worker: years of schooling s , which raise each log-ability with intensity α_k , and an idiosyncratic random draw ν_k for each k . A worker's ability at each age a and for each k is therefore:

$$\rho_k(a) e^{\nu_k + \alpha_k s}. \quad (8)$$

⁶This partition of cohorts—unit labor supply within the working ages, none in retirement—reflects the practically one-for-one relationship between labor force per capita and the working-age share seen in the data. Using a country panel, I regress labor force per capita on the working age share and obtain a coefficient with a 95 percent confidence interval of 0.76–0.83; see Appendix I for details.

The effect of schooling α_k differs across abilities k . It is substantial for fluid-cognitive ability (α_F), very modest for crystallized-and-interpersonal ability (α_C) and close to zero for physical ability ($\alpha_P \approx 0$). Together with the sectoral ability weights introduced in Section 3.3, this implies schooling returns that differ by sector; they are highest in high-skill services and lowest in agriculture. Within a birth cohort c , schooling is Gaussian and independent of the three idiosyncratic draws

$$s \sim \mathcal{N}(\bar{s}_c, \sigma_{s,c}^2). \quad (9)$$

A better-schooled cohort (higher $\bar{s}_c$) enters with higher mean log ability in each category, by an amount proportional to α_k .

The idiosyncratic log-ability draws ν_P , ν_F , and ν_C are jointly Gaussian, with means $\mu_{0,k}$ and covariances $\Sigma_{0,kr}$ for abilities $k, r \in \{P, F, C\}$. Therefore, the permanent log-ability components $\nu_k + \alpha_k s$ (i.e. the idiosyncratic draw plus the effect of the schooling draw) are also jointly Gaussian within each cohort, with distribution:

$$\mu_{c,k} = \mu_{0,k} + \alpha_k \bar{s}_c, \quad \Sigma_{c,kr} = \Sigma_{0,kr} + \alpha_k \alpha_r \sigma_{s,c}^2. \quad (10)$$

3.3 Sectoral ability requirements and worker productivity

Sectors differ in how intensively they use the three abilities. For each sector j let $\theta_{j,k}$ denote its weight on ability k , with

$$\theta_{j,k} \geq 0, \quad \sum_{k \in \{P, F, C\}} \theta_{j,k} = 1. \quad (11)$$

Manufacturing is intensive in physical and fluid-cognitive ability; high-skill services in fluid-cognitive and crystallized-and-interpersonal ability; low-skill services in crystallized-and-interpersonal ability.

Individual efficiency in sector j is the Cobb–Douglas aggregate of a worker’s abilities using the sectoral ability weights $\theta_{j,k}$. For a worker with idiosyncratic abilities ν_P , ν_F , and ν_C and schooling s , write efficiency as $z_j(a)$:

$$\begin{aligned} z_j(a) &= \frac{1}{Z_j} \prod_{k \in \{P, F, C\}} [\rho_k(a) e^{\nu_k + \alpha_k s}]^{\theta_{j,k}}, \\ \log z_j(a) &= -\log Z_j + \sum_{k \in \{P, F, C\}} \theta_{j,k} [\log \rho_k(a) + \nu_k + \alpha_k s], \end{aligned} \quad (12)$$

where Z_j is the sector-specific constant defined in (13). This efficiency has constant returns in the three ability inputs, since each sector’s weights sum to one (11). Separating the age and ability components, define the age profile of sectoral efficiency

$$\tilde{z}_j(a) = \frac{1}{Z_j} \prod_{k \in \{P, F, C\}} \rho_k(a)^{\theta_{j,k}}, \quad Z_j = \frac{1}{45} \sum_{u=20}^{64} \prod_{k \in \{P, F, C\}} \rho_k(u)^{\theta_{j,k}}, \quad (13)$$

so that each sector's profile has unit mean over working ages.⁷ A sector that weights physical or fluid-cognitive ability has a $\tilde{z}_j(a)$ that peaks at early ages, while one that weights crystallized-and-interpersonal ability has a profile that stays high later in life.

3.4 Worker sorting over the life cycle

Working-age individuals supply one unit of time inelastically and choose a sector each year. Moving among manufacturing and the two service sectors is free, as is entering agriculture; leaving agriculture is costly. I first explain how workers choose sectors when moving is costless and how those choices change as they age. I then introduce the cost of leaving agriculture and show how the same approach extends to this case. Aggregating these choices gives sectoral headcount employment; weighting workers by their efficiency in their chosen sector gives effective labor.

With wage per efficiency unit $w_{j,t}$, an age- a worker earns $w_{j,t}z_j(a)$ in each sector j . When every move is free, the worker chooses the sector that pays the most:

$$j^*(a, t) = \arg \max_{j \in \mathcal{J}} \left\{ \underbrace{\log w_{j,t}}_{\text{sector wage}} + \underbrace{\log \tilde{z}_j(a)}_{\text{age profile of efficiency}} + \underbrace{\sum_{k \in \{P, F, C\}} \theta_{j,k}(\nu_k + \alpha_k s)}_{\text{idiosyncratic draw and schooling}} \right\}. \quad (14)$$

The first two terms form the sector return, $\log w_{j,t} + \log \tilde{z}_j(a)$, common to every worker of age a , while the third is heterogeneous across workers. This is a multi-sector Roy problem: workers self-select on comparative advantage, taking the job that best rewards their own mix of talents. A crystallized-and-interpersonal edge pays best in services, a physical edge in manufacturing; and because the profiles $\rho_k(a)$ change with age, the best-paying sector shifts over the life cycle even though the idiosyncratic ability and schooling draws are permanent.

The exact two-dimensional reduction. Ability weights sum to one in every sector, so a proportional increase in all three abilities raises worker earnings equivalently in all sectors. Sorting depends only on pay differences, so permanent abilities affect a worker's sectoral choice only through their relative strengths across abilities, not through their overall ability level.

To isolate these relative strengths, define the worker's average permanent log ability, $\bar{m}$, and each ability's deviation from that average, $\hat{h}_k$:

$$\bar{m} = \frac{1}{3} \sum_{k \in \{P, F, C\}} (\nu_k + \alpha_k s), \quad \hat{h}_k = \nu_k + \alpha_k s - \bar{m}.$$

The component $\hat{h}_k$ measures the worker's strength in ability k relative to her own average permanent

⁷The profile carries how efficiency varies with age; its level is absorbed by sector-specific productivity $A_{j,t}$ for agriculture and the two service sectors, and by the normalization of manufacturing productivity to one in 1993 (Section 4.4).

log ability, not relative to other workers. Since $\nu_k + \alpha_k s = \bar{m} + \hat{h}_k$, log earnings can be written as

$$\log(w_{j,t} z_j(a)) = \underbrace{\bar{m}}_{\text{common across sectors}} + \log w_{j,t} + \log \tilde{z}_j(a) + \sum_{k \in \{P,F,C\}} \theta_{j,k} \hat{h}_k. \quad (15)$$

Overall ability $\bar{m}$ enters earnings identically in every sector and cancels out in sectoral comparisons. The ability mix remains relevant because sectors place different weights on its components.

Moreover, by construction, $\hat{h}_P + \hat{h}_F + \hat{h}_C = 0$; knowing $\hat{h}_P$ and $\hat{h}_F$ determines $\hat{h}_C = -\hat{h}_P - \hat{h}_F$. Two coordinates suffice to describe the permanent worker differences relevant for sectoral choice. Each worker's ability mix can be represented by a fixed point $(\hat{h}_P, \hat{h}_F)$ on a plane. The following proposition formalizes this reduction and characterizes the sectoral choice regions; its proof is in Appendix J.1.2.

Proposition 3.1 (Exact two-dimensional sorting). Suppose each sector's ability weights sum to one, $\sum_{k \in \{P,F,C\}} \theta_{j,k} = 1$ for every $j \in \mathcal{J}$. At any age a and date t :

1. Sectoral choice in (14) depends on the worker's endowments and schooling only through the two ability-mix coordinates $(\hat{h}_P, \hat{h}_F)$.
2. The ability mixes $(\hat{h}_P, \hat{h}_F)$ for which sector j pays at least as much as every other sector form the convex region

$$R_j(a, t) = \{(\hat{h}_P, \hat{h}_F) : \sum_{k \in \{P,F,C\}} (\theta_{j,k} - \theta_{\ell,k}) \hat{h}_k \geq \log \frac{w_{\ell,t} \tilde{z}_{\ell}(a)}{w_{j,t} \tilde{z}_j(a)} \text{ for every } \ell \in \mathcal{J} \setminus \{j\}\}. \quad (16)$$

The regions R_A, R_M, R_L, R_H in (16) describe sectoral choices at a given age and date: a worker whose point $(\hat{h}_P, \hat{h}_F)$ lies in region $R_j(a, t)$ earns the most in sector j and therefore chooses it. On the boundary between two regions, both sectors pay equally. The boundaries are straight lines, with slopes determined by how differently the two sectors reward workers' relative abilities.

Each worker is a fixed point that reflects their permanent ability mix; boundary slopes also remain unchanged because they are determined by the sectoral weights $\theta_{j,k}$. However, aging shifts these boundaries across workers' points as relative earnings change with the life-cycle evolution of abilities. Relative sector wages also adjust in general equilibrium and can reinforce or offset the shift caused by aging. When a boundary passes the worker's point, another sector becomes their best-paying choice. The sequence of these choices over the life cycle constitutes their career.

Cohort employment shares. A cohort's employment share in each sector is the fraction of its workers whose ability mixes lie in that sector's choice region. The Gaussian assumptions on endowments and schooling imply that $\hat{h}_P$ and $\hat{h}_F$ are jointly Gaussian (10). The share of the age- a cohort working in sector j at date t is therefore

$$\varsigma_j(a, t) = \Pr((\hat{h}_P, \hat{h}_F) \in R_j(a, t)). \quad (17)$$

The probability is taken over the cohort's ability-mix distribution; Appendix J.1.3 gives the calculation. Employment shares thus depend on how many workers lie inside each region, not on the region's size.

Figure 4 illustrates this allocation for the cohort entering the labor market in 1976, shown at ages 20 and 64. The same workers appear at the same points in both panels, but the moving boundaries change their sectoral allocation. The percentages report the cohort's employment share in each sector.




Figure 4: The sectoral allocation of a cohort changes as it ages. *Notes:* Dots show the same simulated workers at ages 20 (1976) and 64 (2020). Colors identify the highest-paying sector; percentages are full-cohort employment shares. The simulation uses the simplified version of the model with costless switching being described in the text.

Effective labor. Employment shares describe the allocation of a cohort across sectors. To determine its contribution to production, one must also account for the efficiency of the workers choosing each sector. The overall ability component $\bar{m}$, which was irrelevant in sectoral comparisons, becomes important: workers with the same ability mix choose the same sector, but those with higher overall ability supply more effective labor. Effective labor therefore depends both on how many workers choose a sector and on the average efficiency of those who choose it.

Each cohort member contributes their efficiency in sector j if they choose it, and zero otherwise. The following proposition evaluates the average of these contributions after separating out the age factor $\tilde{z}_j(a)$, which is common to all members of the cohort. To express this calculation using the sectoral choice regions, define a distribution that weights workers by their efficiency in sector j , rather than counting them equally. For any region $\mathcal{R}$ of the ability-mix plane, let

$$\Pr((\tilde{h}_P^j, \tilde{h}_F^j) \in \mathcal{R}) \equiv \frac{\mathbb{E}[z_j(a) \mathbf{1}\{(\hat{h}_P, \hat{h}_F) \in \mathcal{R}\}]}{\mathbb{E}[z_j(a)]}.$$

Both expectations are taken over the entire cohort. The denominator is the average efficiency its members would supply if everyone worked in sector j . When $\mathcal{R} = R_j(a, t)$, the ratio measures the fraction of that potential efficiency supplied by workers who actually choose j . The common age factor $\tilde{z}_j(a)$ cancels from this ratio.

Proposition 3.2 (Exact aggregation of cohort effective labor). Under the Gaussian assumptions on endowments and schooling, effective labor can be calculated using the same two-dimensional sectoral choice regions as employment shares. For birth cohort $c = t - a$, the contribution to sector j per cohort member is

$$\mathbb{E}\left[z_j(a) \mathbf{1}\{(\hat{h}_P, \hat{h}_F) \in R_j(a, t)\}\right] = \tilde{z}_j(a) J_j(c) \Pr((\tilde{h}_P^j, \tilde{h}_F^j) \in R_j(a, t)), \quad (18)$$

where $J_j(c) = \mathbb{E}[e^{\sum_{k \in \{P, F, C\}} \theta_{j,k}(\nu_k + \alpha_k s)}]$ is the cohort's average permanent efficiency in sector j and $\tilde{z}_j(a) J_j(c) = \mathbb{E}[z_j(a)]$ its average potential efficiency at age a . The efficiency-weighted distribution remains jointly Gaussian, with the same variances and covariance as the original distribution. Appendix J.1.3 gives its adjusted means and the proof.

The same two-dimensional sectoral regions therefore determine both employment and effective labor. Both calculations describe the same sectoral allocation; employment shares count each worker equally while effective labor additionally accounts for differences in worker productivity.

Leaving agriculture. The literature estimates that barriers to moving out of agriculture are substantial, and higher for older workers (Adamopoulos, Brandt, Chen, Restuccia, and Wei, 2024); these can even exceed a year of wages (Dix-Carneiro, 2014). To capture these frictions, I introduce a non-monetary cost: a worker who was employed in agriculture in the previous year and switches sector faces

$$\kappa_A(a) w_{A,t} z_A(a), \quad \kappa_A(a) \geq 0. \quad (19)$$

Here, $w_{A,t} z_A(a)$ is what the worker would earn in agriculture in the year of departure. The coefficient $\kappa_A(a)$ varies with age and measures the cost as a multiple of those earnings; it is independent of the chosen destination sector. All other sectoral moves and the initial sector choice are costless. The cost affects workers' sectoral choices but is not deducted from output or the dynasty's income.⁸

Workers evaluate each sector using its current payoff and the value of subsequent career choices. The exit cost can keep a worker in agriculture even when another sector currently pays more; anticipating future exit costs can also discourage entry into agriculture. For workers choosing nonagricultural employment, however, the comparison among manufacturing and the two service sectors still depends only on current earnings. The exit cost is the same for all three destinations, and since there are no switching costs between them, a worker who leaves agriculture will have the same future opportunities whichever destination sector they choose.

The two-dimensional representation in Proposition 3.1 is preserved because the cost of leaving agriculture is proportional to the worker's own earnings. For a given worker, doubling all three

⁸It can be read as a transfer within the dynasty: the mover bears it, but the dynasty as a whole recoups it, so it changes who leaves agriculture without using resources. The welfare comparisons in Section 5.6 accordingly do not count it as a loss.

abilities would double both earnings in every sector and the cost of leaving agriculture at every age. This changes how much the worker earns, but not their chosen sequence of sectoral choices. Within a cohort, workers with the same ability mix choose the same first sector and follow the same career, allowing sectoral employment to remain represented on the same two-dimensional plane with each worker as a fixed point $(\hat{h}_P, \hat{h}_F)$. With costly exit, the boundaries separating agriculture from the other sectors need not be straight. Pairwise earnings comparisons among manufacturing and the two service sectors retain their straight-line form. Appendix Figure 36 shows the allocation in the calibrated baseline economy.

These careers define employment regions $R_j^\kappa(a, t)$ on $(\hat{h}_P, \hat{h}_F)$. As with $R_j(a, t)$, each region contains the ability mixes of workers employed in sector j . Appendix J.1.4 sets out the worker's problem with the agricultural exit cost specified in (19). It proves that the same two ability coordinates still determine workers' careers within each cohort and derives the employment regions $R_j^\kappa(a, t)$. As in (17), the cohort's employment share is the fraction of its workers whose ability mixes lie in the corresponding sectoral region:

$$\varsigma_j(a, t) = \Pr((\hat{h}_P, \hat{h}_F) \in R_j^\kappa(a, t)). \quad (20)$$

The effective-labor formula in Proposition 3.2 remains valid with costly exit from agriculture, using the regions $R_j^\kappa(a, t)$ and retaining the contribution of overall ability. Appendix J.1.4 proves this result.

Aggregate labor supply. Weighting each cohort's contribution by its population share $\pi_t(a)$ and summing over working ages gives headcount employment,

$$L_{j,t} = \sum_{a \in \mathcal{A}^W} \pi_t(a) \varsigma_j(a, t), \quad (21)$$

and effective labor,

$$N_{j,t} = \sum_{a \in \mathcal{A}^W} \pi_t(a) \tilde{z}_j(a) J_j(t - a) \Pr((\tilde{h}_P^j, \tilde{h}_F^j) \in R_j^\kappa(a, t)), \quad (22)$$

where $t - a$ identifies the birth cohort and $\varsigma_j(a, t)$ is given by (20). Sector j 's share of employment is $L_{j,t}/L_t$; effective labor $N_{j,t}$ enters production.

At given wages and cohort sectoral choices, demographic change that places greater population weight $\pi_t(a)$ on cohorts with higher manufacturing shares will raise aggregate manufacturing employment. In general equilibrium, relative wages adjust as well, changing workers' sectoral choices.⁹

Together, these propositions deliver a tractable theory of worker sorting over the life cycle with exact cohort-level aggregation. The framework jointly determines sectoral headcount employment and the productive contribution of the workers who select into each sector. Permanent differences

⁹Each cohort's future population weight also enters workers' decisions through $s_t(a) = \pi_{t+1}(a+1)/\pi_t(a)$ (Appendix J.1.4).

in each ability link a worker's sectoral choices across ages, generating sectoral persistence even when workers are free to move. Aging changes their comparative advantage and can lead them to switch sectors, without requiring taste shocks. Both persistence and life-cycle reallocation can therefore arise from earnings comparisons alone. With costly exit from agriculture, the permanent component of comparative advantage still generates persistence among manufacturing and the two service sectors, even though workers can move freely between them.

The static Gaussian-Roy core follows Ohnsorge and Treffer (2007), with truncated-normal aggregation going back to Tallis (1961) and Heckman and Sedlacek (1985). The contribution here is to extend that core to a life-cycle framework incorporating schooling, which changes cohorts' distributions of permanent log abilities while preserving their Gaussian form; common age profiles, which shift sectoral earnings comparisons; and a cost of leaving agriculture, which, being proportional to the worker's own earnings, preserves the two-dimensional reduction when sectoral choice becomes forward-looking.

3.5 Production, trade, and prices

Production and sectoral wages. Each sector uses the effective labor supplied by workers to produce $Y_{j,t}$ units of a domestic variety under constant returns:

$$Y_{j,t} = A_{j,t}N_{j,t}, \quad j \in \mathcal{J}, \quad (23)$$

where $A_{j,t} > 0$ is sector-specific productivity. Competitive firms take the producer price $p_{j,t}^d$ and the wage per unit of effective labor $w_{j,t}$ as given; zero profits imply

$$w_{j,t} = p_{j,t}^d A_{j,t}, \quad j \in \mathcal{J}. \quad (24)$$

Producer prices and sectoral productivity thus determine the wages entering workers' sectoral choices. I use domestic manufacturing output as the numeraire:

$$p_{M,t}^d \equiv 1 \quad \text{for all } t. \quad (25)$$

Domestic producers sell their output at home or abroad. Let $Y_{j,t}^{\text{abs}}$ denote domestic output sold at home and $X_{j,t}$ denote exports, measured in units of domestic output. Market clearing requires

$$Y_{j,t} = Y_{j,t}^{\text{abs}} + X_{j,t}, \quad j \in \mathcal{J}. \quad (26)$$

Domestic purchases and imports. Households purchase a composite of domestic and imported varieties. Domestic absorption $Q_{j,t}$ therefore need not equal domestic production $Y_{j,t}$. For each sector, the composite is

$$Q_{j,t} = \left[\omega_j^{1/\eta} (Y_{j,t}^{\text{abs}})^{\frac{\eta-1}{\eta}} + (1 - \omega_j)^{1/\eta} m_{j,t}^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}, \quad \eta > 1, \quad (27)$$

where $m_{j,t}$ denotes imports delivered for domestic use, $\omega_j \in (0, 1)$ is the home-bias weight, and η is the elasticity of substitution between domestic and imported varieties. An increase in domestic absorption

can thus be met both by imports and domestic production, and the latter could also come from redirecting output that was previously exported. Therefore, it need not require increasing domestic production. The production response depends on the adjustment of prices and worker allocations in general equilibrium.

World prices $p_{j,t}^*$ are expressed in a foreign numeraire. The exchange rate e_t converts them into domestic manufacturing units. With iceberg import cost $\tau_{j,t}^M$, the import price and the price of the composite purchased by households are

$$\begin{aligned} p_{j,t}^m &= \tau_{j,t}^M e_t p_{j,t}^*, \\ p_{j,t}^Q &= \left[\omega_j (p_{j,t}^d)^{1-\eta} + (1 - \omega_j) (p_{j,t}^m)^{1-\eta} \right]^{\frac{1}{1-\eta}}. \end{aligned} \quad (28)$$

For a given level of absorption, cost minimization determines the quantities supplied by domestic and imported varieties. The division of expenditure between them depends on their relative prices as well as home bias. Appendix J.2.1 gives the corresponding demands.

Exports. For agriculture and manufacturing, I use the small-open-economy approximation whereby exports can be sold at prevailing world prices. When a sector $j \in \{A, M\}$ exports, its domestic producer price equals the world price converted into domestic units and net of the export iceberg cost τ_j^X ; exports absorb the difference between domestic production and domestic use:

$$p_{j,t}^d = \frac{e_t p_{j,t}^*}{\tau_j^X}, \quad X_{j,t} = Y_{j,t} - Y_{j,t}^{\text{abs}}, \quad j \in \{A, M\}. \quad (29)$$

If either of these sectors does not export, all of its output is sold at home and its producer price is at least the net export return $e_t p_{j,t}^* / \tau_j^X$. When manufacturing exports, its normalized producer price fixes the exchange rate at $e_t = \tau_M^X / p_{M,t}^*$; otherwise, the exchange rate adjusts jointly with other prices to satisfy market clearing and external balance. Appendix J.2.1 details these conditions.

Services are also traded, but their differentiation across locations and providers motivates downward-sloping foreign demand (Francois and Hoekman, 2010):

$$X_{j,t} = \Xi_{j,t} \left(\frac{(p_{j,t}^d / e_t) \tau_j^X}{p_{j,t}^*} \right)^{-\varepsilon}, \quad j \in \{L, H\}, \quad (30)$$

where $\Xi_{j,t}$ is an export-demand shifter and $\varepsilon > 0$ is the price elasticity of foreign demand. Exports decline as the delivered price of the domestic variety rises relative to the world price.

External balance. The model does not impose balanced trade. An exogenous wedge $\mathcal{W}_t$ specifies net exports as a share of the value of production:

$$\sum_{j \in \mathcal{J}} \left(p_{j,t}^d X_{j,t} - p_{j,t}^m m_{j,t} \right) = \mathcal{W}_t \sum_{j \in \mathcal{J}} p_{j,t}^d Y_{j,t}. \quad (31)$$

This imbalance is computed for the four model sectors—excluding oil, gas, and mining—and thus differs from Indonesia’s national imbalance. $\mathcal{W}_t < 0$ implies a trade deficit matched by an external

resource inflow to the dynasty's budget (Section 3.7); a positive value implies a surplus and outward transfer. There is no international borrowing and lending.

3.6 Preferences

On the demand side, the dynasty chooses how much each cohort consumes and how that consumption is divided across sectors. I first describe consumption without explicit durable accumulation, then extend the framework to include durable accumulation in Section 3.9. Without explicit accumulation, all four goods, including the manufactured goods that are durable in the data, are treated as consumed in the year of purchase.

Each individual aged a consumes $c_{j,t}(a)$ units of sector j 's non-durable composite at date t . These goods are purchased at absorption prices $p_{j,t}^Q$, which reflect domestic and imported varieties (Section 3.5). Per-period utility is

$$u_t(a) = \log \mathcal{C}_t(a), \quad (32)$$

where the consumption index $\mathcal{C}_t(a) > 0$ is defined implicitly by the non-homothetic CES aggregator (Comin, Lashkari, and Mestieri, 2021):

$$\sum_{j \in \mathcal{J}} \Omega(a, j)^{1/\sigma} \left(\frac{c_{j,t}(a)}{\mathcal{C}_t(a)^{\iota_j}} \right)^{\frac{\sigma-1}{\sigma}} = 1. \quad (33)$$

The elasticity of substitution $\sigma \in (0, 1)$ governs substitution across goods as relative prices change. The parameters $\iota_j > 0$ govern income effects: at given prices, demand grows faster with income for goods with higher ι_j . To capture age differences in spending patterns (Appendix F), I allow sectoral taste weights $\Omega(a, j) > 0$ to vary with age, generating differences in expenditure composition at given income and prices, as in Cravino, Levchenko, and Rojas (2022). In Section 3.9, I extend the framework to model durable-goods accumulation explicitly.

3.7 The dynasty's problem

The dynasty chooses consumption for all cohorts to maximize their population-weighted, discounted utility:

$$\sum_{t=0}^{\infty} \beta^t \sum_{a=20}^{a_{\max}} \pi_t(a) \log \mathcal{C}_t(a), \quad (34)$$

where β is the discount factor and $\pi_t(a)$ weights each age group's per-person utility by its population share. Let $E_t(a) \equiv \sum_{j \in \mathcal{J}} p_{j,t}^Q c_{j,t}(a)$ denote non-durable expenditure. With no borrowing or lending, spending must equal the pooled budget at each date:

$$\sum_{a=20}^{a_{\max}} \pi_t(a) E_t(a) \leq \sum_{j \in \mathcal{J}} w_{j,t} N_{j,t} - \mathcal{W}_t \sum_{j \in \mathcal{J}} p_{j,t}^d Y_{j,t}. \quad (35)$$

The right-hand side is labor income net of the transfer associated with the external-balance wedge $\mathcal{W}_t$ (Section 3.5); $Y_{j,t}$ denotes sectoral output and $p_{j,t}^d$ its producer price.

Non-durable expenditure composition. Different combinations of the four non-durable goods can deliver the same consumption index under preferences (33). At given prices, the dynasty chooses the combination that minimizes expenditure $E_t(a)$, yielding sectoral spending shares $\zeta_{j,t}(a)$:

$$E_t(a) = \left[\sum_{j \in \mathcal{J}} \Omega(a, j) \mathcal{C}_t(a)^{(1-\sigma)\iota_j} (p_{j,t}^Q)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}, \quad (36)$$

$$\zeta_{j,t}(a) \equiv \frac{p_{j,t}^Q c_{j,t}(a)}{E_t(a)} = \frac{\Omega(a, j) \mathcal{C}_t(a)^{(1-\sigma)\iota_j} (p_{j,t}^Q)^{1-\sigma}}{\sum_{\ell \in \mathcal{J}} \Omega(a, \ell) \mathcal{C}_t(a)^{(1-\sigma)\iota_\ell} (p_{\ell,t}^Q)^{1-\sigma}}. \quad (37)$$

The expenditure-share equations show how rising consumption changes the composition of demand. At fixed relative prices, $\partial \log \zeta_{j,t}(a) / \partial \log \mathcal{C}_t(a) = (1 - \sigma)(\iota_j - \bar{\iota}_t(a))$, where $\bar{\iota}_t(a) \equiv \sum_{\ell \in \mathcal{J}} \zeta_{\ell,t}(a) \iota_\ell$ is the share-weighted average of the non-durable ι_ℓ s. A sector's share of spending rises with consumption when its income-effect parameter ι_j lies above this average, and falls when it lies below. At the current calibration, ι_A always lies below the average, ι_H above (Section 4.3). Therefore, an increase in income and consumption shifts non-durable spending out of agriculture and toward high-skill services.

The consumption index. The expressions above give expenditure and its sectoral composition for any given consumption index. The dynasty chooses each cohort's index by allocating its pooled resources so that, at each date, an additional unit of expenditure yields the same utility gain at every age. Let λ_t denote the current-value multiplier on the pooled budget (35); the first-order condition for each cohort's consumption index is

$$\frac{1}{\mathcal{C}_t(a)} = \lambda_t \frac{\partial E_t(a)}{\partial \mathcal{C}_t(a)} = \lambda_t \frac{E_t(a)}{\mathcal{C}_t(a)} \bar{\iota}_t(a). \quad (38)$$

The last equality follows by differentiating the expenditure function (36). When spending is concentrated on goods whose demand grows more strongly with income, $\bar{\iota}_t(a)$ is higher: attaining that increase in the consumption index requires a larger proportional increase in expenditure. Together with the pooled budget, these conditions determine consumption across cohorts.

3.8 Equilibrium

Goods market clearing. Household purchases determine total absorption. In the baseline,

$$Q_{j,t} = \sum_{a=20}^{a_{\max}} \pi_t(a) c_{j,t}(a), \quad j \in \mathcal{J}. \quad (39)$$

Domestic production is allocated between domestic use and exports according to (26).

Given the exogenous paths, parameters, and initial conditions, a competitive equilibrium consists of price and allocation paths such that firms maximize profits, workers choose their sectoral careers optimally, and the dynasty chooses consumption to maximize (34) subject to its pooled budget (35). Labor and goods markets clear, and trade satisfies the conditions above, including the external-balance condition (31). Appendix J.2.2 gives the full definition and equilibrium conditions.

3.9 Extension: manufactured durable goods

I extend the non-homothetic CES demand system to include manufactured durable goods. Durable holdings enter the same consumption index as the four non-durable goods, and the dynasty jointly chooses durable accumulation and non-durable consumption. The dynasty values consumption from durable holdings in current and future periods against the resources spent on new purchases. These purchases depend on the undepreciated stocks carried forward from the previous period but also the same income and relative-price effects that shape structural transformation. Conditional on the consumption index, durable stock, and prices, non-durable expenditure shares retain their standard non-homothetic CES form.

Entering cohorts must purchase any durable stock they hold; older cohorts carry undepreciated holdings acquired earlier in life. Installed durables cannot be sold or reassigned across cohorts. Alongside the age-specific tastes retained from the baseline, this stock-building motive can contribute to the higher durable expenditure shares of younger households documented in Appendix F. The age composition of the population alters the weight of cohorts with different accumulated stocks and purchase needs in aggregate demand for manufactures. Appendices J.3 and J.3.3 present the full demand system and its derivation.

4 Calibration

The model is calibrated to Indonesia over 1976–2025. A period is one year, with labor-market entry at age 20, retirement after age 64, and a terminal age of 89. The calibration combines microdata from the Indonesia Family Life Survey (IFLS) and IPUMS International, sectoral ability weights constructed from O*NET, and age–ability profiles constructed from published evidence on physical and cognitive aging. It also uses population-by-age distributions and aggregate data on sectoral employment, value added, trade flows, and world prices (Section 2.1 and Appendix A describe the underlying sources).

4.1 Demographic inputs

I construct Indonesia’s annual age distribution from the UN WPP. The population weight $\pi_t(a)$ is the share of individuals aged a among those aged 20–89. The cohort-accounting ratios $s_t(a)$ are calculated from successive annual distributions using (6). I set the annual discount factor to $\beta = 0.96$. Although results are reported for 1976–2025, the model is solved over a longer horizon to account for cohorts’ decisions before and after the reporting window.

4.2 Workers: abilities, schooling, and sectoral employment

The worker calibration takes the two central ingredients of the sorting mechanism from external evidence: the age–ability profiles and the sector ability weights. These are the same inputs used in Prediction 4 (Section 2.2). Together, they determine how ability aging changes relative productivity across these sectors. Workers’ responses to these common productivity changes also depend on their schooling, permanent abilities, and costs of leaving agriculture. I construct cohort schooling distributions from IFLS and IPUMS. I then jointly calibrate the schooling loadings α_k , the covariance of permanent abilities Σ_0 , and agricultural exit costs $\kappa_A(a)$ using Indonesian schooling–earnings relationships, sectoral employment profiles by age and changes in sectoral employment shares across IFLS waves for approximately matched cohorts.

Ability profiles and sectoral requirements. I construct the three age–ability profiles using published evidence to discipline both their life-cycle shape and the magnitude of ability changes. The physical profile follows estimated annual declines in flexibility, combined with evidence on grip-strength (Medeiros, de Araújo, and de Araújo, 2013; Dodds, Syddall, Cooper et al., 2014). For fluid-cognitive ability, I use the reported peak ages, annual rates of decline, and cumulative changes in processing speed and reasoning documented by Salthouse (2009). For the crystallized-and-interpersonal bundle, I combine the estimated vocabulary difference between younger and older adults in Verhaeghen (2003) with evidence on vocabulary’s late-life peak (Hartshorne and Germine, 2015; Salthouse, 2014). I express age-related changes in standard deviations across individuals, convert the resulting curves into positive productivity indices, and normalize each to have a mean of one over ages 20–64. Appendix E documents the construction, including the mapping from standardized scores into productivity units. Figure 5 shows the resulting profiles; combined with the O*NET weights in Table 2, they give the age component of efficiency $\tilde{z}_j(a)$ in (13).

Schooling. I allow educational attainment to differ across birth cohorts and among workers within each cohort. Each cohort’s mean years of schooling, $\bar{s}_c$, and variance, $\sigma_{s,c}^2$, are constructed from the IFLS, supplemented by IPUMS for older cohorts. The schooling distributions describe how much education workers have; the loadings α_k determine how that education raises each ability. Because sectors use abilities with different intensities, an additional year of schooling has different earnings effects across sectors. In the model, the effect of schooling on potential log earnings in sector j is $\sum_k \theta_{j,k} \alpha_k$. This reflects the effects of education on all three abilities, weighted by how intensively that sector uses each one.

To assign values to the schooling loadings α_k , I first estimate sector-specific regressions of log hourly earnings on years of schooling in the IFLS, controlling for age and survey wave. The estimated coefficients are 0.100 in manufacturing and 0.069 in agriculture. I use these as targets




Figure 5: Ability profiles used in the quantitative model. Productivity indices normalized to arithmetic mean one over ages 20–64. Appendix E gives their construction.

for the model’s sector-specific schooling coefficients and choose the loadings jointly with the other worker parameters to fit these earnings relationships and the observed sectoral employment profiles by age. The resulting schooling loadings are $(\alpha_P, \alpha_F, \alpha_C) = (0.01, 0.30, 0.03)$ in the baseline and $(\alpha_P, \alpha_F, \alpha_C) = (0.02, 0.31, 0.03)$ in the durable-goods extension.

Permanent ability differences. The employment response to the externally constructed age profiles depends on the size of life-cycle productivity changes relative to workers’ permanent differences in comparative advantage. I calibrate the distribution of these differences through Σ_0 . Together with wages and exit costs, this distribution determines how many workers change sectors as their relative productivity evolves. I calibrate the permanent-ability distribution jointly with the other worker parameters to match how sectoral employment varies across ages and cohorts. The calibrated variances of ν_P , ν_F , and ν_C are $(2.30, 0.52, 3.86)$ in the model without explicit durable accumulation and $(2.37, 0.50, 3.76)$ in the durable extension.

Agricultural employment and exit costs. The ability-based predictions closely reproduce the direction and shape of non-agricultural employment profiles, as shown in Section 2.2. For agriculture, as well as differences in human capital captured by the cohort-specific schooling distributions described above, I also account for barriers to leaving the sector. Porzio, Rossi, and Santangelo (2022) study these forces jointly: cohort differences in human capital shape sectoral choices, and mobility costs discourage agricultural exit, particularly among older workers. Adamopoulos, Brandt, Chen, Restuccia, and Wei (2024) find older workers face higher barriers to leaving agriculture in China. I introduce agricultural exit costs to account for these mobility barriers, allowing them to vary with age.

The agricultural exit-cost schedule $\kappa_A(a)$ in (19) is calibrated jointly with the schooling loadings and the covariance of permanent abilities, using employment patterns across ages and cohorts. The schedule is nonnegative and represented by smooth interpolation through twelve age-specific values. In the baseline without explicit durable accumulation, the calibrated cost rises overall from the equivalent of 1.9 months of agricultural earnings at age 20 to approximately one year at age 64. In the durables extension, it peaks at a little over one year at age 49 and then remains comparatively flat. These increasing profiles are consistent with the literature.

4.3 Households: preferences and demand

The household calibration determines the substitution elasticity σ , the income-effect parameters ι_j and the age-specific taste weights $\Omega(a, j)$, as well as ι_D and $\Omega_D(a)$ in the durables extension.

Income effects and substitution. I estimate the ι_j s and the σ in two steps, using the Engel-curve representation of non-homothetic CES demand in Comin, Lashkari, and Mestieri (2021). Using IFLS consumption data, I express spending on agriculture and each service sector as a ratio to spending on manufactured goods. I then estimate how these ratios vary with total expenditure per adult. Comparisons are made within the same province and survey wave to account for local price differences. The resulting Engel slopes determine the income-effect parameters conditional on σ ; the manufacturing normalization is $\iota_M = 1$. The second step estimates how these spending ratios respond to price changes in Indonesia over 1975–2016. I use BPS national-accounts and IFLS expenditure shares, together with GGDC sectoral price indices. Each sector’s price is expressed as a ratio of the manufacturing price. To isolate the price response, I subtract the changes in spending ratios predicted by expenditure growth using the household Engel curves estimated above. The two steps yield $\sigma = 0.72$ and $(\iota_A, \iota_M, \iota_L, \iota_H) = (0.24, 1, 0.93, 1.47)$. At given prices, these estimates imply that higher income shifts spending toward high-skilled services and away from agriculture, relative to manufacturing.

Age-specific tastes. I choose the taste weights $\Omega(a, j)$ to reproduce the age pattern of sectoral expenditure shares in the 1993 IFLS, after aligning each sector’s level with aggregate expenditure anchors, evaluated at the model’s prices and consumption levels (Appendix K.1.2). Households are divided into twelve groups according to the mean age of members aged 20 and above.

Durable-goods extension. I keep σ unchanged and re-estimate the non-durable income-effect parameters using manufactured non-durables as the reference category. I estimate the durable income-effect parameter ι_D from holdings valued at their annual user cost, rather than purchases. This is the right metric because purchases are affected by stocks accumulated previously. Age-specific tastes are also recalibrated: non-durable weights reproduce the age pattern of sectoral shares of non-durable expenditure, aligned in the same way, while $\Omega_D(a)$ matches IFLS-based targets for durable purchases

as a share of total spending. Appendix K.1.3 describes the calibration of the additional parameters required for the extension, and Appendix K.3 the expenditure targets and fit.

4.4 Sectoral productivity

Sectoral output per worker reflects both sectoral productivity and the efficiency of the workers employed there. I construct annual real value added per worker in manufacturing using the GGDC Economic Transformation Database and historical 10-Sector Database over 1976–2018 (de Vries et al., 2021, 2015), extrapolating through 2025. Under the production function (23), I obtain manufacturing productivity by dividing this series by model-implied effective labor per worker. This adjustment accounts for the age, schooling, and permanent abilities of the workers employed in manufacturing. I retain the resulting manufacturing productivity path in both specifications, normalized to one in 1993. For agriculture and the two service sectors, I calibrate productivity relative to manufacturing to match Indonesia’s annual sectoral employment shares. These relative paths are calibrated separately in the baseline and durable-goods extension; all productivity paths remain fixed across demographic counterfactuals.

4.5 Trade

The trade calibration specifies world prices $p_{j,t}^*$, import-demand parameters $(\eta, \omega_j, \tau_{j,t}^M)$, export-demand parameters $(\varepsilon, \tau_j^X, \Xi_{j,t})$, and the external-balance wedge $\mathcal{W}_t$.

World prices. I normalize manufacturing’s world price and use 1993 as the base year. World relative prices use the World Bank agricultural commodity-price index and GGDC service-sector deflators, each relative to the GGDC manufacturing deflator.

Imports. I set the elasticity of substitution between domestic and imported varieties to $\eta = 8$ (Eaton and Kortum, 2002). I fix home-bias weights ω_j at one minus each sector’s average import share of domestic absorption, using gross trade flows. An import share alone cannot identify both home bias and the import-price wedge; with ω_j fixed, I choose annual import-price wedges $\tau_{j,t}^M$ to match value-added import shares at reference prices using (100).

Exports. I set the foreign-demand elasticity for services to $\varepsilon = 8$ and choose annual export-demand shifters $\Xi_{j,t}$ to match sectoral service exports relative to GDP. I normalize export costs to $\tau_j^X = 1$, since they are not separately identified from the calibrated import wedges and export-demand shifters.

External balance. I calibrate $\mathcal{W}_t$ to net exports of the four model sectors—excluding oil, gas, and mining—as a share of Indonesian GDP (Appendix K.1.5). Negative values imply a trade deficit.

4.6 Model fit

The model reproduces the main features of the Indonesian economy: Figure 6 compares employment shares by age for the baseline and the durable-goods extension while Table 5 summarizes model fit along other dimensions. Appendix K provides additional figures and details about the comparisons.

Employment by age. The baseline economy closely reproduces the magnitude and shape of manufacturing’s employment decline across working ages while holding the age–ability profiles and sectoral ability requirements fixed at their externally constructed values. The common age component of worker productivity is constructed from studies of physical and cognitive aging weighted by O*NET ability requirements, rather than inferred from Indonesian employment patterns. With these inputs held fixed, the baseline generates a decline in manufacturing’s share of total employment of 23.1 percentage points from its early-career peak to age 64, compared with 23.7 points in the data.

The manufacturing profile is closely matched across working life: across single ages, the correlation between model and data is 0.98 in both specifications. A linear fit across ages 20–64 implies a decline of 5.2 percentage points per decade of age in the baseline and 5.4 in the durable-goods extension, compared with 5.5 in the data. The model also captures the broad age patterns in agriculture and services, although the fit is weaker for high-skill services (Figure 6; Table 5). The data estimates for high-skill services at older ages are based on small samples.




Figure 6: Sectoral employment shares by age. Markers: IFLS (1993, 1997, 2000). Solid and dashed lines: baseline and durable-goods extension, pooled over 1976–2025. Details in Appendix K.2.

Employment over time. The model closely matches annual sectoral employment shares over 1976–2025, with a mean absolute gap of 0.38 percentage points in the baseline. Manufacturing’s share of employment peaks in 2019 in both models and in the data, at 21.0 percent in the baseline, against 21.4 percent in the data.

Table 5: Model fit

Moment	Model		Data
	Baseline	Durable-goods extension	
<i>Employment by age</i>			
Manufacturing decline, peak to age 64 (pp)	23.1	23.7	23.7
Manufacturing share, change per ten years of age (pp)	−5.2	−5.4	−5.5
Manufacturing share by age, correlation with data	0.98	0.98	—
Agricultural share by age, correlation with data	0.96	0.96	—
Low-skill services share by age, correlation with data	0.88	0.88	—
High-skill services share by age, correlation with data	0.76	0.75	—
<i>Employment over time, 1976–2025</i>			
Annual sectoral employment shares, mean absolute gap (pp)	0.38	0.17	—
Peak manufacturing employment share (percent)	21.0	21.4	21.4
Year of the manufacturing peak	2019	2019	2019
<i>Expenditure by age, 1993</i>			
Durable-purchase share by age group, correlation with data	—	0.99	—
Manufactured share of spending, under 30 minus 60 and above (pp)	7.3	8.7	7.3

Notes: Employment shares use total employment; the age-profile peak is the maximum at ages 20–25. Manufactured spending includes durable purchases. Expenditure age is mean adult household age in the data and individual age in the models. Durable-purchase fit uses IFLS-based calibration targets. Appendix K gives the moment definitions and construction.

Expenditure by age. In the baseline, durable purchases are part of manufactured consumption, and the manufactured share of spending is 7.3 points higher among households under 30 than among those aged 60 and above, matching the data. The durable-goods extension reproduces the concentration of durable purchases among younger households: across the twelve age groups, the correlation between the model and the IFLS-based durable-purchase targets in 1993 is 0.99.

Earnings and schooling. I also examine sectoral earnings ratios and schooling gaps that are not directly targeted in the calibration. Average monthly earnings in high-skill services are 1.58 times those in manufacturing, compared with 1.64 in the IFLS data. The corresponding ratios for low-skill services are 1.07 and 1.06.¹⁰

The model also closely matches the schooling gap between agricultural workers and the average worker of the same age: 1.86 fewer years of schooling in the model, compared with 1.85 in the IFLS.

¹⁰The IFLS comparison uses full-time government and private employees aged 20–64; model earnings cover all workers. Using broader earnings, including self-employment income, gives data ratios of 1.66 and 0.99, respectively.

Schooling is similar in manufacturing and low-skill services: at comparable ages, low-skill-service workers have 0.4 more years of schooling in the model and 0.2 in the data. High-skill-service workers have the highest schooling in both the model and the data, but the model understates their advantage over manufacturing: 3.0 years, compared with 5.5 in the data.¹¹

5 The Economic Impact of Aging

This section uses the calibrated economy to quantify how the timing of aging affects industrialization and income. The main experiment changes a single primitive—the population-by-age path—and re-solves the general equilibrium, so every difference across demographic paths within a given economy is demographic in origin. Unless otherwise specified, results presented are from the baseline model without explicit durable accumulation.

5.1 Counterfactual design

I solve the calibrated economy under three population-by-age paths aligned at a common stage of structural transformation. For Indonesia and Thailand, I use the last year in which agriculture accounted for at least 60 percent of employment—1976 and 1992, respectively. For the United States, I use 1850, the first available census, when the agricultural share was already slightly below 60 percent. These dates are mapped to the first reporting year, 1976; outcomes are reported over 1976–2025. Each counterfactual replaces only the age distribution $\pi_t(a)$ over ages 20–89; preferences, sectoral productivities, trade objects, cohort schooling, ability weights and the cost of leaving agriculture are held at their baseline Indonesian values. Wages, prices, and allocations adjust in general equilibrium. The objective is therefore to quantify how the Indonesian economy would have evolved under different demographic timing, not to provide accounts of the comparison countries’ histories.

The three paths span the demographic timing of industrialization. The U.S. path has the largest young-adult share throughout the period, and is an example of a young early industrializer. Thailand is initially slightly younger than Indonesia, but its young adult share quickly dips below and it is rapidly aging throughout the period. To illustrate how different the age structures of the three economies are at the same stage of structural transformation, Figure 7 plots the share of young adults (aged 20–34) over the total model population (ages 20–89). The U.S. path comes from historical IPUMS censuses, while the Indonesian and Thai paths use UN WPP 2024.¹²

¹¹Schooling comparisons cover ages 20–64 and use the same observed age weights for model and data within each comparison. All reported earnings and schooling comparisons average the five IFLS waves equally.

¹²The Thai path uses the UN medium-variant projection from model year 2008 onward. Appendix L.1 details the construction of the three paths.




Figure 7: Young-adult (20–34) share of the model population (ages 20–89) under the three demographic paths, 1976–2025.

Decomposing the demographic effects. To separate the contribution of life-cycle worker sorting to the headline results, I solve four versions of both the baseline model and its durable-goods extension under each of the demographic paths (USA, IDN, THA):

1. *Full*: ability aging and age-specific preferences operate as described in Section 3. The extension also retains cohort-specific durable accumulation as per Section 3.9.
2. *Sorting only*: ability aging remains; taste weights governing sectoral demand no longer vary with age. In the durable-goods extension, the durable taste weight is age-independent, and separate cohort stocks are replaced with an economy-wide durable stock, chosen by the dynasty.¹³
3. *Demand only*: age-specific taste weights and, in the extension, cohort-specific durable accumulation remain. I remove ability aging by replacing every sector’s age-efficiency profile $\tilde{z}_j(a)$ with the same constant at every age.
4. *Neither*: both restrictions apply: sectoral age-efficiency profiles are constant and taste weights no longer vary with age. Separate cohort stocks are also replaced by the common durable stock in the extension.

Permanent ability differences, cohort schooling, and the parameters governing agricultural exit costs remain unchanged. Workers continue to choose sectors in response to earnings and demand continues to respond to income and prices; each economy is solved in general equilibrium. Appendix L.2

¹³I choose the common sectoral taste weights to reproduce the calibrated Indonesian economy’s 1993 non-durable expenditure per person and sectoral expenditure shares, evaluated at reference prices and the population-weighted mean consumption index. In the extension, this calculation also uses mean durable holdings, and the durable taste weight is its 1993 population-weighted mean. Weights remain fixed across demographic paths. Appendix L.2 details the construction.

details construction. I use a Shapley decomposition to measure the contributions of life-cycle worker sorting and age-dependent demand to the demographic gaps in both specifications (Table 7).¹⁴

5.2 Demographic timing shapes manufacturing and output

Table 6 reports the full-model outcomes under the three demographic paths.¹⁵ Under U.S. demographics, manufacturing reaches a maximum employment share of 23.2 percent in 2025, the final year of the reporting period. Under both Indonesian and Thai demographics, the peak occurs in 2019, at 21.0 and 18.4 percent, respectively. In the same calibrated economy, manufacturing’s employment share reaches an employment peak that is 4.8 percentage points higher under U.S. than under Thai demographics. This reflects more substantial industrialization in the younger economy throughout the period: manufacturing accounts for 17.1 percent of employment on average under the U.S. path, compared with 14.2 percent under the Thai path.

Figure 8 shows how the manufacturing differences develop over the transition. The U.S.–Thailand employment-share gap grows from 0.4 percentage points in 1976 to 4.9 points in 2025. The Indonesia–Thailand gap instead reaches 3.0 points in 2012 before narrowing to 2.5 points in 2025. It also plots the evolution of the other sectors, to illustrate who absorbs the missing manufacturing employment. In 2025, the smaller manufacturing employment share under the older Thai path is accompanied by greater concentration in agriculture and low-skill services. Low-skill services account for 3.4 percentage points more of total employment under the Thai path than under the U.S. path, and agriculture for 2.8 points more. High-skill services, by contrast, account for 1.3 points less (Appendix L.3).

These manufacturing differences are substantial relative to the historical variation documented in the literature. The baseline model’s 4.8-percentage-point gap between the younger U.S. and older Thai demographic paths amounts to about 43 percent of the difference between the highest recorded manufacturing employment shares reported by Felipe, Mehta, and Rhee (2019, Table 8): 26.4 percent in the United States and 15.1 percent in Thailand. Rodrik (2016b, Table 10) estimates a decline in the fitted peak manufacturing employment share from 21.5 percent before 1990 to 18.9 percent afterward, a difference of 2.6 percentage points. Sposi, Yi, and Zhang (2026, Table III) report a decline in the fitted peak manufacturing value-added share from 31.5 percent through 1990 to 28.0 percent afterward, a difference of 3.5 percentage points. These comparisons involve different sector definitions, outcomes, and empirical exercises, but provide useful benchmarks for the quantitative importance of the model’s

¹⁴For each mechanism, I calculate how its inclusion changes the demographic gap, with the other mechanism present and absent, and average the two changes (Shorrocks, 2013). The two contributions and the gap in the *Neither* economy sum to the full-model gap. Appendix L.4 gives the details.

¹⁵Manufacturing employment is measured as a share of total employment; its peak and mean are the maximum and average over 1976–2025. Unless otherwise stated, real output is valued at each path’s 1976 producer prices. Output per capita refers to the modeled population aged 20–89.




Figure 8: Sectoral employment under the three demographic paths. *Notes:* Baseline model without explicit durable accumulation. Shares of total employment; black lines show Indonesian calibration data, including constructed targets (Appendix K.2). Manufacturing includes utilities and construction; total employment excludes mining.

demographic effects.

Premature deindustrialization involves not only a lower manufacturing peak, but also a peak reached at a lower income; this is the case for the older path. When I compare output at each path's maximum over 1976–2025, real output per capita is 28.9 percent higher under the Indonesian path than under the Thai path, and 56.9 percent higher under the U.S. path than under the Thai path.¹⁶ Moreover, the manufacturing share is still rising in 2025 under U.S. demographics; therefore this difference in output per capita at peak is a lower bound.

The younger U.S. path also supports higher output per worker. Cumulative growth over 1976–2025 is 19.5 percentage points higher than under Thai demographics, while the 2025 level is 8.3 percent higher. The output-per-capita advantage is larger: cumulative growth differs by 99.6 percentage points and the 2025 level differs by 42.8 percent. This reflects the additional difference in the working-age share between the different demographic paths. The durable-goods extension gives similar results:

¹⁶For this comparison, I value production on all three paths at the Indonesian baseline's 1976 producer prices. I compare the Indonesian and Thai paths in 2019 and the U.S. path in 2025, their respective years of maximum manufacturing employment share within the reporting period. Per capita again refers to ages 20–89.

Table 6: Manufacturing employment and output under the three demographic paths

	Indonesia (baseline)	Thailand's path	U.S. path	U.S.–Thailand gap
<i>Panel A. Baseline model</i>				
Peak manufacturing employment share (percent)	21.02	18.41	23.25	4.84
Mean manufacturing employment share (percent)	16.13	14.19	17.06	2.86
Cumulative growth of output per worker (percent)	279.6	259.8	279.3	19.5
Cumulative growth of output per capita (percent)	264.4	171.3	270.9	99.6
2025 output per worker (Indonesia = 100)	100.0	96.9	104.9	8.3%
2025 output per capita (Indonesia = 100)	100.0	76.5	109.2	42.8%
<i>Panel B. Extended model with durable goods</i>				
Peak manufacturing employment share (percent)	21.38	18.76	23.52	4.76
Mean manufacturing employment share (percent)	16.49	14.57	17.39	2.82
Cumulative growth of output per worker (percent)	297.1	274.8	296.9	22.0
2025 output per worker (Indonesia = 100)	100.0	96.5	105.5	9.3%

Notes: Full-model outcomes, 1976–2025. Manufacturing shares use total employment. Output uses each path’s 1976 producer prices; per capita refers to ages 20–89. Valuing all three paths at the Indonesian baseline’s 1976 producer prices instead barely changes the 2025 level rows: by at most 0.2 index points in the baseline model and 0.3 in the extension, where the U.S.–Thailand advantages become 8.8 and 43.5 percent. The last column gives percentage-point gaps for shares and cumulative growth, but percentage U.S.-path advantages for 2025 output levels. Gaps use unrounded values.

the manufacturing-peak gap is 4.8 points, cumulative output-per-worker growth differs by 22.0 points, and its 2025 level is 9.3 percent higher under the U.S. path.

5.3 Worker sorting drives most of the manufacturing results

Life-cycle worker sorting accounts for most of the manufacturing response to demographic change under both the baseline and extended models. In the baseline, it accounts for 3.4 of the 4.8-percentage-point U.S.–Thailand difference in the peak manufacturing employment share and 2.2 of the 2.9-point mean gap. Age-dependent demand contributes only 0.16 and 0.06 points, respectively. The durable-goods extension gives the same ranking: sorting accounts for 3.3 of the 4.8-point peak gap, while age-dependent demand contributes 0.1 points (Table 7, Panels A and B).

Table 7: Decomposition of demographic gaps, baseline model and extended model with durable goods

Pair	Baseline model				Extended model with durable goods			
	Full	Sorting	Demand	Neither	Full	Sorting	Demand	Neither
<i>Panel A. Peak manufacturing employment share (percentage points)</i>								
U.S.–Thailand	4.84	3.35	0.16	1.32	4.76	3.27	0.09	1.39
Indonesia–Thailand	2.61	1.58	0.14	0.89	2.62	1.62	0.11	0.89
<i>Panel B. Mean manufacturing employment share (percentage points)</i>								
U.S.–Thailand	2.86	2.20	0.06	0.60	2.82	2.17	0.03	0.62
Indonesia–Thailand	1.94	1.46	0.04	0.44	1.92	1.43	0.03	0.45
<i>Panel C. Cumulative real output-per-worker growth (percentage points)</i>								
U.S.–Thailand	19.5	11.1	-1.1	9.6	22.0	12.4	-1.7	11.3
Indonesia–Thailand	19.9	9.5	-1.3	11.7	22.3	10.1	-1.9	14.0
<i>Panel D. Cumulative real output-per-capita growth (percentage points)</i>								
U.S.–Thailand	99.6	6.2	-1.0	94.4	105.4	6.7	-1.5	100.3
Indonesia–Thailand	93.1	4.8	-1.2	89.5	98.6	4.7	-1.6	95.5

Notes: *Sorting* and *Demand* are Shapley contributions; *Neither* is the gap with ability aging and age-dependent demand removed. Outcome definitions: Table 6; decomposition: Appendix L.4.

Removing ability aging makes manufacturing’s employment share flatter across ages, weakening

the manufacturing advantage of a younger workforce. Under Indonesian demographics, the decline in this share between ages 20 and 64 shrinks from 23.1 to 6.5 percentage points once ability aging is switched off. Figure 9 plots the trajectory of the manufacturing sector under the three demographic paths, showing results for the full model as well as the *Demand only* economy. Without ability aging, the U.S.–Thailand peak gap narrows from 4.8 to 1.5 percentage points, while the US–Indonesia gap practically disappears.¹⁷




Figure 9: Manufacturing paths with and without ability aging. *Notes:* Baseline model without explicit durable accumulation. Panel (a): *Full*; panel (b): *Demand only*, with ability aging removed. Both panels show the same Indonesian calibration data, including constructed targets. Data definitions follow Figure 8; restrictions are described in Appendix L.2.

The baseline *Neither* economy still has a U.S.–Thailand manufacturing-peak gap of 1.3 percentage points. In this economy, cohorts continue to differ in schooling, while the costs of leaving agriculture still vary with age. Both favor a larger agricultural sector in an older economy. Changes in age composition therefore continue to affect sectoral employment even without ability aging or age-dependent demand.

Ability aging also accounts for much of the older path’s greater concentration in low-skill services. In the baseline model, its Shapley contribution is 2.1 of the 3.4-percentage-point difference in this sector’s employment share between the Thai and U.S. paths in 2025. By contrast, the U.S. path’s higher high-skill-services share is already present in the *Neither* economy, which retains schooling differences across cohorts. This is consistent with the younger workforce’s greater weight on better-schooled cohorts, favoring high-skill services—the sector with the highest return to schooling in the model. In

¹⁷The figure compares the *Full* and *Demand only* economies. The Shapley contribution in Table 7 averages this effect with the corresponding comparison when age-dependent demand is absent.

fact, ability aging has a negative effect on the younger economy’s high-skill-services advantage, but not enough to offset the effects of schooling (Appendix L.4).

Ability aging affects both workers’ sectoral allocation and their productivity, so its consequences extend beyond manufacturing employment. It accounts for 11.1 of the 19.5-percentage-point U.S.–Thailand difference in cumulative output-per-worker growth over the fifty-year period (Table 7, Panel C). Output per capita differences largely reflect the importance of the working-age share, which remains higher under the U.S. path. Much of the growth advantage thus persists in the *Neither* economy, although sorting contributes 6.2 percentage points (Panel D).

Pronounced age differences in spending need not generate a large manufacturing employment response. In an open economy, additional spending on manufactured goods can be met by imports or by redirecting exports to the domestic market, rather than by expanding domestic production. The next subsection examines how opportunities to export and specialize shape the manufacturing response to demographic change.

5.4 Export opportunities and manufacturing specialization

A younger workforce has a comparative advantage in manufacturing. Trade openness allows it to specialize beyond what domestic demand alone would support. The cross-country evidence is consistent with this relationship: the positive association between the young-adult share and industrial employment is stronger in more open economies (Appendix G.7, Table 34).

Given Indonesia’s limited share of global goods exports, the benchmark treats Indonesia as a small open economy in goods markets: agriculture and manufacturing can export at prevailing world prices. I examine more restrictive export conditions by replacing price-taking exports in these sectors with downward-sloping foreign demand. Selling more abroad then requires lower prices relative to foreign goods. Import parameters and foreign demand for the two service sectors are unchanged. I calibrate the new schedules to reproduce the full Indonesian economy’s export quantities at its equilibrium prices in every year, preserving its allocation. These schedules remain fixed across demographic paths (Appendix L.5).

More restrictive export conditions reduce the manufacturing advantage of having a younger workforce. In the baseline model, the U.S.–Thailand manufacturing-peak gap falls by a third: from 4.8 percentage points in the baseline to 3.1 under restrictive trade. The mean employment-share gap also falls, from 2.9 to 2.0 percentage points (Table 8, Panels A and B; Appendix L.5).

Export opportunities also shape the relative importance of sorting and demand. Sorting’s contribution to the peak gap falls from 3.35 to 2.14 percentage points under more restrictive export conditions, but still accounts for around 70 percent of the gap. Age-dependent demand contributes 0.62 points or 20 percent of the gap, compared to 0.16 in the baseline. When export expansion is

more constrained, differences in domestic spending play a larger role in determining manufacturing employment.

Table 8: Decomposition of demographic gaps under different export assumptions

Pair	Price-taking exports in agriculture and manufacturing				Downward-sloping export demand in all sectors			
	Full	Sorting	Demand	Neither	Full	Sorting	Demand	Neither
<i>Panel A. Peak manufacturing employment share (percentage points)</i>								
U.S.–Thailand	4.84	3.35	0.16	1.32	3.11	2.14	0.62	0.35
Indonesia–Thailand	2.61	1.58	0.14	0.89	1.64	1.05	0.46	0.13
<i>Panel B. Mean manufacturing employment share (percentage points)</i>								
U.S.–Thailand	2.86	2.20	0.06	0.60	1.99	1.44	0.27	0.27
Indonesia–Thailand	1.94	1.46	0.04	0.44	1.32	0.96	0.20	0.16
<i>Panel C. Cumulative real output-per-worker growth (percentage points)</i>								
U.S.–Thailand	19.5	11.1	-1.1	9.6	21.0	12.1	-1.5	10.4
Indonesia–Thailand	19.9	9.5	-1.3	11.7	20.5	9.9	-1.6	12.1
<i>Panel D. Cumulative real output-per-capita growth (percentage points)</i>								
U.S.–Thailand	99.6	6.2	-1.0	94.4	100.9	7.4	-1.3	94.8
Indonesia–Thailand	93.1	4.8	-1.2	89.5	93.6	5.5	-1.3	89.4

Notes: Baseline model without explicit durable accumulation. Decomposition and outcome definitions follow Tables 7 and 6. Export assumptions: Appendix L.5.

The durable goods extension shows the same reduction in the manufacturing gap and greater contribution of age-dependent demand under more restrictive export conditions (Appendix Table 56).

5.5 Can higher schooling offset the effects of aging?

The preceding comparisons hold each cohort’s schooling fixed, yet fertility differences may also affect schooling if parents with fewer children invest more in each child (Becker and Lewis, 1973).¹⁸ I first examine whether allowing schooling to vary with these fertility differences changes the manufacturing and output comparisons. I use the twin-based estimate for Indonesia in Feng (2021): a loss of 0.107 years of schooling per additional sibling. For each birth cohort, I compare total fertility in the corresponding birth year on the counterfactual path with that in Indonesia. I adjust each cohort’s mean schooling by applying the estimated effect to this fertility difference, treating each additional child per woman as one additional sibling (Appendix L.7).

In 2025, this adjustment raises average schooling among workers by 0.22 years on the Thai path and lowers it by 0.19 years on the U.S. path. Manufacturing results remain largely unaffected by the adjustment: the peak gap falls from 4.84 to 4.75 percentage points and the mean gap from 2.86 to 2.79 points. In the baseline, output per capita in 2025 was 42.8 percent higher under the younger U.S. path than under the Thai path; the schooling adjustment reduces this advantage to 36.7 percent (Appendix L.7, Table 58, Panel A).

I next return to baseline cohort schooling and ask whether a larger schooling expansion on the

¹⁸Estimates of this trade-off range from no detectable effect in Norway and Israel (Black, Devereux, and Salvanes, 2005; Angrist, Lavy, and Schlosser, 2010) to moderate negative effects in China (Li, Zhang, and Zhu, 2008; Rosenzweig and Zhang, 2009).

Thai path can offset its output-per-capita disadvantage. An increase of 3.2 years in mean schooling for each Thai-path cohort born after 1956 approximately equalizes 2025 output per capita with the U.S. path. These cohorts enter the workforce from 1977 onward. This is a substantial increase: it amounts to more than half the rise in average schooling between Indonesian cohorts born in 1956 and 2005.¹⁹ Moreover, the offset works through an increase in productivity and a shift in employment towards high-skill services, where the model’s return to schooling is highest. Its effect on industrialization is modest: even five years of additional schooling for these cohorts close only one quarter of the manufacturing-peak gap (Table 9; Appendix L.7). This five-year increase raises high-skill services’ share of total employment by 4.8 percentage points in 2025, which is more than one third of its baseline size that year (Appendix L.7).

Table 9: Schooling increases on the Thai demographic path

Additional schooling (years)	Thai output per capita in 2025 (U.S. = 100)	Thai manufacturing peak (percent)	U.S.–Thailand peak gap (pp)
0	70.0	18.41	4.84
3.2	100.0	19.17	4.08
5.0	123.0	19.59	3.66

Notes: Baseline model. Schooling increases apply to Thai-path cohorts born from 1957 onward; the U.S. path is unchanged. Manufacturing peaks cover 1976–2025. Appendix L.7 gives the construction.

5.6 Aging lowers consumption at all ages

This final section examines the consumption and welfare implications of the higher output per capita under the younger path. To do so, I look at the additional spending required, at Thai prices, to give each age group the same utility they have under the U.S. path. Using the expenditure function in (36), I calculate the minimum cost for each age. I then compare the total with spending under the Thai path, using the Thai age distribution for both.

In 2025, matching U.S.-path utility at every age requires 44.4 percent more total expenditure. The required increases range from 43.7 to 44.9 percent across ages in 2025 and are positive at every age and in every year. Over 1976–2025, matching U.S.-path consumption welfare requires 22.1 percent more expenditure on average, giving each year equal weight (Appendix L.8).

The younger path also has more children to support. I therefore examine how much of its output advantage remains when children are included in the calculation of output per capita (Cutler, Poterba, Sheiner, and Summers, 1990; Lee, Mason, and members of the NTA Network, 2014). I divide output by the population aged 20–89 plus children weighted to reflect their consumption needs.²⁰ This changes only the population used in the calculation; model output and consumption allocations remain

¹⁹The exercise holds workforce entry at age 20 and working time fixed. In practice, a schooling expansion of this size could entail foregone work and production, costs not captured in the reported output gains.

²⁰Each child is counted as 0.3 of an adult (Hagenaars, de Vos, and Zaidi, 1994).

unchanged. On this basis, child-adjusted output per person is 21.9 percent higher under the U.S. path than under the Thai path in 2025, and 5.4 percent higher on average over 1976–2025.

6 Conclusion

This paper shows that demographic timing shapes the path of industrialization. Manufacturing relies relatively heavily on abilities that peak early in working life, while services place greater weight on abilities that develop for longer. As workers age, these differences shift their comparative advantage toward services. Workforce aging weakens manufacturing expansion even when sectoral productivity paths and trade conditions remain unchanged. Demographic change affects development through workers' sectoral allocation and productivity, not only through the share of the population that works.

The quantitative effects on industrialization and output are substantial. Holding the calibrated Indonesian economy otherwise fixed, the younger demographic path of the historical United States produces a peak manufacturing employment share over 1976–2025 that is 4.8 percentage points higher than under the older Thai path. This raises cumulative real output-per-worker growth over the same period by 19.5 percentage points. Opportunities to export and specialize amplify the manufacturing advantage of a younger workforce: the manufacturing-peak gap falls by about one third under more restrictive export conditions. These results identify workforce aging as a source of premature deindustrialization and slower output-per-worker growth in many of the low- and middle-income economies that are industrializing with an older workforce.

These findings are particularly relevant for current economies that are aging before they have established a substantial manufacturing sector. Future research could also examine how production adapts to the changing abilities of the workforce. Automation may reduce reliance on abilities that decline with age, altering the relationship between aging and sectoral specialization. Migration can also reshape comparative advantage by changing the age and ability composition of the workforce. In particular, the emigration of younger workers could compound demographic pressures faced by low- and middle-income origin economies by reducing the supply of workers with a comparative advantage in manufacturing.

References

- Acemoglu, D., Autor, D., Beirne, K., and Scott, A. (2026). Baby busts and growth booms: Demographic change and the macroeconomy. NBER Working Paper 35401, National Bureau of Economic Research.
- Acemoglu, D. and Guerrieri, V. (2008). Capital deepening and nonbalanced economic growth. *Journal of Political Economy*, 116(3):467–498.
- Acemoglu, D. and Restrepo, P. (2022). Tasks, automation, and the rise in U.S. wage inequality. *Econometrica*, 90(5):1973–2016.
- Adamopoulos, T., Brandt, L., Chen, C., Restuccia, D., and Wei, X. (2024). Land security and mobility frictions. *The Quarterly Journal of Economics*, 139(3):1941–1987.
- Angrist, J., Lavy, V., and Schlosser, A. (2010). Multiple experiments for the causal link between the quantity and quality of children. *Journal of Labor Economics*, 28(4):773–824.
- Angrist, J. D. and Pischke, J.-S. (2009). *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton University Press, Princeton, NJ.
- Artuç, E., Chaudhuri, S., and McLaren, J. (2010). Trade shocks and labor adjustment: A structural empirical approach. *American Economic Review*, 100(3):1008–1045.
- Autor, D. H. and Dorn, D. (2009). This job is “getting old”: Measuring changes in job opportunities using occupational age structure. *American Economic Review*, 99(2):45–51.
- Bárány, Z. L. and Siegel, C. (2018). Job polarization and structural change. *American Economic Journal: Macroeconomics*, 10(1):57–89.
- Becker, G. S. and Lewis, H. G. (1973). On the interaction between the quantity and quality of children. *Journal of Political Economy*, 81(2, Part 2):S279–S288.
- Black, S. E., Devereux, P. J., and Salvanes, K. G. (2005). The more the merrier? The effect of family size and birth order on children’s education. *Quarterly Journal of Economics*, 120(2):669–700.
- Bloom, D. E., Canning, D., and Sevilla, J. (2003). *The Demographic Dividend: A New Perspective on the Economic Consequences of Population Change*. RAND, Santa Monica, CA.
- Bloom, D. E. and Williamson, J. G. (1998). Demographic transitions and economic miracles in emerging asia. *The World Bank Economic Review*, 12(3):419–455.
- Bolt, J. and van Zanden, J. L. (2024). Maddison project database. Version 2023 [dataset]. Data-verseNL.

- Boppart, T. (2014). Structural change and the kaldor facts in a growth model with relative price effects and non-gorman preferences. *Econometrica*, 82(6):2167–2196.
- Buera, F. J. and Kaboski, J. P. (2012). The rise of the service economy. *American Economic Review*, 102(6):2540–2569.
- Bustos, P., Caprettini, B., and Ponticelli, J. (2016). Agricultural productivity and structural transformation: Evidence from Brazil. *American Economic Review*, 106(6):1320–1365.
- Cai, J. and Stoyanov, A. (2016). Population aging and comparative advantage. *Journal of International Economics*, 102:1–21.
- Caliendo, L., Dvorkin, M., and Parro, F. (2019). Trade and labor market dynamics: General equilibrium analysis of the china trade shock. *Econometrica*, 87(3):741–835.
- Caselli, F. and Coleman, II, W. J. (2001). The U.S. structural transformation and regional convergence: A reinterpretation. *Journal of Political Economy*, 109(3):584–616.
- Comin, D., Lashkari, D., and Mestieri, M. (2021). Structural change with long-run income and price effects. *Econometrica*, 89(1):311–374.
- Cravino, J., Levchenko, A. A., and Rojas, M. (2022). Population aging and structural transformation. *American Economic Journal: Macroeconomics*, 14(4):479–498.
- Cravino, J. and Sotelo, S. (2019). Trade-induced structural change and the skill premium. *American Economic Journal: Macroeconomics*, 11(3):289–326.
- Cutler, D. M., Poterba, J. M., Sheiner, L. M., and Summers, L. H. (1990). An aging society: Opportunity or challenge? *Brookings Papers on Economic Activity*, 1990(1):1–73.
- Dauth, W., Findeisen, S., Suedekum, J., and Woessner, N. (2021). The adjustment of labor markets to robots. *Journal of the European Economic Association*, 19(6):3104–3153.
- de Vries, G., Arfelt, L., Drees, D., Godemann, M., Hamilton, C., Jessen-Thiesen, B., Kaya, A. I., Kruse, H., Mensah, E., and Woltjer, P. (2021). The economic transformation database (etd): Content, sources, and methods. WIDER Technical Note 2/2021, UNU-WIDER.
- de Vries, G., de Vries, K., Gouma, R., Pahl, S., and Timmer, M. P. (2015). Ggdc 10-sector database: Contents, sources and methods. Technical report, Groningen Growth and Development Centre, University of Groningen. Version: January 2015.
- Delventhal, M. J., Fernández-Villaverde, J., and Guner, N. (2024). Demographic transitions across time and space. Working paper, June 2.

- Deming, D. J. and Noray, K. (2020). Earnings dynamics, changing job skills, and STEM careers. *The Quarterly Journal of Economics*, 135(4):1965–2005.
- Dix-Carneiro, R. (2014). Trade liberalization and labor market dynamics. *Econometrica*, 82(3):825–885.
- Dodds, R. M., Syddall, H. E., Cooper, R., et al. (2014). Grip strength across the life course: Normative data from twelve british studies. *PLoS ONE*, 9(12):e113637.
- Donovan, K., Lu, W. J., and Schoellman, T. (2023). Labor market dynamics and development. *The Quarterly Journal of Economics*, 138(4):2287–2325.
- Duarte, M. and Restuccia, D. (2010). The role of the structural transformation in aggregate productivity. *The Quarterly Journal of Economics*, 125(1):129–173.
- Duernecker, G. and Herrendorf, B. (2022). Structural transformation of occupation employment. *Economica*, 89(356):789–814.
- Eaton, J. and Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, 70(5):1741–1779.
- Fan, T., Peters, M., and Zilibotti, F. (2023). Growing like India—the unequal effects of service-led growth. *Econometrica*, 91(4):1457–1494.
- Feenstra, R. C., Inklaar, R., and Timmer, M. P. (2015). The next generation of the penn world table. *American Economic Review*, 105(10):3150–3182.
- Felipe, J., Mehta, A., and Rhee, C. (2019). Manufacturing matters...but it’s the jobs that count. *Cambridge Journal of Economics*, 43(1):139–168.
- Feng, N. (2021). The effect of sibling size on children’s educational attainment: Evidence from Indonesia. *ECNU Review of Education*, 4(4):830–856.
- Fernández-Villaverde, J. and Krueger, D. (2011). Consumption and saving over the life cycle: How important are consumer durables? *Macroeconomic Dynamics*, 15(5):725–770.
- Feyrer, J. (2007). Demographics and productivity. *Review of Economics and Statistics*, 89(1):100–109.
- Foellmi, R. and Zweimüller, J. (2008). Structural change, Engel’s consumption cycles and Kaldor’s facts of economic growth. *Journal of Monetary Economics*, 55(7):1317–1328.
- Francois, J. and Hoekman, B. (2010). Services trade and policy. *Journal of Economic Literature*, 48(3):642–692.
- Fujiwara, I. and Matsuyama, K. (2024). A technology-gap model of “premature” deindustrialization. *American Economic Review*, 114(11):3714–3745.

- García-Santana, M., Pijoan-Mas, J., and Villacorta, L. (2021). Investment demand and structural change. *Econometrica*, 89(6):2751–2785.
- Gathmann, C. and Schönberg, U. (2010). How general is human capital? a task-based approach. *Journal of Labor Economics*, 28(1):1–49.
- Guvenen, F., Kuruscu, B., Tanaka, S., and Wiczer, D. (2020). Multidimensional skill mismatch. *American Economic Journal: Macroeconomics*, 12(1):210–244.
- Hagenaars, A. J. M., de Vos, K., and Zaidi, M. A. (1994). *Poverty Statistics in the Late 1980s: Research Based on Micro-data*. Office for Official Publications of the European Communities, Luxembourg.
- Han, J. and Suen, W. (2011). Age structure of the workforce in growing and declining industries: Evidence from Hong Kong. *Journal of Population Economics*, 24(1):167–189.
- Hartshorne, J. K. and Germine, L. T. (2015). When does cognitive functioning peak? the asynchronous rise and fall of different cognitive abilities across the life span. *Psychological Science*, 26(4):433–443.
- Heckman, J. J. and Sedlacek, G. (1985). Heterogeneity, aggregation, and market wage functions: An empirical model of self-selection in the labor market. *Journal of Political Economy*, 93(6):1077–1125.
- Herrendorf, B., Rogerson, R., and Valentinyi, Á. (2013). Two perspectives on preferences and structural transformation. *American Economic Review*, 103(7):2752–2789.
- Huneus, F. and Rogerson, R. (2024). Heterogeneous paths of industrialization. *The Review of Economic Studies*, 91(3):1746–1774.
- International Labour Organization (2026). ILOSTAT database. Accessed 28 September 2026.
- Kongsamut, P., Rebelo, S., and Xie, D. (2001). Beyond balanced growth. *The Review of Economic Studies*, 68(4):869–882.
- Kopecky, J. (2026). The comparative advantage of age. TEP Working Paper 0726, Department of Economics, Trinity College Dublin.
- Lagakos, D., Moll, B., Porzio, T., Qian, N., and Schoellman, T. (2018a). Life-cycle human capital accumulation across countries: Lessons from US immigrants. *Journal of Human Capital*, 12(2):305–342.
- Lagakos, D., Moll, B., Porzio, T., Qian, N., and Schoellman, T. (2018b). Life cycle wage growth across countries. *Journal of Political Economy*, 126(2):797–849.
- Lagakos, D. and Waugh, M. E. (2013). Selection, agriculture, and cross-country productivity differences. *American Economic Review*, 103(2):948–980.

- Lee, R., Mason, A., and members of the NTA Network (2014). Is low fertility really a problem? Population aging, dependency, and consumption. *Science*, 346(6206):229–234.
- Li, H., Zhang, J., and Zhu, Y. (2008). The quantity-quality trade-off of children in a developing country: Identification using Chinese twins. *Demography*, 45(1):223–243.
- Lindenlaub, I. (2017). Sorting multidimensional types: Theory and application. *The Review of Economic Studies*, 84(2):718–789.
- Lise, J. and Postel-Vinay, F. (2020). Multidimensional skills, sorting, and human capital accumulation. *American Economic Review*, 110(8):2328–2376.
- Maestas, N., Mullen, K. J., and Powell, D. (2023). The effect of population aging on economic growth, the labor force, and productivity. *American Economic Journal: Macroeconomics*, 15(2):306–332.
- Matsuyama, K. (2009). Structural change in an interdependent world: A global view of manufacturing decline. *Journal of the European Economic Association*, 7(2–3):478–486.
- Medeiros, H. B. O., de Araújo, D. S. M. S., and de Araújo, C. G. S. (2013). Age-related mobility loss is joint-specific: an analysis from 6,000 Flexitest results. *AGE*, 35:2399–2407. PMID 23529505.
- Minnesota Population Center (2025). Ipums international: Integrated public use microdata series, international. Version 7.7 [dataset]. Minneapolis, MN: IPUMS.
- National Center for O*NET Development (2024). O*NET database. Release 29.0 [dataset]. U.S. Department of Labor, Employment and Training Administration. (Update year/release if different in your extraction.).
- Ngai, L. R. and Petrongolo, B. (2017). Gender gaps and the rise of the service economy. *American Economic Journal: Macroeconomics*, 9(4):1–44.
- Ngai, L. R. and Pissarides, C. A. (2007). Structural change in a multisector model of growth. *American Economic Review*, 97(1):429–443.
- Ohnsorge, F. and Trefler, D. (2007). Sorting it out: International trade with heterogeneous workers. *Journal of Political Economy*, 115(5):868–892.
- Peters, M. and Walsh, C. (2026). Population growth and firm dynamics. *Journal of Political Economy Macroeconomics*, 4(2):256–296.
- Pollak, R. A. and Wales, T. J. (1981). Demographic variables in demand analysis. *Econometrica*, 49(6):1533–1551.

- Porzio, T., Rossi, F., and Santangelo, G. (2022). The human side of structural transformation. *American Economic Review*, 112(8):2774–2814.
- Rodrik, D. (2016a). Premature deindustrialization. *Journal of Economic Growth*, 21(1):1–33.
- Rodrik, D. (2016b). Premature deindustrialization. *Journal of Economic Growth*, 21(1):1–33.
- Rosenzweig, M. R. and Zhang, J. (2009). Do population control policies induce more human capital investment? Twins, birth weight and China’s “one-child” policy. *Review of Economic Studies*, 76(3):1149–1174.
- Salthouse, T. A. (2009). When does age-related cognitive decline begin? *Neurobiology of Aging*, 30(4):507–514.
- Salthouse, T. A. (2014). Quantity and structure of word knowledge across adulthood. *Intelligence*, 46:122–130.
- Shorrocks, A. F. (2013). Decomposition procedures for distributional analysis: A unified framework based on the Shapley value. *Journal of Economic Inequality*, 11(1):99–126.
- Sposi, M., Yi, K.-M., and Zhang, J. (2026). The role of trade in premature deindustrialization and industry polarization. *Econometrica*. Forthcoming.
- Strauss, J., Witoelar, F., and Sikoki, B. (2016). The fifth wave of the indonesia family life survey: Overview and field report. volume 1. Working Paper WR-1143/1-NIA/NICHD, RAND Corporation.
- Tallis, G. M. (1961). The moment generating function of the truncated multi-normal distribution. *Journal of the Royal Statistical Society: Series B (Methodological)*, 23(1):223–229.
- Teignier, M. (2018). The role of trade in structural transformation. *Journal of Development Economics*, 130:45–65.
- United Nations, Department of Economic and Social Affairs, Population Division (2022). World population prospects. Revision year can be updated to match the vintage used in the paper.
- Verhaeghen, P. (2003). Aging and vocabulary scores: A meta-analysis. *Psychology and Aging*, 18(2):332–339.
- World Bank (2025). World development indicators. Washington, DC. (Use access date / download vintage in the final paper.).

Appendix

A Data sources, samples, and harmonization

This appendix documents the sources, samples, and harmonization procedures underlying the empirical analysis and the data-based inputs to the quantitative model.

A.1 Cross-country census data

IPUMS International. The worker-sorting analysis uses harmonized IPUMS International census microdata (Minnesota Population Center, 2025). The main age-profile estimation sample covers 88 countries from 1960 to 2020. The original unit of observation is an individual census record. For the main age profiles, I use IPUMS person weights to calculate sectoral employment shares and demographic composition for each single year of age within each country and census. The regressions cover ages 20–80 and weight these observations by total employment with a valid industry classification. The alternative age–period–cohort specifications retain more detailed groups defined by sex, nativity, schooling, and rural residence. The census data are repeated cross-sections: they allow five-year birth-cohort groups to be observed at different ages across successive censuses, but do not follow the same individuals over time.

Sector definitions. Industries are harmonized using the IPUMS INDGEN classification. Table 10 reports the concordance. For the main age-profile exercise, the manufacturing aggregate combines manufacturing, utilities, and construction and excludes mining. Exercises explicitly described as narrow manufacturing retain manufacturing alone. The denominator in the main profiles is total non-agricultural employment with a valid industry classification.

Table 10: IPUMS sector concordance

Sector	INDGEN codes	Activities included
Agriculture	10	Agriculture, forestry, and fishing
Manufacturing aggregate	30, 40, 50	Manufacturing; electricity, gas, water, and waste management; construction
Narrow manufacturing	30	Manufacturing only
Low-skill services	60, 70, 80, 110, 114, 120	Trade; hotels and restaurants; transportation, storage, and communications; unspecified, other, and private-household services
High-skill services	90, 100, 111, 112, 113	Finance and insurance; public administration and defense; business services and real estate; education; health and social work

Notes: Mining (INDGEN=20) and residual other industry (INDGEN=130) are included in non-agricultural employment but are not assigned to the three displayed non-agricultural sectors. Broad-industry exercises combine mining, manufacturing, utilities, construction, and residual other industry.

The age-profile controls are constructed among all workers with a valid industry classification.

They are the male share among workers with known sex; the urban share among workers with reported urban status; the share recorded as having moved from another first-level region or country among workers with reported prior-residence information; indicators for cells with no reported urban or prior-residence information; the primary-, secondary-, and university-educated shares among workers with known educational attainment, with less than primary omitted; and the share with missing or not-in-universe educational attainment.

Occupations and ability requirements. The occupation-based exercises use census samples reporting three-digit occupations under either the 1988 or 2008 revision of the International Standard Classification of Occupations (ISCO-88 and ISCO-08). Because a census may report either classification, both, or neither, the two occupation samples are analyzed separately. The occupation-gradient sample contains 112 ISCO-08 occupations observed across 21 country-years and 102 ISCO-88 occupations observed across 54 country-years; the samples overlap in one country-year.

To examine whether workers move toward less physically demanding occupations within manufacturing as they age, I also follow birth-cohort groups across successive censuses reporting ISCO-88 occupations. This sample covers 76 occupations in 13 countries and 15 census pairs. Occupation shares are measured relative to all employment in manufacturing, utilities, and construction within the cohort. Appendix D.3 describes the cohort matching and estimation.

O*NET Release 30.2 provides occupation-level ratings from its *Abilities* and *Generalized Work Activities* modules (National Center for O*NET Development, 2024). I combine these into three requirement measures: physical, such as strength and manual dexterity; fluid-cognitive, such as reasoning and problem-solving; and crystallized-and-interpersonal, such as communication and applying acquired knowledge. The measures range from zero to one and sum to one within each occupation, describing its relative requirements. They are matched separately to three-digit ISCO-08 and ISCO-88 occupations for the empirical exercises and averaged within sector assignments to construct the model’s fixed ability weights.

Age-ability profiles. I construct the age-ability profiles from published studies of physical performance and cognitive test scores. The physical profile draws on evidence on flexibility and grip strength (Medeiros, de Araújo, and de Araújo, 2013; Dodds, Syddall, Cooper et al., 2014), while the fluid-cognitive profile draws on processing speed and reasoning (Salthouse, 2009). For the crystallized-and-interpersonal bundle, I use vocabulary as a measure of accumulated knowledge (Verhaeghen, 2003; Hartshorne and Germine, 2015; Salthouse, 2014). The same profiles enter the empirical analysis and the quantitative model. Appendix E explains how these findings are used to construct the annual profiles and convert them into productivity indices.

Employment within manufacturing industries. ILOSTAT provides employment by age and two-digit manufacturing industry, together with separate tabulations of employment by occupation and industry (International Labour Organization, 2026). The age tabulations measure the share of manufacturing workers aged 15–24 relative to workers aged 15 and above. The occupation tabulations combine with O*NET requirements to measure each industry’s physical demands. This permits a comparison of younger-worker employment across manufacturing industries with different requirements. I retain manufacturing divisions under ISIC Revision 3 and Revision 4, treating the classifications separately and excluding utilities and construction. The baseline estimation sample covers 63 countries from 1993 to 2025. Appendix D.4 reports the sample requirements and estimation.

Census inputs used in the quantitative analysis. Indonesian IPUMS censuses provide complementary sector-by-age evidence and supplement the schooling distribution for the oldest cohorts. The full four-sector classification is available for 1976, 1980, 1985, 1990, 1995, 2005, and 2010.

The historical U.S. demographic path combines IPUMS USA Full Count census schedules with the Human Mortality Database. The census component uses the 1850, 1860, 1870, 1880, and 1900–1950 full counts. The 1850 and 1860 enslaved-person schedules are added when constructing both the population distribution and the agricultural-employment trigger, with enslaved workers assigned to agriculture. The trigger classifies industry using the harmonized IND1950 codes. Because the 1890 census schedules are unavailable, the series between 1880 and 1900 is interpolated. Age heaping in the early censuses is smoothed within five-year age groups while preserving each group’s total population. The Human Mortality Database supplies the annual population-by-age series from 1933 onward.

A.2 Panel data: Indonesia Family Life Survey

The Indonesia Family Life Survey (IFLS) is a longitudinal household survey conducted by RAND and SurveyMETER (Strauss, Witoelar, and Sikoki, 2016). Its five survey waves were fielded in 1993, 1997, 2000, 2007, and 2014. The first wave interviewed approximately 7,200 households and 33,000 individuals in 13 provinces representing 83 percent of Indonesia’s population; subsequent waves recontacted more than 90 percent of baseline-dynasty households.

At the individual level, the IFLS records employment, sector, occupation, earnings, schooling, and demographic characteristics. Stable individual identifiers allow workers to be followed across waves. Sector assignments for 1988 and 1993 are recovered from employment histories collected in 1997; assignments from 1997 onward use current employment records. Earnings for 1993 come from the contemporaneous 1993 interview. The main transition and age-earnings exercises use observations from 1993 to 2014, while the employee earnings-change exercise uses the 1997–2000, 2000–2007, and 2007–2014 wave pairs. The labor, wage, and schooling records provide the worker-allocation, wage, schooling-return, and cohort-schooling moments used to discipline the model’s worker block.

For the transition breakdown, high-skill services comprise finance, insurance, real estate and business services, together with selected professional, educational, health, government and defense occupations in community, social and personal services. Low-skill services comprise trade, transport and the remaining workers in that mixed industry. Appendix C.2 gives the corresponding industry and occupation codes.

Earnings and prices. Hourly earnings are calculated by dividing monthly earnings by weekly hours times 52/12. The age-earnings profiles use salary including benefits where available and business earnings otherwise. Business earnings are measured by net profits in later waves and by the combined profit/gross-income measure available in 1993. Monetary amounts reported in thousands of rupiah in 1993 are first converted to rupiah. All earnings are expressed in constant 2010 rupiah using Indonesia’s consumer price index. The age-earnings profiles use prices for the survey-wave year; the employee earnings-change exercise uses prices for the recorded interview year, with the wave year used when that date is unavailable.

The age-earnings profiles follow full-time workers aged 20–64 with usable earnings in the same sector in at least two waves. The employee earnings-change exercise instead follows employees observed at consecutive waves, using salary earnings at both interviews. The sample restrictions and estimation procedures are described in Appendix C.4 and Appendix C.5.

At the household level, the expenditure modules provide the Indonesian evidence on durable demand. For the descriptive profiles, I use all five waves, group households by the mean age of members aged 20 and above, and calculate survey-weighted expenditure shares within each wave. Expenditures are mapped into durables, manufactured non-durables, agriculture, services, and other expenditure using a common item-level concordance. The internally consistent 1993 expenditure categories provide the demand moments used in the model’s demand block, including the expenditure profile used to initialize durable stocks; the 1993 household asset module (Book 2, Section HR) records, for the same households, ownership and current values of vehicles and household appliances, which identify the durable income-effect parameter ι_D and the translation parameter $\bar{d}$.

A.3 Cross-country data

Sectoral employment, demography, and income. The principal aggregate employment outcomes come from the World Bank’s *World Development Indicators* (WDI) (World Bank, 2025). WDI reports agriculture, industry, and services as exhaustive employment categories; industry combines manufacturing, mining, utilities, and construction. The main within-country analysis uses 3,650 observations from 147 countries over 1991–2016. The young-adult variable in these regressions is the share of the total population aged 25–34, constructed from the 2024 revision of the United Nations’ *World Population Prospects* (WPP) (United Nations, Department of Economic and Social Affairs,

Population Division, 2022).

The Groningen Growth and Development Centre’s Economic Transformation Database and 10-Sector Database provide measures that isolate manufacturing and also supply Indonesian sectoral employment and real value added for the quantitative model (de Vries et al., 2021, 2015). These sectoral accounts anchor the model’s productivity calibration; the productivity paths themselves are recovered within the model.

The birth-rate instrument is the mean crude birth rate over years $t - 34$ through $t - 25$, using the historical series assembled by Delventhal, Fernández-Villaverde, and Guner (2024). WPP population stocks and survival probabilities are used to construct the surviving-child-stock instrument. GDP per capita comes from the Maddison Project Database in the long-run and census exercises (Bolt and van Zanden, 2024). Unless stated otherwise, income groups are defined using 2010 GDP per capita.

Macroeconomic controls. Female labor-force participation and imports relative to GDP come from WDI. Mean years of schooling among people aged 25 and above comes primarily from the United Nations Development Programme’s Human Development Report series, with the Wittgenstein Centre series used where the former is unavailable. The capital–output ratio, investment share, and total factor productivity are constructed primarily from Penn World Table national accounts (Feenstra, Inklaar, and Timmer, 2015). For countries outside its coverage, corresponding WDI series are used to extend the measures. These extensions preserve the 147-country estimation sample used in the main specification.

Eurostat household expenditure. Complementary evidence on durable demand comes from Eurostat Household Budget Survey tabulations for 34 European countries observed in an unbalanced set of seven rounds between 1988 and 2020. The source reports expenditure by the Classification of Individual Consumption According to Purpose (COICOP) and by the age of the household reference person. I retain the four reported age groups—under 30, 30–44, 45–59, and 60 or above—and convert expenditure reported per thousand into budget shares. The manufactured-durable aggregate combines the durable and semi-durable categories assigned to manufacturing in Table 11.

Trade, sectoral accounts, and world prices. Indonesian merchandise trade from 1995 onward is constructed from CEPII’s BACI product-level trade database; the preceding series uses the NBER–United Nations World Trade Flows and is joined to BACI at their common boundary. OECD Inter-Country Input–Output tables provide the utilities-and-construction components and sector-accounting denominators needed to construct the model’s manufacturing aggregate, supplemented by United Nations Industrial Development Organization manufacturing accounts and WDI national-accounts series.

Services trade uses World Trade Organization International Trade Statistics before 1995 and the

Table 11: COICOP categories classified as durables

Durable goods	Semi-durable goods treated as durable
Motor cars; motorcycles; bicycles; furniture and furnishings; household appliances; major tools and equipment; audio-visual, photographic, and information-processing equipment; telephone equipment; musical instruments and major recreational durables; therapeutic appliances and equipment; jewelry, clocks, and watches; electrical appliances for personal care.	Clothing and clothing materials; garments and clothing accessories; footwear; household textiles; carpets and other floor coverings; books; games, toys, and hobbies.

Notes: The empirical aggregate combines durable and semi-durable COICOP categories. Every included category is assigned to manufacturing in the paper’s concordance.

OECD–WTO Balanced Trade in Services database thereafter. Manufacturing’s world price is normalized, with 1993 as the base year. Agricultural prices use the World Bank agricultural commodity-price index relative to the GGDC manufacturing deflator, while service prices use GGDC sectoral deflators relative to manufacturing. These data are used to construct the model’s fixed trade wedges, export-demand shifters, external-balance terms, and world-price paths.

B Prediction 1: Manufacturing age profiles and cohort reallocation

This appendix develops the evidence for Prediction 1. Manufacturing’s employment share is higher among younger workers under several treatments of age, cohort, and period effects. The young–old difference also appears under alternative sector definitions and in most countries. Separate estimates show the same young-worker concentration among men, women, urban workers and rural workers. The final exercises examine manufacturing growth and hiring.

B.1 Age, cohort, and period specifications

The first set of exercises asks whether manufacturing’s concentration among younger workers persists under different ways of separating age, cohort and period variation. The five specifications below group census workers by country, census year, birth year, sex, foreign-born status, schooling and rural residence. Each resulting group is one observation. Regressions include the corresponding demographic controls, weight observations by the cell’s total employment, and cluster standard errors by country. Estimation covers ages 20–80; the figures display ages 20–70. Manufacturing includes manufacturing, utilities, and construction. The outcome is each sector’s share of nonagricultural employment, with total-employment shares also reported for the first two specifications.

Since calendar year equals age plus birth year, separating all three components requires restrictions. Each specification states its treatment of cohort and period variation explicitly. Table 12 summarizes the manufacturing difference between ages 20 and 64. It ranges from 10.3 to 19.1 percentage points

in the full sample, and the difference remains substantial in the non-high-income and middle-income samples.

Table 12: Manufacturing employment-share differences across age specifications

	All countries	Non-high-income	Middle-income
Country-cohort effects and GDP controls	18.16	18.95	20.43
Country-cohort effects and country trends	19.10	20.82	19.61
Quintic age polynomial	10.29	9.75	11.73
Five-year age indicators	10.96	12.12	12.88
Country-year effects	19.07	21.92	23.28

Notes: Entries are the manufacturing share of nonagricultural employment at age 20 minus the share at age 64, in percentage points, implied by each estimated age profile. For the age-indicator specification, the comparison is between ages 20–24 and 60–64. All specifications use the census-cell construction described in the text. The non-high-income sample excludes Austria, Canada, Finland, France, Germany, Greece, Ireland, Israel, Italy, Portugal, Spain, Switzerland, and the United States. The middle-income sample has 2010 GDP per capita above 0.10 and at most 0.45 times the United States level in the Maddison data. The full-sample GDP specification has 401,866 cells in 75 countries; the country-trend, country-year, polynomial, and age-group specifications have 409,019 cells in 78 countries. The figures add the estimation-sample weighted outcome mean to the estimated age component, preserving all age contrasts. Their vertical levels use this additive normalization.

B.1.1 Country-cohort effects and GDP controls

For demographic cell n in country c and census year t , the first specification is

$$s_n^j = \sum_{a \neq 50} \alpha_a^j \mathbf{1}\{a_n = a\} + X_n' \delta^j + \gamma_1^j \log y_{ct} + \gamma_2^j (\log y_{ct})^2 + \mu_{c,b(n)}^j + \varepsilon_n^j, \quad (40)$$

where s_n^j is sector j 's employment share, X_n contains the demographic controls, y_{ct} is real GDP per capita, and $b(n)$ is the five-year birth-cohort group. Country-cohort effects absorb persistent differences across cohorts within each country, while the GDP quadratic controls for period variation associated with the level of development. The specification assumes that remaining period changes are not systematically related to age after accounting for the included controls. The manufacturing difference between ages 20 and 64 is 18.2 percentage points in the full sample and 20.4 points in middle-income countries.

B.1.2 Country-cohort effects and country trends

The second specification replaces the GDP quadratic with a separate linear calendar-time trend for each country:

$$s_n^j = \sum_{a \neq 50} \alpha_a^j \mathbf{1}\{a_n = a\} + X_n' \delta^j + \mu_{c,b(n)}^j + \tau_c^j t + \varepsilon_n^j. \quad (41)$$

This permits secular period changes to differ across countries, while maintaining five-year cohort effects and single-year age indicators. It specifies period effects as linear within each country. Manufacturing's young-old difference is 19.1 percentage points in the full sample, 20.8 in non-high-income countries, and 19.6 in middle-income countries.




Figure 10: Country-cohort effects and GDP controls: All countries

Notes: Solid lines use nonagricultural employment as the denominator; dashed lines use total employment. Each line shows the estimated difference relative to age 50, added to the sector's employment-weighted average share in the estimation sample. Manufacturing includes utilities and construction. Ages 20–70 are shown; estimation covers ages 20–80.




Figure 11: Country-cohort effects and GDP controls: Non-high-income countries

Notes: Solid lines use nonagricultural employment as the denominator; dashed lines use total employment. Each line shows the estimated difference relative to age 50, added to the sector's employment-weighted average share in the estimation sample. Manufacturing includes utilities and construction. Ages 20–70 are shown; estimation covers ages 20–80.




Figure 12: Country-cohort effects and GDP controls: Middle-income countries

Notes: Solid lines use nonagricultural employment as the denominator; dashed lines use total employment. Each line shows the estimated difference relative to age 50, added to the sector's employment-weighted average share in the estimation sample. Manufacturing includes utilities and construction. Ages 20–70 are shown; estimation covers ages 20–80.




Figure 13: Country-cohort effects and country trends: All countries

Notes: Solid lines use nonagricultural employment as the denominator; dashed lines use total employment. Each line shows the estimated difference relative to age 50, added to the sector's employment-weighted average share in the estimation sample. Manufacturing includes utilities and construction. Ages 20–70 are shown; estimation covers ages 20–80.




Figure 14: Country-cohort effects and country trends: Non-high-income countries

Notes: Solid lines use nonagricultural employment as the denominator; dashed lines use total employment. Each line shows the estimated difference relative to age 50, added to the sector's employment-weighted average share in the estimation sample. Manufacturing includes utilities and construction. Ages 20–70 are shown; estimation covers ages 20–80.




Figure 15: Country-cohort effects and country trends: Middle-income countries

Notes: Solid lines use nonagricultural employment as the denominator; dashed lines use total employment. Each line shows the estimated difference relative to age 50, added to the sector's employment-weighted average share in the estimation sample. Manufacturing includes utilities and construction. Ages 20–70 are shown; estimation covers ages 20–80.

B.1.3 A quintic age polynomial

Following the age-profile approach of Porzio, Rossi, and Santangelo (2022), I allow common calendar-year effects and restrict the age function to a quintic polynomial centered at $\bar{a} = 23$, with no linear term:

$$s_n^j = \sum_{k=2}^5 \beta_k^j \left(\frac{a_n - \bar{a}}{\bar{a}} \right)^k + X_n' \delta^j + \mu_{c,b(n)}^j + \psi_t^j + \varepsilon_n^j. \quad (42)$$

The age profile is smooth and is constrained to have a zero slope at age 23. Setting its level to zero at that age fixes the reference point. Country-cohort effects capture persistent cohort differences and common year effects capture shared period changes. The estimated manufacturing share falls between ages 20 and 64 by 10.3 percentage points in the full sample, 9.8 in non-high-income countries, and 11.7 in middle-income countries.




Figure 16: Quintic age polynomial: All countries

Notes: Shares of nonagricultural employment. Each line shows the estimated difference relative to age 50, added to the sector's employment-weighted average share in the estimation sample. Manufacturing includes utilities and construction. Ages 20–70 are shown; estimation covers ages 20–80.

B.1.4 Five-year age indicators

I also replace the polynomial by five-year age indicators:

$$s_n^j = \sum_{g \neq 50} \alpha_g^j \mathbf{1}\{5 \lfloor a_n/5 \rfloor = g\} + X_n' \delta^j + \mu_{c,b(n)}^j + \psi_t^j + \varepsilon_n^j. \quad (43)$$

The omitted age group is 50–54. Age effects are constant within five-year age groups, cohort effects are constant within five-year birth-cohort groups, and period effects are common across countries. These grouping restrictions permit estimation of the age steps alongside cohort and year effects; the steps allow the profile's shape to vary across age groups without a polynomial restriction. The




Figure 17: Quintic age polynomial: Non-high-income countries

Notes: Shares of nonagricultural employment. Each line shows the estimated difference relative to age 50, added to the sector's employment-weighted average share in the estimation sample. Manufacturing includes utilities and construction. Ages 20–70 are shown; estimation covers ages 20–80.




Figure 18: Quintic age polynomial: Middle-income countries

Notes: Shares of nonagricultural employment. Each line shows the estimated difference relative to age 50, added to the sector's employment-weighted average share in the estimation sample. Manufacturing includes utilities and construction. Ages 20–70 are shown; estimation covers ages 20–80.

manufacturing difference between ages 20–24 and 60–64 is 11.0 percentage points in the full sample and 12.9 in middle-income countries.




Figure 19: Five-year age indicators: All countries

Notes: Shares of nonagricultural employment. Each line shows the estimated difference relative to ages 50–54, added to the sector’s employment-weighted average share in the estimation sample. Manufacturing includes utilities and construction. Ages 20–70 are shown; estimation covers ages 20–80.




Figure 20: Five-year age indicators: Non-high-income countries

Notes: Shares of nonagricultural employment. Each line shows the estimated difference relative to ages 50–54, added to the sector’s employment-weighted average share in the estimation sample. Manufacturing includes utilities and construction. Ages 20–70 are shown; estimation covers ages 20–80.

B.1.5 Country–year effects

The final specification compares ages within the same country and census year:

$$s_n^j = \sum_{a \neq 50} \alpha_a^j \mathbf{1}\{a_n = a\} + X_n' \delta^j + \psi_{ct}^j + \varepsilon_n^j. \quad (44)$$




Figure 21: Five-year age indicators: Middle-income countries

Notes: Shares of nonagricultural employment. Each line shows the estimated difference relative to ages 50–54, added to the sector’s employment-weighted average share in the estimation sample. Manufacturing includes utilities and construction. Ages 20–70 are shown; estimation covers ages 20–80.

Country–year effects absorb all common sectoral conditions in each census. This specification compares workers of different ages within the same country and census, holding the demographic controls fixed; it does not include birth-cohort effects. Manufacturing’s young–old difference is 19.1 percentage points in the full sample and 23.3 in middle-income countries. Together with the cohort-controlled estimates, these profiles establish a substantial young-worker concentration under the five treatments of age and time variation.

B.2 Sector definitions, employment denominators, and workforce groups

The next exercise asks whether the young–old difference depends on which industries are included in manufacturing or on whether employment in agriculture enters the denominator. Table 13 applies equation (1) to the main census cells, changing the manufacturing boundary or employment denominator. The manufacturing difference between ages 20 and 64 is 20.9 percentage points with the main aggregate and 18.7 points when manufacturing excludes utilities and construction. Relative to total employment, including agriculture, the difference is 12.1 points. The 20.9-point difference in the main specification corresponds to about 21 more manufacturing workers per 100 nonagricultural workers at age 20 than at age 64, after accounting for the regression controls. Figure 25 presents the corresponding age profiles.

Figure 26 shows the age profiles for agriculture, manufacturing, and both service sectors as shares of total employment, using equation (1). Including agriculture in the denominator leaves a pronounced manufacturing age gradient: the recentered manufacturing share falls from 23.3 percent at age 20 to




Figure 22: Country-year effects: All countries

Notes: Shares of nonagricultural employment. Each line shows the estimated difference relative to age 50, added to the sector's employment-weighted average share in the estimation sample. Manufacturing includes utilities and construction. Ages 20–70 are shown; estimation covers ages 20–80.




Figure 23: Country-year effects: Non-high-income countries

Notes: Shares of nonagricultural employment. Each line shows the estimated difference relative to age 50, added to the sector's employment-weighted average share in the estimation sample. Manufacturing includes utilities and construction. Ages 20–70 are shown; estimation covers ages 20–80.

11.2 percent at age 64, a decline of 12.1 percentage points (s.e. 2.1).

Separate estimates for men, women, urban workers and rural workers assess whether manufacturing's concentration among younger workers is shared across these groups. Table 14 applies equation (1) separately to each group and compares manufacturing's share of nonagricultural employment at ages 20 and 64. The difference is positive in all four groups, ranging from 14.9 to 20.4 percentage points.




Figure 24: Country-year effects: Middle-income countries

Notes: Shares of nonagricultural employment. Each line shows the estimated difference relative to age 50, added to the sector's employment-weighted average share in the estimation sample. Manufacturing includes utilities and construction. Ages 20–70 are shown; estimation covers ages 20–80.

Table 13: Manufacturing's young-old difference by sector definition and denominator

Sector / employment denominator	Age 20 minus 64 (pp)	Cells	Countries
Manufacturing / nonagricultural	20.91 (2.67)	16,965	88
Narrow manufacturing / nonagricultural	18.67 (1.93)	16,965	88
Manufacturing / total	12.09 (2.08)	16,978	88

Notes: Country-year-single-age census cells; equation (1), with country-five-year-cohort effects, single-year age indicators, a quadratic in log GDP per capita, sex, urban residence, migration, and education-composition controls, and indicators for missing composition controls. Estimation uses ages 20–80 and total-employment weights. Standard errors in parentheses are clustered by country. Manufacturing includes utilities and construction; narrow manufacturing includes manufacturing alone. The total-employment denominator includes agriculture. This table uses the main census-cell sample; Table 12 uses the more disaggregated demographic cells described above.

The urban and rural estimates use a smaller sample of censuses and workers with reported residence status.

B.3 The geographic breadth of within-cohort changes

Repeated censuses also allow the manufacturing share of individual birth cohorts to be followed directly within each country. For five-year birth cohort b observed in consecutive censuses t_0 and t_1 , define

$$d_{cb,t_0t_1} = \frac{1000}{t_1 - t_0} \left(\frac{E_{cb,t_1}^M}{E_{cb,t_1}^{NA}} - \frac{E_{cb,t_0}^M}{E_{cb,t_0}^{NA}} \right), \quad \bar{d}_c = \frac{\sum_{b,t_0,t_1} \bar{E}_{cb,t_0t_1}^{NA} d_{cb,t_0t_1}}{\sum_{b,t_0,t_1} \bar{E}_{cb,t_0t_1}^{NA}}, \quad (45)$$

where $\bar{E}_{cb,t_0t_1}^{NA} = (E_{cb,t_0}^{NA} + E_{cb,t_1}^{NA})/2$. The factor 1000 expresses a share change in percentage points per decade of age. I retain census intervals of four to fifteen years and cohorts whose five birth years




Figure 25: Manufacturing age profiles under alternative sector definitions and denominators

Notes: Estimated age coefficients relative to age 50, in percentage points. Specifications and samples are those in Table 13.




Figure 26: Within-cohort sectoral age profiles: Shares of total employment

Notes: Profiles are recentered so that each line's employment-weighted mean equals the corresponding estimation-sample employment share. Estimation covers ages 20–80; ages 20–70 are displayed. Shaded areas show 95 percent confidence intervals for age effects relative to age 50, shifted by the same recentering constant. Standard errors are clustered by country. Manufacturing includes utilities and construction. Mining and residual industry remain in the employment denominator but are not plotted.


Table 14: Young-worker concentration across workforce groups

	Men	Women	Urban	Rural
Age 20 minus age 64 (pp)	20.4 (3.2)	20.2 (3.0)	14.9 (3.5)	16.5 (4.8)
Cells	16,958	16,876	11,475	11,457
Countries	88	88	73	73

Notes: Each column uses country–census–single-age observations constructed for the indicated group, with employment shares and composition controls recalculated within that group. The specification follows equation (1), omitting the controls that are constant within the group. Manufacturing includes utilities and construction; the denominator is nonagricultural employment. Men and women cover the 88-country main sample. Urban and rural estimates use workers with reported residence status and cover 73 countries: fourteen countries have no reported urban–rural status in the censuses retained for the main analysis, and Israel has insufficient observations within country–cohort groups after restricting to reported status. Estimation covers ages 20–80 and weights each observation by the group’s total employment with a valid industry classification. Entries compare single ages 20 and 64; standard errors in parentheses are clustered by country.

lie within ages 20–64 at both endpoints. The country statistic averages the observed cohort changes using midpoint nonagricultural employment as weights.

The manufacturing share declines in 59 of the 68 countries with usable census pairs. The cross-country median is -4.11 percentage points per decade of age. Figures 27 and 28 show the distribution and individual country estimates, respectively. Each country enters the distribution once. The widespread decline therefore appears directly in observed cohort shares as well as in the pooled regression profiles.

**Figure 27:** Within-cohort manufacturing-share changes across countries

Notes: Distribution of the 68 country statistics in equation (45). Manufacturing includes utilities and construction; the denominator is each cohort’s nonagricultural employment. Each cohort’s change between censuses is expressed per decade of age and averaged within its country, using the mean of its employment at the two censuses as the weight.

Countries receive equal weight in the histogram.




Figure 28: Country estimates of within-cohort manufacturing-share changes
Notes: Country statistics underlying Figure 27, in percentage points per decade of age.

B.4 Manufacturing growth and workforce age

B.4.1 Age profiles in expanding and contracting manufacturing

A sector's growth can affect its workforce age through hiring and retention (Autor and Dorn, 2009). To examine how the manufacturing age profile varies with growth, I classify each country–census year using the trailing ten-year slope of the manufacturing employment share in the GGDC series. Expanding observations have slopes above +0.5 percentage points per decade, and contracting observations have slopes below −0.5. I estimate equation (1) separately in the two groups, using the main census cells and total-employment weights.

Figure 29 shows that manufacturing's share declines over the early working life in both environments. Between ages 20 and 35, the estimated decline is 12.72 percentage points (s.e. 1.58) in expanding manufacturing and 9.11 points (2.33) in contracting manufacturing. Over ages 20–64, the

differences are 21.41 (3.93) and 25.78 (3.85) points, respectively. Thus, the young-worker tilt is present even where manufacturing's aggregate employment share is increasing.




Figure 29: Manufacturing age profiles in expanding and contracting sectors

Notes: Age coefficients relative to age 50, in percentage points of nonagricultural employment, with pointwise 95 percent confidence intervals using country-clustered standard errors and country-cluster degrees of freedom. The expanding sample contains 3,704 country-year-age cells in 32 countries; the contracting sample contains 3,829 cells in 29 countries. Both specifications include the controls and country-five-year-cohort effects of equation (1), estimate ages 20–80, and use total-employment weights. The figure displays ages 20–64. The regime threshold is ± 0.5 percentage points per decade in the external manufacturing-share series; observations between the thresholds or without coverage are excluded.

B.4.2 Employment growth and hiring

The hiring channel formalized by Han and Suen (2011) provides a further benchmark: expanding sectors recruit disproportionately from young entrants and therefore tend to have younger workforces. For outcome q_{ctj} —sector j 's mean workforce age or its share of workers aged 20–34—I estimate

$$q_{ctj} = \beta g_{ctj} + \lambda_{ct} + \varepsilon_{ctj}, \quad g_{ctj} = 10 \frac{\log E_{ctj} - \log E_{c,t^-,j}}{t - t^-}, \quad (46)$$

where E_{ctj} is sector employment and t^- is the preceding census with usable employment data. Regressions include country-year effects, weight observations by sector employment, and cluster standard errors by country.

Across agriculture, manufacturing, and services, employment growth is associated with a younger workforce: the mean-age coefficient is -3.83 years (s.e. 0.98) per unit of decadal log employment growth. Yet the observed growth ordering predicts the opposite of manufacturing's youth advantage over services. Services employment grows faster, by 0.122 log points per decade. Applying the three-sector coefficient therefore predicts services to be 0.47 years younger than manufacturing, whereas

they are 1.05 years older. The young-worker-share comparison gives the same result: manufacturing’s observed advantage is 3.4 percentage points, while the growth benchmark predicts a 1.5-point disadvantage. The hiring-growth benchmark therefore does not account for the observed age ordering of manufacturing and services.

Table 15: Employment growth and the manufacturing youth advantage

	Mean age (years)	Young-worker share (fraction)
<i>A. Coefficient on decadal log employment growth</i>		
Agriculture, manufacturing, and services	−3.83 (0.98)	0.122 (0.030)
Manufacturing and services	−0.74 (1.65)	0.035 (0.060)
	Observed difference	Growth prediction
<i>B. Workforce-age differences</i>		
Mean age: services minus manufacturing	+1.05 years	−0.47 years
Young share: manufacturing minus services	+3.4 pp	−1.5 pp

Notes: Ages 20–64; 217 country–census observations with preceding censuses in 77 countries. The three-sector regression has 651 observations and the manufacturing–services regression has 434. Mean age uses five-year age-band midpoints; the young share covers ages 20–34. Panel A reports equation (46), with country-clustered standard errors in parentheses. Panel B multiplies the three-sector coefficients by the observed services-minus-manufacturing growth difference of 0.122; the young-share row reverses the sector order and converts fractions to percentage points. Excluding Germany, which has a 1970–71 coverage break, gives a mean-age coefficient of −4.58 (0.65) and a manufacturing young-share advantage of 3.5 points. Winsorizing growth at the first and ninety-ninth percentiles gives a mean-age coefficient of −4.64 (0.61).

C Prediction 2: Individual sectoral transitions and earnings profiles

The Indonesian panel links the decline in manufacturing’s employment share within birth cohorts to individual workers’ sectoral transitions. The 1993 sector assignment uses employment histories collected in 1997, while 1993 earnings are observed in the 1993 interview. This appendix presents the construction of the predicted probabilities, their stability across specifications, the destinations within services, and complementary evidence on sectoral persistence and earnings.

C.1 Predicted transition probabilities

This exercise asks whether the age pattern of sectoral transitions remains when local changes in employment opportunities are allowed to differ across provinces. The sample comprises workers aged 20–64 at the origin interview who are observed employed outside agriculture at both interviews. The four consecutive interview pairs are 1993–1997, 1997–2000, 2000–2007 and 2007–2014. Manufacturing covers the manufacturing industry alone; services combine trade, transport, finance and other services. Other nonagricultural destinations remain in the estimation sample: the outcome for a manufacturing-to-services transition, for example, is one for a services destination and zero for manufacturing or other industry. Age is restricted at origin; the destination observation is retained when the worker is older than 64.

For each direction d , equation (2) includes origin-age indicators, sex, educational level, grade completed, years of schooling, and urban residence. The three education variables enter linearly; missing values are replaced by their medians in the corresponding origin-sector sample. The urban indicator is zero when unreported. Estimation uses positive origin-wave person weights and observations with a recorded origin province. Standard errors are clustered by worker. Province-by-wave-pair fixed effects allow transition probabilities to differ across origin provinces in each interview interval. Wave-pair fixed effects allow probabilities to differ across intervals, with a common effect across provinces.

Let $\mathcal{I}_d$ denote the origin-sector estimation sample, ω_i the origin person weight, and $\hat{T}_i^d(a)$ the fitted value when worker i is assigned age band a , retaining that worker's other characteristics and fixed effects. The predicted probability reported for age a is

$$\hat{P}_d(a) = \frac{\sum_{i \in \mathcal{I}_d} \omega_i \hat{T}_i^d(a)}{\sum_{i \in \mathcal{I}_d} \omega_i}. \quad (47)$$

Within each transition direction, predictions for every age use the same weighted distribution of worker characteristics, origin provinces and interview intervals. These predictions include the estimated fixed effects from the corresponding specification. Confidence intervals account for the joint estimation uncertainty in all terms entering the prediction, with errors clustered by worker.

Table 16 shows closely aligned age patterns under the two specifications. Manufacturing-to-services probabilities rise from the beginning of working life through middle age, while the reverse transition becomes progressively less frequent. Allowing transition probabilities to differ across provinces leaves these patterns intact. Evaluated over the same distribution of characteristics and province-interval observations, the predictions correspond to about 29 of every 100 manufacturing-origin workers moving to services at ages 20–24, compared with 39 at ages 45–54.

Table 16: Predicted transition probabilities across fixed-effect specifications

Age at origin	$M \rightarrow S$		$S \rightarrow M$	
	Interview pair	Province $\times$ pair	Interview pair	Province $\times$ pair
20–24	27.9 (1.6)	28.7 (1.6)	12.3 (0.9)	12.2 (0.9)
25–34	32.1 (1.1)	32.3 (1.1)	9.1 (0.4)	9.2 (0.4)
35–44	36.4 (1.4)	36.0 (1.4)	6.4 (0.4)	6.4 (0.4)
45–54	39.5 (2.0)	38.9 (2.1)	4.6 (0.5)	4.6 (0.5)
55–64	35.1 (3.5)	34.4 (3.4)	4.6 (0.9)	4.4 (0.9)
Observations	4,826	4,826	16,459	16,459
Workers	3,197	3,197	9,485	9,485

Notes: Probabilities are in percent; worker-clustered standard errors of the probability levels are in parentheses. The origin-sector samples and controls are identical across the two fixed-effect specifications. Predictions follow equation (47). Intervals span 3–7 years.

C.2 Transitions by service sector

Separating low- and high-skill services identifies which part of services receives the age-related movement out of manufacturing. The increase through middle age is concentrated in low-skill services. I distinguish low-skill services (L) from high-skill services (H) using the classification in the quantitative analysis. High-skill services comprise finance, insurance, real estate and business services, together with selected professional, educational, health, government and defense occupations in the mixed community, social and personal-services industry. Low-skill services comprise trade, transport and the remaining observations in that mixed industry.²¹

Table 17 applies the province-by-wave-pair specification separately to each origin sector and destination. The probability of moving from manufacturing to low-skill services rises from 26.0 percent at ages 20–24 to 36.5 percent at ages 45–54. Movement in the reverse direction falls from 13.1 percent to 5.7 percent. Transitions from manufacturing into high-skill services are much smaller, at approximately 2–4 percent across age groups.

Table 17: Predicted transitions between manufacturing and service sectors

Transition	Age at origin				
	20–24	25–34	35–44	45–54	55–64
$M \rightarrow L$	26.0 (1.6)	29.8 (1.1)	33.4 (1.4)	36.5 (2.1)	30.7 (3.4)
$M \rightarrow H$	2.7 (0.6)	2.5 (0.4)	2.6 (0.5)	2.4 (0.6)	3.7 (1.1)
$L \rightarrow M$	13.1 (1.0)	10.1 (0.5)	7.6 (0.5)	5.7 (0.6)	5.4 (0.9)
$L \rightarrow H$	4.6 (0.7)	4.8 (0.3)	5.5 (0.4)	6.5 (0.6)	5.5 (0.6)
$H \rightarrow M$	4.7 (1.6)	3.0 (0.6)	2.1 (0.5)	1.3 (0.6)	2.8 (2.5)
$H \rightarrow L$	31.5 (3.2)	21.8 (1.3)	18.3 (1.3)	16.8 (1.8)	36.9 (5.8)
M -origin observations	4,826 (3,197 workers)				
L -origin observations	13,496 (8,219 workers)				
H -origin observations	250	1,021	1,064	546	82
H -origin total	2,963 (1,653 workers)				

Notes: Entries are predicted probabilities in percent, with worker-clustered standard errors in parentheses. The sample, controls and prediction procedure follow Appendix C.1. The two manufacturing-origin destination probabilities add to the pooled-services probability before rounding. Observation counts refer to origin-sector interview pairs, including unchanged and other eligible destinations. H accounts for 11.9 percent of origin-weighted employment in the full nonagricultural transition sample of 23,359 interview-pair observations, including other industry.

High-skill services have a relatively small employment base in this panel, accounting for 11.9 percent of weighted nonagricultural origin employment. This provides context for the low entry probabilities into H . Its origin sample is also smaller, particularly at the youngest and oldest ages; the table reports the age-specific observation counts and the resulting sampling uncertainty.

²¹In the IFLS classification, H includes industry 8 and workers in industry 9 with occupation codes 6, 7, 13, 14, 20, 31, M1, M2, MM or SS. L includes industries 6 and 7 and the remaining industry-9 observations. Manufacturing is industry 3. The service split uses both industry and occupation information.

C.3 Sectoral persistence and local employment change

The next two exercises examine the probability of remaining in a sector. The first asks whether manufacturing persistence falls with age after controlling for local manufacturing employment changes. The second asks whether its age pattern differs from that of services within the same province and interview interval. In both, the reference age group is 25–34 and D_{ia} indicates the worker’s age band at origin.

First, among workers originating in manufacturing, I estimate

$$R_i^M = \sum_{a \neq a_{\text{ref}}} \beta_a^R D_{ia} + X_i' \gamma^R + \delta \Delta M_{p(i),q(i)}^{-i} + \lambda_{q(i)}^R + \varepsilon_i^R, \quad (48)$$

where R_i^M indicates persistence in manufacturing and $\Delta M_{p(i),q(i)}^{-i}$ is the change in manufacturing’s share of total employment in the worker’s origin province over the same wave pair, calculated excluding that worker. The age coefficients therefore compare manufacturing persistence after conditioning on both national changes between interviews and contemporaneous manufacturing reallocation within the origin province.

Second, I compare persistence in manufacturing with persistence in services:

$$R_i = \sum_{a \neq a_{\text{ref}}} \beta_a^{MS} (M_i D_{ia}) + X_i' \gamma^{MS} + \phi_{s(i),p(i),q(i)} + \psi_{a(i),p(i),q(i)} + \varepsilon_i^{MS}, \quad (49)$$

where R_i indicates persistence in the origin sector and M_i identifies a manufacturing origin. The sample contains workers whose origin and destination sectors are manufacturing or services. Origin-sector $\times$ province $\times$ wave-pair fixed effects absorb each sector’s average persistence in a given province and interval, while age-group $\times$ province $\times$ wave-pair fixed effects absorb changes common to workers of the same age in that province and interval. The coefficients thus measure how the manufacturing-minus-services persistence differential at each age differs from that among workers aged 25–34.

Table 18: Alternative specifications of sectoral persistence by age

Outcome	Age at origin, relative to ages 25–34				Observations
	20–24	35–44	45–54	55–64	
Persist in manufacturing	2.5 (2.0)	−5.0*** (1.7)	−7.7*** (2.3)	−8.4** (3.5)	5,455
Manufacturing–services persistence differential	5.8** (2.3)	−5.5*** (1.9)	−8.9*** (2.4)	−2.7 (3.6)	20,506

Notes: Entries are percentage-point coefficients relative to ages 25–34. The first specification follows equation (48); the second follows equation (49). Both include controls for sex, education, and urban residence, use the IFLS person weight from the origin wave, and cluster standard errors by worker. The first sample contains workers originating in manufacturing and permits any observed destination sector. The second restricts origins and destinations to manufacturing and services. The sample covers consecutive IFLS wave pairs from 1993 to 2014.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Both specifications reproduce the age ordering of the directional flows. Manufacturing persistence is 5.0, 7.7, and 8.4 percentage points lower at ages 35–44, 45–54, and 55–64 than at ages 25–34, even

after conditioning on local manufacturing reallocation. In the second specification, the gap between manufacturing and services in the probability of remaining in the sector is 5.8 percentage points larger at ages 20–24 than at ages 25–34. It is 5.5 and 8.9 points smaller at ages 35–44 and 45–54, respectively. The individual-career pattern therefore survives comparisons that absorb sector-specific and age-specific changes within provinces and survey intervals.

C.4 Within-worker earnings profiles

The earnings profiles ask whether the sector workers leave as they age also has a less favorable age-related earnings profile. The IFLS provides a direct view of these earnings changes within individual workers. By observing the same worker repeatedly in a sector, the panel separates these changes from fixed differences in earnings across workers. I estimate the following equation separately by sector j :

$$\log w_{it} = \alpha_i^j + \lambda_t^j + \sum_{g \neq 20-24} \beta_g^j \mathbf{1}[a_{it} \in g] + \varepsilon_{it}^j. \quad (50)$$

Here w_{it} is real hourly earnings, α_i^j is a worker fixed effect, λ_t^j is a sector-specific survey-wave effect, and g indexes five-year age bands. Each five-year age band has its own earnings coefficient, allowing the profile to rise or fall across age groups. Within each sector, a separate effect for each survey wave captures earnings movements shared across workers. The plotted coefficients show the remaining age-related earnings profile.

The sample covers full-time workers aged 20–64 observed with usable earnings in the sector in at least two of the 1993, 1997, 2000, 2007 and 2014 waves. Full-time work is at least 30 hours per week. Earnings comprise salary including benefits or, for the self-employed, business earnings; monthly amounts are converted to constant 2010 rupiah using Indonesia’s consumer price index and divided by weekly hours times 52/12. Positive earnings are trimmed at the unweighted first and ninety-ninth percentiles within wave among full-time workers, before applying the age restriction. Regressions use person weights and worker-clustered standard errors. The profiles concern repeated observations in the same sector; continuous sector tenure is not required.

Figure 30 shows a marked difference in the age-related earnings profiles. Manufacturing earnings are approximately flat early in working life and decline thereafter, reaching 0.55 of the age-20–24 index by ages 60–64. The index for all services rises to 1.38 at ages 45–49 and remains above its age-20–24 value through ages 55–59. For example, at ages 45–49 the estimated age component is 38 percent above its age-20–24 value in services and 9 percent below it in manufacturing. These within-worker patterns complement the sectoral transitions: the sector from which workers move as they age has the more steeply declining age-related earnings profile.




Figure 30: Within-worker earnings profiles by age, Indonesia

Notes: Each index is $\exp(\hat{\beta}_g^j)$ from equation (50), relative to ages 20–24. Manufacturing aggregates manufacturing, utilities and construction; the second line comprises trade, transport and community, social and personal services; all services additionally include finance, insurance, real estate and business services. Manufacturing uses 5,640 observations from 2,262 workers; trade, transport and other services 13,326 observations from 5,222 workers; pooled services 14,018 observations from 5,497 workers. The age profiles hold the survey-wave component fixed.

Table 19: Within-worker earnings profiles: coefficients and indices

Age band	Manufacturing aggregate		All services	
	Log coefficient	Index	Log coefficient	Index
25–29	0.012 (0.060)	1.01	0.148 (0.058)	1.16
30–34	0.000 (0.082)	1.00	0.217 (0.071)	1.24
35–39	−0.004 (0.113)	1.00	0.295 (0.094)	1.34
40–44	−0.022 (0.145)	0.98	0.321 (0.119)	1.38
45–49	−0.094 (0.180)	0.91	0.325 (0.143)	1.38
50–54	−0.204 (0.215)	0.82	0.261 (0.169)	1.30
55–59	−0.417 (0.248)	0.66	0.124 (0.195)	1.13
60–64	−0.601 (0.305)	0.55	−0.146 (0.240)	0.86

Notes: Coefficients are relative to ages 20–24. Worker-clustered standard errors are in parentheses. The sample and specification correspond to the two indicated series in Figure 30; each index is the exponential of the age coefficient.

C.5 Employee earnings around sectoral transitions

This exercise asks whether employees moving between manufacturing and services experience higher or lower earnings growth than employees observed in their original sector at the next interview. The comparison is estimated separately for manufacturing and service origins:

$$\frac{\log w_{i,t'} - \log w_{it}}{t' - t} = \delta_j \mathbf{1}[s_{i,t'} \neq j] + f_j(a_{it}) + X'_{it} \gamma_j + \lambda_{p(i,t),q(t,t')}^j + u_{i,t,t'}. \quad (51)$$

The sample includes government and private employees with positive salaries and valid hours, aged 20–64 at both endpoints, in the 1997–2000, 2000–2007 and 2007–2014 interview pairs. Both endpoints must be in manufacturing or services. Here manufacturing is industry 3, services are industries 6–9,

and hourly earnings use employee salary including benefits. The age function includes five-year origin-age indicators; X_{it} contains sex, years of schooling and rural residence at origin. The fixed effects are origin-province by interview pair. Regressions use origin person weights and worker-clustered standard errors.

Monthly salaries are divided by weekly hours times 52/12. Within each wave, hourly salaries are trimmed at the unweighted first and ninety-ninth percentiles among positive salary observations aged 20–64 with positive person weights. Earnings are then expressed in constant 2010 rupiah using annual consumer prices for the year of the first recorded interview for the employment questionnaire, with the wave year used when the interview year is unavailable. The denominator $t' - t$ is the three- or seven-year survey-wave interval. Missing schooling is replaced by the origin-sample median with a missing-value indicator; missing sex and rural residence are retained as separate categories. Origin province must be observed.

Table 20: Employee earnings growth and sectoral transitions, Indonesia

	$M \rightarrow S$ relative to $M \rightarrow M$	$S \rightarrow M$ relative to $S \rightarrow S$
Annual log hourly earnings-growth difference	−1.79	−0.22
95 percent confidence interval	[−3.64, 0.05]	[−1.72, 1.28]
Interview-pair observations	1,721	4,615
Workers	1,337	3,254
Sector changes	552	431

Notes: Entries are $100\hat{\delta}_j$ from equation (51): differences in annualized log real hourly salary growth, in log points per year. Confidence intervals use worker-clustered standard errors and the t distribution with the number of workers minus one degrees of freedom. Employees are government or private employees at both interviews; casual workers and self-employed workers are excluded. Hours must lie between 1 and 112 per week. The comparison is between observed movers and workers in the same sector at both endpoints.

Both estimated earnings-growth differences are negative and statistically indistinguishable from zero at the 5 percent level. The estimates therefore do not identify a statistically significant earnings-growth premium for employees changing sectors over the observed interview intervals.

D Prediction 3: Occupational requirements and employment profiles

This appendix examines how occupational requirements relate to employment at different ages. It relates occupations' employment profiles to their ability requirements, measures how much of the manufacturing age gap is accounted for by occupational composition, follows that composition as cohorts age, and compares young workers' employment shares across manufacturing divisions using ILOSTAT data.

D.1 Occupation age gradients and ability requirements

I begin by asking whether occupations requiring early-peaking abilities are more common among younger workers. The main text presents this relationship for five groups of occupations; here I quantify it across individual three-digit occupations.

Employment profiles by occupation. First, I estimate a separate age profile for each three-digit occupation o , using the specification in equation (1):

$$s_{octa} = \sum_{a' \neq 50} \alpha_{a'}^o \mathbf{1}[a = a'] + X'_{cta} \beta^o + \gamma_1^o \ln \text{GDP}_{ct} + \gamma_2^o (\ln \text{GDP}_{ct})^2 + \mu_{c,b(t,a)}^o + u_{octa}. \quad (52)$$

Here, s_{octa} is occupation o 's share of non-agricultural employment among workers of age a in country c and census t . Both the numerator and denominator retain only workers with valid occupation codes. The denominator includes all such non-agricultural workers, including those in occupations outside the set with matched ability requirements.

As in Prediction 1, $\mu_{c,b(t,a)}^o$ denotes country–five-year-birth-cohort fixed effects, where $b(t, a) = 5\lfloor(t - a)/5\rfloor$, and the GDP terms use income per capita. The vector X_{cta} contains the same sex, urban-residence, migration, and educational-attainment shares, together with the same missing-value indicators. These controls retain their values from the original Prediction 1 country–census–age cells; employment shares and regression weights are calculated using workers with valid occupation codes. The interpretation therefore follows the same age–period–cohort assumptions as Prediction 1.

Regressions cover single ages 20–80, with age 50 omitted, and the reported comparisons concern ages 20–64. Each observation is weighted by total employment among workers with valid occupation codes in its country–census–age cell, including agricultural employment. Within each occupational classification, every occupation is estimated on the same observations and with the same weights. An occupation's absence from an observed cell is recorded as a zero share and retained in estimation.

The ISCO-08 sample contains 112 occupations, 1,276 country–census–age cells, and 21 censuses from 17 countries. The ISCO-88 sample contains 102 occupations, 3,269 cells, and 54 censuses from 36 countries. The occupational sets are held fixed before the updated regressions are estimated.²²

Ability requirements and age-related employment changes. I then ask whether these estimated profiles vary systematically with occupations' ability requirements. Let $\theta_o^{PF} = \theta_o^P + \theta_o^F = 1 - \theta_o^C$

²²The sets retain occupations with matched requirements and sufficient coverage in the initial five-year-age-band construction: at least ten cells across three country–censuses, with survey-weighted employment of at least 50 in the occupation and 1,000 in the non-agricultural denominator. Armed-forces occupations and occupations in ISCO major group 6 are excluded from these sets, but not from the all-valid-coded non-agricultural denominator. These initial coverage restrictions determine the occupational sets; they are not imposed on every single-age cell in the updated regressions, which retain zero occupation shares.

be the combined physical and fluid-cognitive share of occupation o 's requirements. I summarize each occupation's age profile by

$$\widehat{\Delta}_o = 100(\widehat{\alpha}_{64}^o - \widehat{\alpha}_{20}^o),$$

the estimated change in its employment share between ages 20 and 64, measured in percentage points. A negative value means that the occupation is less common among older workers after the adjustments in equation (52). This comparison does not impose a constant decline at every age.

Across occupations, I estimate

$$\widehat{\Delta}_o = \kappa + \chi\theta_o^{PF} + v_o. \quad (53)$$

Each occupation contributes one observation: its estimated employment-share change from equation (52). Each receives equal weight. The mechanism predicts $\chi < 0$: occupations placing greater weight on early-peaking abilities should have more negative age-related changes in employment share.

The estimated relationship is negative in both classifications. Panel A of Table 21 reports $\widehat{\chi} = -3.284$ in ISCO-08 and -2.963 in ISCO-88. To express these magnitudes in terms of observed differences in requirements, consider occupations at the 75th and 25th percentiles of θ_o^{PF} . The fitted change in employment share between ages 20 and 64 is 1.21 percentage points lower for the higher-requirement occupation in ISCO-08 and 1.27 points lower in ISCO-88. These are differences in employment-share changes, not percentage changes in the ratio of the occupations' employment shares.

Uncertainty in the estimated relationship. Table 21 reports two standard-error calculations. Across-occupation heteroskedasticity-robust standard errors describe uncertainty in the second regression using occupations as observations. They allow residual variation to differ across occupations, but do not propagate the covariance between first-step estimates obtained from the same countries.

The country-clustered calculation instead accounts for that shared estimation uncertainty, holding the occupational set and requirement scores fixed. Because all occupations use the same observations, regressors, and weights, the second-step coefficient can also be obtained from a single regression. Define

$$q_o = \frac{\theta_o^{PF} - \bar{\theta}^{PF}}{\sum_{\ell} (\theta_{\ell}^{PF} - \bar{\theta}^{PF})^2}, \quad S_{cta} = 100 \sum_o q_o s_{octa},$$

where $\bar{\theta}^{PF}$ is the unweighted mean across occupations. Estimating equation (52) with S_{cta} as the outcome yields an age-64 minus age-20 contrast equal to $\widehat{\chi}$. Clustering this regression by country allows estimation errors to be correlated across ages and censuses within a country, including their effects on different occupations.

The ISCO-08 relationship is statistically distinguishable from zero under both calculations. In ISCO-88, it is distinguishable under country clustering but imprecise under the across-occupation calculation. The latter reflects substantial variation in employment changes among occupations with

similar requirement scores. The two calculations describe different sources of uncertainty and are therefore reported separately.

Occupation size and the age comparison. Absolute employment-share changes can be larger for occupations that account for more employment. I therefore also divide each occupation’s estimated change by its mean employment share:

$$\hat{R}_o = \frac{\hat{\alpha}_{64}^o - \hat{\alpha}_{20}^o}{\bar{s}_o}, \quad \bar{s}_o = \frac{\sum_{c,t,a} E_{cta} s_{octa}}{\sum_{c,t,a} E_{cta}},$$

where E_{cta} is the total-employment weight used in the first regression. I repeat equation (53) with $\hat{R}_o$ as the dependent variable. This measures the age-related change relative to an occupation’s fixed mean share; it is not the ratio of its employment shares at ages 64 and 20.

The relationship remains negative in both classifications after this normalization. ISCO-08 remains precise under both standard-error calculations, while ISCO-88 remains imprecise under the across-occupation calculation. For country-clustered inference, the normalized composite outcome is $\sum_o q_o s_{octa} / \bar{s}_o$; the mean-share normalizers are treated as fixed.

The findings also do not depend on comparing two single ages. Panel B replaces the age-64 minus age-20 contrast with the difference between the average estimated coefficients at ages 60–64 and 20–24. The coefficients remain negative in both classifications, with the same distinction in precision between the two standard-error calculations.

Table 21: Occupational employment profiles and physical and fluid-cognitive requirements

	Employment-share change (percentage points)		Change relative to mean employment share	
	ISCO-08	ISCO-88	ISCO-08	ISCO-88
<i>Panel A. Age 64 minus age 20</i>				
Requirement-share coefficient	−3.284	−2.963	−2.047	−1.739
Country-clustered standard error	(0.541)	(0.481)	(0.376)	(0.471)
Across-occupation robust standard error	(1.203)	(3.424)	(0.526)	(1.597)
R^2	0.124	0.041	0.090	0.026
<i>Panel B. Average ages 60–64 minus average ages 20–24</i>				
Requirement-share coefficient	−2.507	−2.539	−1.568	−1.690
Country-clustered standard error	(0.469)	(0.409)	(0.348)	(0.311)
Across-occupation robust standard error	(0.978)	(3.185)	(0.478)	(1.347)
Occupations	112	102	112	102
Countries	17	36	17	36

Notes: Each coefficient relates an occupation’s estimated age-related employment change to its requirement share θ_o^{PF} , giving occupations equal weight. All first-step regressions use equation (52). The last two columns divide changes by each occupation’s fixed employment-weighted mean share. Across-occupation standard errors use HC1; country-clustered standard errors use the equivalent composite-outcome regression described in the text. The reported R^2 values refer to the across-occupation regressions in Panel A.

Grouped profiles in the main text. Figure 2 uses the same occupational sets. Within each classification, I rank occupations by θ_o^{PF} and assign whole occupations to five groups using cumula-

tive pooled-employment thresholds of 20 percent. Group membership is fixed across ages. Because occupations remain intact, the groups need not contain exactly one-fifth of employment.

For group g , the outcome is $s_{cta}^g = \sum_{o \in g} s_{octa}$. I estimate equation (52) with this group share as the dependent variable and plot the age coefficients in percentage points relative to age 50. The denominator remains all valid-coded non-agricultural employment, rather than employment only in the selected occupations. The selected occupational sets cover 95.7 percent of this employment in ISCO-08 and 92.2 percent in ISCO-88.

The estimated change between ages 20 and 64 becomes progressively more negative across the five requirement groups in both classifications. The difference between the highest- and lowest-requirement groups' changes is -36.6 percentage points in ISCO-08 (s.e. 6.8) and -23.3 points in ISCO-88 (s.e. 4.7). Comparing average coefficients at ages 60–64 with those at 20–24 gives the same qualitative pattern. Thus, the grouped contrast is pronounced in both samples, while the occupation-level regressions also reveal variation among jobs with similar ability requirements.

D.2 Occupational composition and the manufacturing age gradient

This exercise examines how much of the manufacturing age gradient in Prediction 1 reflects differences in the occupations workers hold. I apply the specification underlying Figure 1 to the occupation-coded census samples and then add three-digit occupation fixed effects. Within each occupational classification, both regressions use exactly the same workers and weights. I retain censuses satisfying the Figure 1 sample restrictions and require a valid three-digit occupation code. The resulting samples contain 21 censuses from 17 countries for ISCO-08 and 54 censuses from 36 countries for ISCO-88.

Specification and weighting. The original regression uses manufacturing's share of non-agricultural employment in each country–census–single-age cell. To introduce occupation fixed effects, I divide these cells by three-digit occupation and manufacturing status. Each observation n therefore identifies a country, census year, single age, occupation, and manufacturing indicator M_n . As in Prediction 1, M_n equals one for manufacturing, utilities, and construction, and zero for other non-agricultural industries.

Let $g(n)$ denote observation n 's country–census–age cell and e_n its survey-weighted employment. Within that cell, let E_g denote total employment, including agriculture, and N_g non-agricultural employment, both calculated among workers with valid occupation codes. I assign observation n the regression weight

$$w_n = e_n \frac{E_{g(n)}}{N_{g(n)}}.$$

These weights sum to E_g within each country–census–age cell, and the weighted mean of M_n equals its manufacturing share of non-agricultural employment. Consequently, without occupation controls,

the expanded regression reproduces the age coefficients from the Figure 1 regression on the restricted sample.

Separately for ISCO-08 and ISCO-88, I estimate

$$M_n = \sum_{a' \neq 50} \alpha_{a'} \mathbf{1}[a(n) = a'] + Z'_{g(n)} \beta + \mu_{c(n), b(n)} + \varepsilon_n, \quad (54)$$

$$M_n = \sum_{a' \neq 50} \tilde{\alpha}_{a'} \mathbf{1}[a(n) = a'] + Z'_{g(n)} \tilde{\beta} + \tilde{\mu}_{c(n), b(n)} + \phi_{o(n)} + \tilde{\varepsilon}_n. \quad (55)$$

Here, $\mu_{c,b}$ denotes country–birth-cohort fixed effects, with five-year birth cohorts defined by $b(n) = 5\lfloor[t(n) - a(n)]/5\rfloor$. The vector Z_g contains the Figure 1 controls: male, urban, and migrant shares; indicators for missing urban and migration information; education shares and the share with missing education information; and log GDP per capita and its square. These controls retain their original country–census–age values; employment shares and regression weights are recalculated for workers with observed occupations.

The second regression adds three-digit occupation fixed effects $\phi_{o(n)}$, with all other coefficients re-estimated. Both specifications use single ages from 20 to 80, omit age 50, and cluster standard errors by country. Table 22 reports selected single-age coefficients over ages 20–64 and the contrast between ages 20 and 64.

Results. Define the manufacturing age gap as $G = \alpha_{20} - \alpha_{64}$ before occupation controls and $\tilde{G} = \tilde{\alpha}_{20} - \tilde{\alpha}_{64}$ after them. A positive gap indicates that manufacturing’s employment share is lower at age 64 than at age 20. The proportional reduction in this estimated gap is

$$\mathcal{A} = 1 - \frac{\tilde{G}}{G} = 1 - \frac{\tilde{\alpha}_{20} - \tilde{\alpha}_{64}}{\alpha_{20} - \alpha_{64}}. \quad (56)$$

In ISCO-08, the gap falls from 31.22 percentage points (standard error 5.48) to 4.66 points (standard error 0.90), a reduction of 85.1 percent. Thus, a substantially smaller but statistically significant age gradient remains after occupation controls. In ISCO-88, the gap falls from 14.32 points (standard error 6.77) to -0.48 points (standard error 4.60). The occupation-controlled point estimate is close to zero, although it is imprecisely estimated. Together, the estimates show that differences in detailed occupational composition account for much of the estimated manufacturing age gradient in these samples.

D.3 Occupational composition within manufacturing cohorts

Following age groups across censuses allows me to assess whether their occupational mix within manufacturing shifts away from physically demanding jobs as they grow older. Figure 31 shows this relationship. It follows five-year age groups across successive censuses, retaining observations between ages 20 and 64. Manufacturing here includes utilities and construction. For cohort–census pair h ,

Table 22: Manufacturing age profiles before and after controlling for occupations

Age	ISCO-08 sample		ISCO-88 sample	
	No occupation fixed effects	Three-digit occupation effects	No occupation fixed effects	Three-digit occupation effects
20	19.46 (4.53)	2.72 (0.84)	10.05 (5.08)	−0.52 (3.04)
25	10.98 (3.11)	2.10 (0.56)	4.98 (3.81)	−1.45 (2.75)
30	6.59 (2.35)	1.33 (0.41)	3.80 (3.05)	−1.17 (2.22)
35	5.13 (1.94)	0.86 (0.36)	2.78 (2.21)	−0.95 (1.69)
40	4.05 (1.13)	0.54 (0.22)	1.82 (1.51)	−0.64 (1.15)
45	2.55 (0.77)	0.29 (0.14)	0.82 (0.75)	−0.43 (0.54)
50 (reference)	0.00	0.00	0.00	0.00
55	−3.60 (0.69)	−0.57 (0.18)	−1.61 (0.80)	0.02 (0.64)
60	−7.82 (1.34)	−1.14 (0.25)	−2.79 (1.44)	0.29 (1.16)
64	−11.76 (1.97)	−1.93 (0.30)	−4.27 (2.02)	−0.04 (1.64)
Age 20 minus age 64	31.22 (5.48)	4.66 (0.90)	14.32 (6.77)	−0.48 (4.60)
Countries	17		36	
Censuses	21		54	
Three-digit occupations	173		126	
Country–census–age cells	1,276		3,269	
Regression cells	172,210		375,932	

Notes: Age coefficients are in percentage points relative to age 50; parentheses report country-clustered standard errors. Regressions include single-age effects over ages 20–80, country–birth-cohort fixed effects, and the Figure 1 controls. Within each classification, samples and weights are identical across columns. Manufacturing includes utilities and construction; shares use non-agricultural employment. Contrasts are calculated before rounding.

let $E_{oh\tau}$ denote employment in occupation o and $E_{h\tau}^M$ total manufacturing employment at endpoint $\tau \in \{0, 1\}$. The change in occupation o 's share of the cohort's manufacturing employment is

$$\Delta q_{oh} = \frac{E_{oh1}}{E_{h1}^M} - \frac{E_{oh0}}{E_{h0}^M}. \quad (57)$$

I average these changes using initial cohort manufacturing employment:

$$\overline{\Delta q_o} = 100 \frac{\sum_h E_{h0}^M \Delta q_{oh}}{\sum_h E_{h0}^M}, \quad \overline{\Delta q_o} = \zeta + \psi \theta_o^P + \nu_o. \quad (58)$$

The first expression gives each occupation's average change in percentage points, placing more weight on cohorts with larger initial manufacturing employment. The regression then places more weight on occupations with larger pooled initial employment, $\sum_h E_{oh0}$. Circle areas use these same occupation weights.

The sample covers 76 three-digit ISCO-88 occupations and 15 pairs of consecutive censuses in 13 countries. Across these census pairs, 109 age groups can be followed from the first census to the second. Intervals range from four to fifteen years. Age groups are advanced by the nearest multiple of five




Figure 31: Physical requirements and changes in occupational shares within manufacturing

Notes: Each circle is a three-digit ISCO-88 occupation. The vertical axis is the employment-weighted average change in the occupation's share of its cohort's manufacturing employment across successive censuses, in percentage points. The horizontal axis is the physical-requirement share, measured from 0 to 1. Circle areas and the fitted line use pooled initial occupation employment. The sample includes 76 occupations, 13 countries and 15 census pairs; manufacturing includes utilities and construction. Equations (57)–(58) define the plotted quantities and weights.

years to the census interval; ten pairs match the same birth years exactly, while the other five retain four of the original five birth years. Occupations must have an O*NET requirement measure, positive initial employment in at least five cohort–census-pair observations, and pooled initial employment of at least 5,000 survey-weighted workers. The manufacturing denominator includes all occupations, including those without a task score.²³

Occupational-share changes are more negative for jobs with stronger physical requirements. A 0.1 increase in physical requirements is associated with a 0.125 percentage-point lower change in occupational share over the observed census interval; the country-bootstrap 95 percent interval is $[-0.217, -0.028]$. These changes are measured relative to total manufacturing employment within a cohort, so a one-percentage-point change can be large for an individual occupation. For example, textile, fur and leather-products machine operators have a physical-requirement score of 0.81. Their average share falls from 3.22 to 2.05 percent across the observed census intervals: a decline of 1.17 percentage points, or about 36 percent of their initial share. General managers, with a physical-requirement score of 0.06, rise from 2.04 to 3.04 percent. The shares at both dates use the same initial-cohort manufacturing employment weights. These are observed changes for the named occupations; the

²³The countries are Bolivia, Botswana, Cambodia, El Salvador, Iran, Mauritius, Mozambique, Nicaragua, Portugal, Romania, South Africa, Trinidad and Tobago, and Uruguay. Switzerland is excluded because of incomplete occupation coding in its 2000 census.

regression summarizes the relationship across all 76 occupations. The figure thus shows occupational reallocation within manufacturing alongside the sectoral reallocation documented in Predictions 1 and 2.

Table 23 examines country influence, the manufacturing definition, cohort alignment and interval length. The relationship remains negative under every country omission, within manufacturing alone, among exactly matched cohorts, and after expressing each cohort change per decade. On the 73 occupations common to the two manufacturing definitions, the estimates are -0.125 and -0.137 percentage points per 0.1 increase in physical requirements.

Table 23: Physical requirements and within-manufacturing occupational change: supporting specifications

Specification	Estimate	95% interval	Occupations
Observed interval change	-0.125	$[-0.217, -0.028]$	76
Omit each country in turn	$[-0.142, -0.109]$	—	76
Common occupations: manufacturing, utilities and construction	-0.125	$[-0.218, -0.028]$	73
Common occupations: manufacturing alone	-0.137	$[-0.249, -0.078]$	73
Exactly matched cohorts	-0.150	$[-0.247, +0.075]$	76
Changes expressed per decade	-0.227	$[-0.404, -0.067]$	76
Subtract aggregate occupation-share change	-0.044	$[-0.074, -0.011]$	76

Notes: Entries are percentage-point changes in the occupation’s share of cohort manufacturing employment associated with a 0.1 increase in physical requirements. Changes cover the observed census interval unless expressed per decade. Confidence intervals are percentile intervals from 1,999 country-bootstrap draws, resampling all observations of a country together and recomputing both pooled changes and occupation weights, keeping the same set of occupations in every draw. The country-omission row reports the minimum and maximum estimates across all 13 omissions. The common-occupation rows retain the same 73 occupations, with denominators and employment weights appropriate to each manufacturing definition. Exact cohort matching retains ten census pairs in ten countries and 71 cohort–census-pair observations. The decade-rate row multiplies each cohort change by ten divided by its census interval before pooling. The aggregate-change row subtracts the occupation’s change in its share of national manufacturing employment at ages 20–64 for the same census pair.

The final specification asks whether the cohort relationship extends beyond changes in the country’s overall manufacturing occupation mix. For each occupation and census pair, I subtract its change in the share of manufacturing employment among all workers aged 20–64 from each cohort’s change in equation (57), before pooling. The estimated relationship remains negative, at -0.044 percentage points per 0.1 increase in physical requirements, with a country-bootstrap interval of $[-0.074, -0.011]$. The relationship between physical requirements and occupational-share changes is therefore more negative within cohorts than in the corresponding changes in the overall manufacturing occupation mix.

D.4 Physical requirements and youth intensity across manufacturing divisions

The ILOSTAT exercise asks whether manufacturing divisions with stronger physical requirements also employ a larger share of young workers. Let Y_{kctv} and O_{kctv} denote employment at ages 15–24 and 25 or older in division k , country c , year t and industry-classification version v . Define the division’s

youth share and its deviation from the manufacturing-wide share as

$$y_{kctv} = \frac{Y_{kctv}}{Y_{kctv} + O_{kctv}}, \quad d_{kctv} = 100 \left(y_{kctv} - \frac{\sum_k Y_{kctv}}{\sum_k (Y_{kctv} + O_{kctv})} \right).$$

The manufacturing-wide share sums divisions with both age groups observed. When employment at ages 25 or older is unavailable directly, it is obtained by subtracting employment at ages 15–24 from employment at ages 15 or older. I retain country–year–classification observations covering at least fifteen manufacturing divisions and exclude division observations whose reported youth or adult employment is flagged as unreliable. The sample covers 1993–2025 and the 24 manufacturing divisions in ISIC Revision 4 and 23 in Revision 3; utilities and construction are excluded.

Each division’s physical-requirement measure, θ_{kv}^P , combines O*NET occupational requirements with ILOSTAT occupation employment shares. Three-digit requirements are first averaged within one-digit ISCO major groups using employment in manufacturing alone as weights in both ISCO-08 and ISCO-88. These major-group scores are then averaged using the division’s occupation shares and pooled across country–years with division employment weights. The occupational-mix sample requires at least six of the nine civilian major groups and, whenever a division total is reported, coverage of at least 90 percent of that total. This gives a fixed requirement measure for each division and industry-classification version.

I estimate

$$d_{kctv} = \delta_{ctv} + \gamma \theta_{kv}^P + \eta g_{kctv} + \epsilon_{kctv}, \quad (59)$$

using division employment weights and country-clustered standard errors. These fixed effects compare divisions within the same country, year and industry classification. The first specification omits g_{kctv} ; the second controls for trailing division employment growth relative to total manufacturing employment growth. Specifically,

$$g_{kctv} = \frac{10}{L} \left[\log \frac{E_{kctv}}{E_{kc,t-L,v}} - \log \frac{E_{ctv}^M}{E_{c,t-L,v}^M} \right],$$

where the lag L is the longest available interval from two to five years within the same industry classification, and growth intervals do not cross a reported break in the division’s series. Employment weights use reported employment at ages 15 or older, or the sum of the two age groups when that total is unavailable.

Table 24 shows that more physically intensive manufacturing divisions employ a larger share of young workers. A 0.1 increase in physical requirements is associated with a 0.928 percentage-point increase in the division’s young-worker share relative to the manufacturing-wide share. The estimate is 0.928 percentage points when relative employment growth is included. The positive relationship therefore also holds after accounting for recent differences in division employment growth.

Table 24: Physical requirements and the youth intensity of manufacturing divisions

	Baseline	Controlling for relative growth
Physical requirements: effect per 0.1 increase	0.928 (0.309)	0.928 (0.316)
Relative division employment growth	—	1.112 (0.510)
Country–year–classification fixed effects	Yes	Yes
Division employment weights	Yes	Yes
Observations	12,411	7,744
Countries	63	50

Notes: The outcome is the share of division employment at ages 15–24 minus the corresponding share for manufacturing, in percentage points. Physical-requirement coefficients and standard errors are scaled to a 0.1 increase in the 0–1 score. Relative division growth is the difference between division and manufacturing log employment growth, expressed per decade over a preceding two- to five-year interval. The second column uses observations with this growth measure available. Country-clustered standard errors are in parentheses. Division definitions are specific to each ISIC revision. Source: ILOSTAT and O*NET.

E Prediction 4: Ability aging and relative sectoral productivity

E.1 Sources and age patterns

I construct the common age profiles $\rho_k(a)$ from published evidence on physical and cognitive aging. The underlying studies measure physical performance and cognitive test scores rather than productivity in particular sectors. I first express their age-related differences in standard deviations across individuals, making changes in the different abilities comparable, and construct a curve for each ability at every working age from 20 to 64.

The three abilities are physical (P), fluid-cognitive (F), and crystallized-and-interpersonal (C). The construction has two uses. In the quantitative model, the profiles determine how the productivity of each ability changes with age. In the empirical exercise, combining them with sectoral ability requirements gives relative-productivity profiles that can be compared with observed employment patterns. The steps below distinguish these two uses.

For physical ability, I draw on studies of flexibility and grip strength. Medeiros, de Araújo, and de Araújo (2013) report age-related declines in flexibility of approximately 0.6–0.8 percent per year. Grip-strength evidence also shows substantial deterioration during later working life, with mean strength falling by roughly 20 percent from its early-adult peak to age 65 (Dodds, Syddall, Cooper et al., 2014). I use the midpoint of the flexibility estimates to construct a linear decline of 0.7 percent of the reference level per year. To express this profile in units comparable to the cognitive measures, I use a standard deviation equal to 20 percent of the reference mean, informed by the dispersion in grip strength. This gives an annual decline of 0.035 standard deviations.

For fluid-cognitive ability, Salthouse (2009) reports declines from young adulthood to age 60 of approximately one standard deviation for processing speed and spatial visualization, compared with 0.6–0.7 standard deviations for reasoning and memory. I place the peak at age 24.5, the midpoint

of the peak ages 22–27 reported across the study’s twelve cognitive measures. I use an intermediate magnitude of 0.85 standard deviations for the decline from this peak to age 60. This corresponds to an annual decline of approximately 0.024 standard deviations, within the 0.02–0.03 range reported in the same study. The study also documents faster declines after age 60. To capture this acceleration, I increase the annual decline by 50 percent between 60 and 65, to approximately 0.036 standard deviations.

For the crystallized-and-interpersonal bundle, I construct the age profile using vocabulary as a measure of accumulated knowledge. Verhaeghen (2003) reports an average vocabulary advantage of 0.80 standard deviations for older relative to younger adults (average ages of 70 and 21, respectively). I use this difference to anchor the increase from young adulthood to age 65, the retirement age in the model. This late peak is supported by Hartshorne and Germine (2015), who find that vocabulary peaks around 65 in their sample, and Salthouse (2014), whose estimates across four vocabulary measures place the turning point between ages 64 and 66. To capture this rise and eventual flattening, I use a concave quadratic in standardized scores that reaches a plateau at age 65.

E.2 Construction of the annual profiles

Annual scores. Let $\mathcal{A}^W = \{20, \dots, 64\}$ denote the working-age grid. For each ability $k \in \{P, F, C\}$, let $\tilde{g}_k(a)$ denote the constructed score at age a , expressed in standard-deviation units before normalization. The following formulas specify the values between the reference ages described above.

For physical ability, let $d_P = 0.007$ denote the annual decline as a fraction of the reference level and let $v_P = 0.20$ denote the standard deviation divided by that reference level. Dividing the former by the latter converts the annual decline into standard-deviation units:

$$\tilde{g}_P(a) = -\frac{d_P}{v_P}(a - 20) = -0.035(a - 20). \quad (60)$$

The score is zero at age 20 and falls by 0.035 standard deviations per year. The decline is linear in the score: it is not a percentage reduction compounded each year.

For fluid cognition, let $p_F = 24.5$ denote the peak age and $d_F = 0.85$ the decline from the peak to age 60. The annual decline over this interval is

$$b_F = \frac{d_F}{60 - p_F}.$$

Before the peak, I use a linear increase of 0.015 standard deviations per year. After age 60, the annual decline increases to $1.5b_F$. The full profile is

$$\tilde{g}_F(a) = \begin{cases} 0.015(a - p_F), & a \leq p_F, \\ -b_F(a - p_F), & p_F < a \leq 60, \\ -d_F - 1.5b_F(a - 60), & a > 60. \end{cases} \quad (61)$$

These expressions join continuously at the peak and at age 60. The peak is defined at age 24.5, but the constructed profile is evaluated at integer ages; its highest annual observation is therefore at age 24.

For crystallized-and-interpersonal ability, let $p_C = 65$ denote the plateau age and $d_C = 0.80$ the increase from young adulthood to that plateau. Define

$$x_C(a) = \min \left\{ 1, \max \left\{ 0, \frac{a - 22}{p_C - 22} \right\} \right\}.$$

The variable $x_C(a)$ measures progress from age 22 to the plateau age, with values restricted to the interval from zero to one. The score is

$$\tilde{g}_C(a) = d_C \left[1 - (1 - x_C(a))^2 \right]. \quad (62)$$

It is flat through age 22, increases thereafter, and becomes flat again at age 65. Between these ages,

$$\frac{\partial \tilde{g}_C(a)}{\partial a} = \frac{2d_C}{p_C - 22} \left(1 - \frac{a - 22}{p_C - 22} \right), \quad \frac{\partial^2 \tilde{g}_C(a)}{\partial a^2} = -\frac{2d_C}{(p_C - 22)^2}.$$

Thus the gains diminish continuously and reach zero at the plateau. Age 65 anchors the curve, but only ages 20–64 enter the working-age calculations.

Converting scores into positive ability factors. The model represents abilities as positive factors that multiply a worker's permanent endowments and schooling component. I obtain these factors by exponentiating the constructed scores and dividing by their working-age arithmetic mean:

$$\rho_k(a) = \frac{\exp\{\tilde{g}_k(a)\}}{\frac{1}{45} \sum_{u=20}^{64} \exp\{\tilde{g}_k(u)\}}, \quad k \in \{P, F, C\}. \quad (63)$$

The same profiles are denoted by $\rho_k^{\text{pub}}(a)$ in the empirical comparison.

Exponentiation uses a common unit conversion from standardized scores to log ability productivity. Consequently,

$$\log \rho_k(a) - \log \rho_k(a') = \tilde{g}_k(a) - \tilde{g}_k(a'). \quad (64)$$

A one-standard-deviation difference in the constructed score therefore becomes a one-unit difference in log ability productivity. This conversion is a modeling assumption.

The denominator in (63) sets each profile's arithmetic mean to one:

$$\frac{1}{45} \sum_{a=20}^{64} \rho_k(a) = 1.$$

Unlike the conversion into productivity units, this division only fixes the level of the index. It does not change proportional differences between ages. Values above one indicate productivity above that ability's working-age average; they do not compare the absolute productivity of different abilities.

Figure 32 shows the resulting profiles.




Figure 32: Ability profiles over the working life

Notes: Ability-productivity factors constructed from the physical, fluid-cognitive, and crystallized-and-interpersonal score curves using equation (63). Each factor has arithmetic mean one over the 45 integer ages from 20 to 64. The same unit conversion from standardized scores to log productivity is applied to all three abilities.

E.3 From ability profiles to worker productivity

Sectoral ability requirements. Sectors use the three abilities with different intensities. I measure their requirements using occupation-level ratings from the Abilities and Generalized Work Activities modules of the Occupational Information Network, O*NET (National Center for O*NET Development, 2024). Physical requirements include strength, dexterity, and manual activities. Fluid-cognitive requirements include reasoning, attention, information processing, and problem solving. Crystallized-and-interpersonal requirements include acquired knowledge, verbal comprehension, communication, teaching, care, and coordination.

Let $\vartheta_{o,k}$ denote occupation o 's normalized requirement for ability k , with

$$\vartheta_{o,k} \geq 0, \quad \sum_{k \in \{P,F,C\}} \vartheta_{o,k} = 1.$$

Appendix A describes the underlying O*NET measures and occupational mappings.

For the empirical sector groups, I average these occupation-level shares:

$$\theta_{j,k} = \frac{1}{|\mathcal{O}_j|} \sum_{o \in \mathcal{O}_j} \vartheta_{o,k}, \quad (65)$$

where $\mathcal{O}_j$ is the set of occupations assigned to sector j . Each included occupation receives equal weight; the average is not weighted by employment. The sectoral weights therefore also sum to one.

Table 25 reports the weights used in the employment comparison. Manufacturing includes utilities and construction. Manufacturing places more weight on physical requirements than either service sector, while both service sectors place more weight on crystallized-and-interpersonal requirements.

Table 25: Ability requirements in the non-agricultural sectors

Sector	Physical	Fluid-cognitive	Crystallized-and-interpersonal
Manufacturing	0.555	0.315	0.130
Low-skill services	0.252	0.302	0.446
High-skill services	0.044	0.455	0.501

Notes: O*NET requirement weights $\theta_{j,k}$, calculated as unweighted averages of normalized occupation-level requirements within each sector group. The weights sum to one within each sector before rounding. These are the weights used to construct the relative-productivity profiles for manufacturing and the two service sectors.

Combining the ability profiles. For a sector with ability weights $\theta_{j,k}$, the common age component of log efficiency is

$$\ell_j(a) = \sum_{k \in \{P,F,C\}} \theta_{j,k} \log \rho_k(a). \quad (66)$$

Equivalently, its efficiency component in levels is

$$\exp\{\ell_j(a)\} = \prod_{k \in \{P,F,C\}} \rho_k(a)^{\theta_{j,k}}.$$

Thus a sector's profile combines the same three ability curves as every other sector, but gives them different weights. For example, the decline in physical ability has a larger effect on manufacturing because manufacturing places more weight on that ability.

Sectoral normalization in the quantitative model. The quantitative model normalizes each sector's common efficiency profile to have arithmetic mean one over working ages:

$$\tilde{z}_j(a) = \frac{\exp\{\ell_j(a)\}}{\frac{1}{45} \sum_{u=20}^{64} \exp\{\ell_j(u)\}} = \frac{\prod_{k \in \{P,F,C\}} \rho_k(a)^{\theta_{j,k}}}{\frac{1}{45} \sum_{u=20}^{64} \prod_{k \in \{P,F,C\}} \rho_k(u)^{\theta_{j,k}}}. \quad (67)$$

This additional normalization is needed because a product of unit-mean ability profiles has mean below one. It fixes the units of the sectoral age component; the corresponding level is absorbed into sectoral productivity for agriculture and services and into the 1993 normalization of manufacturing productivity.

Substituting (63) into (67) gives

$$\tilde{z}_j(a) = \frac{\exp \left\{ \sum_{k \in \{P, F, C\}} \theta_{j,k} \tilde{g}_k(a) \right\}}{\frac{1}{45} \sum_{u=20}^{64} \exp \left\{ \sum_{k \in \{P, F, C\}} \theta_{j,k} \tilde{g}_k(u) \right\}}. \quad (68)$$

The ability-specific normalization constants cancel between the numerator and denominator. The model can therefore construct exactly the same sectoral profiles either from the normalized ability indices or directly from the constructed scores.

This is only the common age component of worker productivity. Individual workers also differ in their permanent endowments of each ability and in schooling, which affects the three abilities with different intensities.

E.4 Comparison with sectoral employment by age

The empirical exercise asks whether the age-related changes in relative productivity implied by these inputs match the direction and shape of the employment profiles in Prediction 1. It uses the common age component constructed above.

Relative productivity across sectors. Let $\mathcal{J}^{NA} = \{M, L, H\}$ denote manufacturing, low-skill services, and high-skill services. To compare sectors at the same age, subtract the equally weighted mean of their log-efficiency profiles:

$$r_j(a) = \ell_j(a) - \frac{1}{3} \sum_{h \in \mathcal{J}^{NA}} \ell_h(a), \quad j \in \mathcal{J}^{NA}. \quad (69)$$

The reference includes all three sectors, including sector j itself. A rise in $r_j(a)$ means that the sector's ability-implied efficiency improves relative to the three-sector average.

This subtraction removes changes common to all three sectors. In particular, defining

$$\bar{\theta}_k = \frac{1}{3} \sum_{h \in \mathcal{J}^{NA}} \theta_{h,k},$$

gives

$$r_j(a) - r_j(a') = \sum_{k \in \{P, F, C\}} (\theta_{j,k} - \bar{\theta}_k) [\tilde{g}_k(a) - \tilde{g}_k(a')]. \quad (70)$$

A sector's relative age profile therefore depends on whether it uses each ability more or less intensively than the sectoral average, and on how that ability changes with age. Normalization constants that do not vary with age disappear from this comparison.

The employment profiles. I use the schooling-controlled, single-age employment coefficients from Prediction 1, equation (1), estimated on the 88-country census sample. Let $\hat{\alpha}_a^j$ denote the estimated coefficient for age a in sector j .

Putting both profiles on the same age basis. I compare nine five-year bands:

$$\mathcal{B} = \{\{20, \dots, 24\}, \{25, \dots, 29\}, \dots, \{60, \dots, 64\}\}.$$

Each annual age receives equal weight within a band. Both the relative-productivity profiles and the estimated employment profiles are expressed relative to ages 50–54.

For relative productivity, define the annual profile relative to this reference band as

$$\check{r}_j(a) = r_j(a) - \frac{1}{5} \sum_{u=50}^{54} r_j(u).$$

Its band average is

$$r_{jB} = \frac{1}{5} \sum_{a \in B} \check{r}_j(a) = \frac{1}{5} \sum_{a \in B} r_j(a) - \frac{1}{5} \sum_{u=50}^{54} r_j(u). \quad (71)$$

For the employment coefficients, the corresponding quantity is

$$e_{jB} = \frac{1}{5} \sum_{a \in B} \hat{\alpha}_a^j - \frac{1}{5} \sum_{u=50}^{54} \hat{\alpha}_u^j. \quad (72)$$

Both profiles therefore have a band average of zero at ages 50–54. This common reference allows their changes with age to be compared without requiring productivity levels and employment-share levels to have the same meaning.

Figure 3 shows the employment estimates $100e_{jB}$ as dots, measured in percentage points on the left axis, and the annual relative-productivity profiles $\check{r}_j(a)$ as lines, measured in log units on the right axis. Both are expressed relative to ages 50–54.

Measuring agreement across age groups. For each sector, I calculate the Pearson correlation between r_{jB} and e_{jB} across the nine age bands:

$$\text{Corr}_j = \frac{\sum_{B \in \mathcal{B}} (r_{jB} - \bar{r}_j)(e_{jB} - \bar{e}_j)}{\sqrt{\sum_{B \in \mathcal{B}} (r_{jB} - \bar{r}_j)^2 \sum_{B \in \mathcal{B}} (e_{jB} - \bar{e}_j)^2}}, \quad (73)$$

where $\bar{r}_j$ and $\bar{e}_j$ are the respective averages across the nine bands. Each band receives equal weight. These are correlations of the profile values, not rank correlations.

The correlations are 0.997 for manufacturing, 0.905 for low-skill services, and 0.981 for high-skill services. Between ages 20–24 and 60–64, the estimated manufacturing employment profile declines by 18.0 percentage points, while low- and high-skill services rise by 8.0 and 8.8 points. The ability-implied profiles move in the same directions. The underlying ability measurements and occupational requirements thus generate relative-productivity patterns that closely track the age patterns measured separately in the census employment data.

E.5 Sensitivity to the construction of the ability profiles

The sensitivity exercises vary the timing and magnitude of the constructed ability changes. They assess the stability of the relative-productivity comparison, rather than re-solving the quantitative model or re-estimating the census regressions.

I vary the cognitive and physical profiles separately to distinguish how assumptions about each affect the results. In each exercise, the other profiles remain at their baseline values.

Cognitive profiles. I consider every combination of four parameters:

$$\begin{aligned} p_F &\in \{19.5, 24.5, 29.5\}, \\ p_C &\in \{60, 65, 70\}, \\ d_F &\in \{0.425, 0.850, 1.275\}, \\ d_C &\in \{0.40, 0.80, 1.20\}. \end{aligned}$$

These choices give $3^4 = 81$ combinations. The peak and plateau ages vary by five years on either side of the baseline values; the fluid-cognitive decline and crystallized-and-interpersonal increase vary between one-half and one-and-a-half times their baseline magnitudes.

For each combination, I reconstruct the fluid-cognitive and crystallized-and-interpersonal score curves using equations (61) and (62). The fluid pre-peak increase remains 0.015 standard deviations per year, and its decline after age 60 remains 1.5 times the preceding annual decline. The physical profile and sectoral requirement weights are unchanged.

Physical profile. In a separate exercise, I vary the annual physical decline and the standard deviation used to convert it into standardized units:

$$d_P \in \{0.006, 0.007, 0.008\}, \quad v_P \in \{0.15, 0.20, 0.25\}.$$

These choices give nine combinations. For each, the physical score is $\tilde{g}_P(a) = -(d_P/v_P)(a - 20)$. The cognitive profiles and sectoral weights remain at their baseline values. This exercise varies both the age gradient taken from the flexibility evidence and the assumed dispersion used to express that gradient in standard-deviation units.

Comparing the alternative profiles with employment. Each parameter combination produces a different set of ability curves. I use these curves to recalculate relative sectoral productivity, keeping the O*NET weights and census employment estimates unchanged. I then check the change in relative productivity between ages 20 and 64 and calculate its correlation with employment across the nine age bands. The exercise therefore asks whether the predicted manufacturing decline, the service-sector increases, and the agreement with employment persist under different assumptions about ability aging.

The direction check asks whether

$$r_M(64) - r_M(20) < 0, \quad r_L(64) - r_L(20) > 0, \quad r_H(64) - r_H(20) > 0.$$

It concerns the change between ages 20 and 64, not whether every annual movement has the same sign.

These three directions hold in all 81 cognitive combinations and all nine physical combinations. Table 26 reports the correlation ranges, and Figure 33 shows the full sets of profiles.

Table 26: Sensitivity of relative-productivity profiles to the ability curves

	Combinations with predicted directions	Correlation with employment profile		
		Manufacturing	Low-skill services	High-skill services
Baseline	1 of 1	0.997	0.905	0.981
Cognitive profiles	81 of 81	0.994–0.998	0.893–0.914	0.973–0.982
Physical profile	9 of 9	0.996–0.997	0.904–0.907	0.980–0.981

Notes: Predicted directions refer to the change in relative log productivity between ages 20 and 64: negative for manufacturing and positive for both service sectors. Correlations are Pearson correlations across nine five-year age bands, using the schooling-controlled employment profiles from the 88-country census sample. Ranges report the minimum and maximum across the parameter combinations in each exercise. The cognitive and physical exercises are conducted separately.

F Household age and demand for manufactured goods

Household demand for manufactures also varies with age. Younger households devote a larger share of expenditure to durable goods, such as vehicles, appliances, furniture, and household equipment. These purchases build stocks that remain useful for years, so holdings acquired early in adulthood can reduce the need for new purchases later in life. This complements the evidence in Cravino, Levchenko, and Rojas (2022) on the rising services share of expenditure with age.

Using Eurostat data on household consumption by COICOP category (Classification of Individual Consumption According to Purpose) and age of the reference person, I find that the manufactured-durable budget share is about 4 percentage points higher for households whose reference person is under 30 than for those aged 60 and above.²⁴ Restricting the Eurostat sample to middle-income country-years yields very similar results.

Indonesia provides additional evidence for a middle-income economy undergoing structural transformation. Using all five IFLS waves from 1993 to 2014 (Appendix A.2), I group households by the mean age of their members aged 20 and above and calculate survey-weighted expenditure shares within

²⁴The empirical aggregate combines durable and semi-durable COICOP categories; every included category is assigned to manufacturing in the paper’s concordance. Appendix A, Table 11, lists the categories.




Figure 33: Relative sectoral productivity under alternative ability profiles

Notes: The upper panels show the 81 cognitive-profile combinations; the lower panels show the nine physical-profile combinations. Thick lines show the baseline relative-productivity profiles, thin lines the alternatives, and dots the estimated employment profiles. Lines use annual ages; dots are five-year averages. All profiles are expressed relative to ages 50–54. Employment coefficients use the left axis, in percentage points; relative-productivity profiles use the right axis, in log units.

each wave. In every IFLS wave, households with a mean adult age below 30 devote a larger share of expenditure to manufactured goods than those with a mean adult age of 60 or above. The gap averages 6.1 percentage points, of which 5.5 points—nine-tenths—come from durable goods.²⁵ Younger households also spend more overall, so the difference is proportionately larger in expenditure levels: households with a mean adult age below 30 spend, on average, 3.5 times as much on durables as those with a mean adult age of 60 or above.²⁶ The baseline model incorporates age differences in sectoral expenditure shares through age-specific preferences. The concentration of the manufacturing gap in durables motivates the extension in Section 3.9, which explicitly models how stocks accumulated over the life cycle affect the need for new purchases.

G Robustness of the within-country employment relationship

This appendix examines the robustness of the relationship between young-adult cohort size and sectoral employment. It documents the cohort instruments and the age pattern of the first stage, assesses economic conditions at birth and regional annual shocks, and tests stability across controls and countries. Manufacturing-specific and subnational data provide complementary evidence. It also

²⁵Manufactured non-durables account for the remaining 0.6 points.

²⁶This figure averages the young–old expenditure ratios across the five IFLS waves, giving each wave equal weight.

examines whether trade openness strengthens the demographic employment relationship.

G.1 Full baseline and second-stage estimates

Table 27 reports every non-fixed-effect coefficient from the specifications summarized in the main text.

G.2 Cohort instruments and the first-stage age fingerprint

The main specification uses the mean crude birth rate 25–34 years earlier to instrument the current population share aged 25–34. The surviving-child-stock instrument reconstructs the same cohort from population stocks and survival probabilities instead. Its numerator ages the population observed at ages 5–14 in $t-20$ forward to ages 25–34; its denominator ages the entire $t-20$ population forward and therefore predicts the population aged 20 and above in t . Table 28 reports country-clustered estimates together with wild-cluster-bootstrap and weak-instrument-robust inference. Each instrument is strong on its own and gives a positive, statistically significant coefficient for both employment outcomes. Using both instruments jointly produces estimates of 0.730 for industry’s share of total employment and 0.694 for its share within non-agriculture; the overidentifying restrictions are not rejected. The two instruments reconstruct the same cohorts from different underlying measurements, so their agreement tests measurement robustness rather than exclusion threats common to both.

The relevance of the birth-rate instrument has a sharply age-specific form. I estimate

$$young_{it} = \sum_{\ell \in \mathcal{L}} b_{\ell} BR_{i,t-\ell} + \alpha_i + \gamma_t + \lambda_1 \log GDPpc_{it} + \lambda_2 (\log GDPpc_{it})^2 + \mathbf{X}'_{it}\theta + u_{it}, \quad (74)$$

where $\mathcal{L} = \{10-14, \dots, 40-44\}$ and each $BR_{i,t-\ell}$ is the mean birth rate over the indicated five-year lag window. All seven windows enter jointly, so each coefficient isolates the predictive content of births at that cohort age conditional on the other windows and the full baseline controls.

Figure 34 shows the resulting fingerprint. The coefficients are positive and precisely estimated for the two windows corresponding to the current 25–34 population; they are negative or indistinguishable from zero at shorter lags and indistinguishable from zero at longer lags. Thus, the positive first-stage coefficients are concentrated exactly where historical births map into the endogenous cohort share, with the offsetting signs for younger cohorts implied by population-share arithmetic. This is a diagnostic for cohort mapping and relevance, not a separate test of exclusion.

G.3 Control sensitivity and reduced-form estimates

Several baseline controls may themselves respond to demographic change. Capital accumulation and productivity, in particular, may adjust to changes in the age composition of the workforce. Conditioning on these responses can remove part of the relationship of interest and introduce selection

Table 27: Full estimates for young-adult cohort size and sectoral employment shares

	OLS	Birth-rate 2SLS	Child-stock 2SLS
<i>Panel A. Agriculture / total employment</i>			
Young-adult share (25–34)	−0.307 (0.138)	−0.258 (0.323)	−0.260 (0.207)
Log GDP per capita	−0.320 (0.056)	−0.321 (0.055)	−0.321 (0.056)
(Log GDP per capita) ²	0.018 (0.003)	0.017 (0.003)	0.017 (0.003)
Female labor-force participation	0.145 (0.080)	0.147 (0.080)	0.147 (0.080)
Mean years of schooling, ages 25+	−0.001 (0.004)	−0.001 (0.004)	−0.001 (0.004)
Imports / GDP	−0.021 (0.015)	−0.020 (0.015)	−0.020 (0.014)
Log capital–output ratio	−0.054 (0.019)	−0.053 (0.019)	−0.053 (0.019)
Investment share	−0.030 (0.024)	−0.031 (0.023)	−0.031 (0.023)
Log TFP	−0.061 (0.025)	−0.060 (0.025)	−0.060 (0.024)
<i>Panel B. Industry / total employment</i>			
Young-adult share (25–34)	0.863 (0.184)	0.746 (0.313)	0.730 (0.180)
Log GDP per capita	0.144 (0.064)	0.146 (0.066)	0.146 (0.066)
(Log GDP per capita) ²	−0.006 (0.004)	−0.006 (0.004)	−0.006 (0.004)
Female labor-force participation	−0.183 (0.052)	−0.187 (0.054)	−0.187 (0.054)
Mean years of schooling, ages 25+	−0.005 (0.004)	−0.005 (0.004)	−0.005 (0.004)
Imports / GDP	−0.010 (0.014)	−0.012 (0.015)	−0.013 (0.014)
Log capital–output ratio	0.013 (0.014)	0.012 (0.014)	0.012 (0.014)
Investment share	0.087 (0.020)	0.090 (0.021)	0.091 (0.021)
Log TFP	−0.011 (0.017)	−0.014 (0.019)	−0.015 (0.019)
<i>Panel C. Services / total employment</i>			
Young-adult share (25–34)	−0.556 (0.211)	−0.487 (0.336)	−0.470 (0.184)
Log GDP per capita	0.176 (0.060)	0.175 (0.061)	0.175 (0.060)
(Log GDP per capita) ²	−0.012 (0.003)	−0.012 (0.003)	−0.012 (0.003)
Female labor-force participation	0.038 (0.068)	0.040 (0.068)	0.040 (0.069)
Mean years of schooling, ages 25+	0.007 (0.004)	0.006 (0.004)	0.006 (0.004)
Imports / GDP	0.031 (0.012)	0.032 (0.013)	0.033 (0.012)
Log capital–output ratio	0.041 (0.014)	0.041 (0.015)	0.042 (0.014)
Investment share	−0.057 (0.020)	−0.059 (0.021)	−0.060 (0.021)
Log TFP	0.072 (0.024)	0.074 (0.026)	0.075 (0.027)
<i>Panel D. Industry / non-agricultural employment</i>			
Young-adult share (25–34)	0.840 (0.210)	1.027 (0.350)	0.696 (0.226)
Log GDP per capita	0.090 (0.078)	0.087 (0.074)	0.092 (0.080)
(Log GDP per capita) ²	−0.003 (0.004)	−0.003 (0.004)	−0.003 (0.004)
Female labor-force participation	−0.110 (0.059)	−0.104 (0.060)	−0.114 (0.061)
Mean years of schooling, ages 25+	−0.007 (0.004)	−0.007 (0.004)	−0.007 (0.004)
Imports / GDP	−0.011 (0.016)	−0.007 (0.018)	−0.013 (0.015)
Log capital–output ratio	0.011 (0.017)	0.013 (0.017)	0.010 (0.018)
Investment share	0.083 (0.024)	0.078 (0.027)	0.087 (0.026)
Log TFP	−0.018 (0.025)	−0.012 (0.025)	−0.022 (0.027)
Country and year fixed effects	Yes	Yes	Yes
Observations	3,650	3,650	3,650
Countries	147	147	147
Kleibergen–Paap rk Wald F	—	87.1	186.4

Notes: The table reports all non-fixed-effect coefficients from equation (3); the IV columns report second-stage estimates. Within each cell, the coefficient is followed by its country-clustered standard error in parentheses. All specifications use the same observations and controls as Table 3. Employment and population shares enter as fractions.

Table 28: Alternative cohort instruments and weak-IV-robust inference

Outcome	Instrument(s)	2SLS coefficient	K–P F	Wild p	AR p	Hansen p
Industry / total employment	Lagged birth rate	0.746 (0.313)	87.1	0.025	0.038	—
	Surviving child stock	0.730 (0.180)	186.4	< 0.001	< 0.001	—
	Both	0.730 (0.179)	110.3	< 0.001	0.002	0.946
Industry / non-agricultural employment	Lagged birth rate	1.027 (0.350)	87.1	0.004	0.007	—
	Surviving child stock	0.696 (0.226)	186.4	0.005	0.007	—
	Both	0.694 (0.226)	110.3	0.006	0.007	0.261

Notes: The dependent variables and covariates are those in equation (3). All regressions use the same 3,650 country–year observations from 147 countries and include country and year fixed effects. Standard errors clustered by country are in parentheses. K–P F is the cluster-robust Kleibergen–Paap first-stage statistic. Wild p is the null-imposed wild-cluster-bootstrap p -value; AR p is its Anderson–Rubin counterpart. Both use Rademacher weights, 9,999 replications, and country clusters. Hansen p is reported for the overidentified specifications.




Figure 34: The first-stage age fingerprint of lagged birth rates. *Notes:* The figure plots $100\hat{b}_\ell$ from equation (74); the dependent variable is the share of the total population aged 25–34, and birth rates are measured per 1,000 population. All seven five-year birth-rate windows enter jointly. The shaded area marks the two windows that map into ages 25–34. The specification includes the full controls in equation (3) and country and year fixed effects. The estimation sample contains 3,128 country–year observations from 147 countries. Bars are 95 percent confidence intervals based on country-clustered standard errors.

bias, rather than simply separate a direct effect from an indirect one (Angrist and Pischke, 2009, Section 3.2.3). This motivates repeating the analysis without contemporaneous economic controls. Because these variables may also capture differences in development conditions, the exercise assesses sensitivity to their inclusion. Omitting them does not establish instrument validity or automatically identify a total causal effect.

I remove the quadratic in log GDP per capita, female labor-force participation, mean schooling, imports relative to GDP, the log capital–output ratio, the investment share, and log TFP from equation (3), retaining the two cohort instruments in separate regressions.

For these comparisons, I retain the same 3,650 country–year observations from 147 countries over 1991–2016. The outcome is industry’s share of total employment. Standard errors are clustered by country.

Table 29 shows that the positive demographic coefficient does not depend on these controls. Without them, a one-percentage-point larger young-adult share predicts an approximately one-percentage-point higher industry employment share, compared with 0.746 and 0.730 points in the fully controlled specifications.

Reduced forms and their units. The reduced form relates employment directly to the historical cohort instrument, rather than to the current young-adult share. For each instrument Z_{it} , I estimate

$$s_{it}^j = r_j Z_{it} + \mathbf{C}_{it}' \psi_j + \alpha_i + \gamma_t + u_{it}, \quad (75)$$

$$\text{Young}_{it} = f Z_{it} + \mathbf{C}_{it}' \psi_F + \alpha_i^F + \gamma_t^F + v_{it}. \quad (76)$$

The first equation is the reduced form; the second is the first stage. The vector $\mathbf{C}_{it}$ contains all baseline controls, including the GDP quadratic, in the fully controlled regressions and is omitted in the fixed-effects-only regressions. Within each specification, both equations use exactly the same observations and controls as the corresponding 2SLS regression.

The two instruments have different units. The birth-rate instrument is the average annual number of births per 1,000 people over years $t - 34, \dots, t - 25$. I report the employment and young-adult-share changes associated with a five-unit increase in this average, for example from 25 to 30 births per 1,000. For the child-stock instrument, I report a one-percentage-point increase in the reconstructed cohort's share of the projected population aged 20 and above (Appendix G.2). This is distinct from a one-point increase in the observed young-adult share of the total population. Since employment and population shares enter the regressions as fractions, the displayed birth-rate coefficients are multiplied by 500; the child-stock coefficients numerically equal the percentage-point changes for a one-point increase in that instrument.

Table 29 reports reduced forms alongside their corresponding first stages. With full controls, five additional historical births per 1,000 are associated with a 1.106-percentage-point higher industry employment share and a 1.483-point larger current young-adult share. For a one-point increase in the child-stock instrument, these changes are 0.425 and 0.582 points.

With one endogenous regressor and one excluded instrument, the 2SLS point estimate equals the reduced-form coefficient divided by the first-stage coefficient. Thus, expressing these reduced forms per one-percentage-point larger current young-adult share recovers the baseline IV estimates:

$$\hat{\beta}_{\text{birth}} = \frac{1.106494}{1.483353} \simeq 0.746, \quad \hat{\beta}_{\text{child}} = \frac{0.425232}{0.582315} \simeq 0.730.$$

The denominators are estimated first-stage changes for the stated instrument increments, not fixed conversion factors. They change when the controls or sample change. The standard errors of the ratios are those from 2SLS, which account for estimation of the first stage; they are not obtained by dividing the reduced-form standard errors by these denominators.

The birth-rate reduced form is almost unchanged when controls are removed: 1.120 rather than 1.106 employment-share points. Its first stage is smaller, however, so the corresponding IV coefficient is larger. The change in the IV coefficient therefore should not be interpreted as the amount of the demographic effect operating through the controls.

Table 29: Reduced forms, first stages, and IV estimates

Instrument	Controls	Industry share	Young-adult share	2SLS
		reduced form	first stage	
Birth rate	Full	1.106 (0.492)	1.483 (0.159)	0.746 (0.313)
	None	1.120 (0.454)	1.154 (0.173)	0.970 (0.377)
Child stock	Full	0.425 (0.112)	0.582 (0.043)	0.730 (0.180)
	None	0.485 (0.114)	0.485 (0.040)	1.001 (0.236)

Notes: Reduced-form and first-stage entries are percentage-point changes associated with five additional historical births per 1,000 people or a one-percentage-point increase in the child-stock instrument. 2SLS entries are employment-share percentage points per one-point larger observed young-adult share. All rows retain country and year fixed effects and use the same 3,650 observations from 147 countries. “Full” includes all baseline controls; “None” omits them. Parentheses contain country-clustered standard errors, estimated separately for each regression. Ratios use unrounded estimates.

Table 30 adds the capital and technology block sequentially while holding the 3,650 observations fixed. The coefficient barely moves: for industry’s share of total employment it falls from 0.788 to 0.746, and within non-agriculture from 1.059 to 1.027. The remaining rows ask whether conditioning on measured birth-era health attenuates the estimate. It does not: controlling for birth-window crude death rates, under-five mortality, adult mortality, or both child and adult mortality leaves the coefficient at least as large as in the baseline.

Table 30: Coefficient stability on the identical WDI sample

Controls	Industry employment share		K-P F
	Total employment	Non-ag. employment	
Development, trade, labor supply, and human capital	0.788 (0.321)	1.059 (0.355)	75.7
+ capital-output ratio and investment	0.755 (0.314)	1.035 (0.350)	78.0
+ TFP (baseline)	0.746 (0.313)	1.027 (0.350)	87.1
Baseline + birth-era crude death rate	0.831 (0.272)	1.061 (0.322)	114.3
Baseline + birth-era under-five mortality	0.828 (0.272)	1.067 (0.325)	124.1
Baseline + birth-era adult mortality	0.877 (0.271)	1.063 (0.320)	110.9
Baseline + child and adult mortality	0.856 (0.260)	1.047 (0.316)	129.1

Notes: Entries are 2SLS coefficients on the population share aged 25–34, instrumented with the mean birth rate 25–34 years earlier. All rows use the same 3,650 observations from 147 countries and include country and year fixed effects, a quadratic in log GDP per capita, female labor-force participation, mean years of schooling, and imports relative to GDP. The capital block adds the log capital-output ratio and the investment share; the baseline additionally includes log TFP. Mortality variables are averaged over the same 25–34-year birth window as the instrument. Adult mortality is the probability of dying between ages 15 and 60. Standard errors clustered by country are in parentheses.

Nor does a single country drive the result. Re-estimating the baseline after deleting each of the 147 countries in turn leaves every coefficient positive. The total-employment estimate ranges from 0.569 to 0.871, around the full-sample value of 0.746; the non-agricultural estimate ranges from 0.839 to 1.114, around 1.027.

G.4 Economic conditions at birth and regional shocks

Historical economic conditions may influence both fertility and subsequent industrialization. To account directly for these conditions, I augment equation (3) with income measured during the years in

which the young-adult cohorts were born. The controls are mean log GDP per capita over $t-34, \dots, t-25$ and its square; a further specification also includes income growth over this birth window, measured by the slope of log GDP per capita on calendar year. All ten annual income observations must be available. The resulting sample contains 3,353 country-year observations from the same 147 countries.

Table 31 compares these estimates with the baseline on the identical sample. For industry's share of total employment, the coefficient changes from 0.664 to 0.658 after including birth-window income and its square, and to 0.681 after also including income growth. The corresponding estimates for industry's share of non-agricultural employment are 0.923, 0.866, and 0.911. Both augmented specifications retain positive, statistically significant demographic coefficients and strong first stages.

A separate specification allows countries in the same geographic region to experience common annual shocks. It retains country fixed effects and replaces common year effects with World Bank region-year fixed effects. Thus, a regional expansion or contraction shared across countries in a given year is absorbed before estimating the demographic relationship. On the full sample, the coefficients are 0.820 for industry's share of total employment and 1.171 for its share of non-agricultural employment. The relationship therefore remains positive when accounting directly for economic conditions at birth and for annual developments shared within regions.

Table 31: Birth-era economic conditions and regional annual shocks

Specification	Industry employment share		Observations	K-P F
	Total employment	Non-ag. employment		
Baseline: historical-income sample	0.664 (0.321)	0.923 (0.338)	3,353	78.5
Birth-window income and its square	0.658 (0.306)	0.866 (0.333)	3,353	81.8
+ birth-window income growth	0.681 (0.313)	0.911 (0.337)	3,353	81.6
Baseline: full sample	0.746 (0.313)	1.027 (0.350)	3,650	87.1
Region-year fixed effects	0.820 (0.325)	1.171 (0.382)	3,650	83.6

Notes: Entries are 2SLS coefficients on the share of the total population aged 25–34, instrumented with the mean crude birth rate 25–34 years earlier. All specifications retain the baseline economic controls and country fixed effects. The historical-income specifications add mean log GDP per capita over $t-34, \dots, t-25$ and the square of that mean; the growth control is the slope of log GDP per capita on calendar year within the same window. All ten annual income observations must be available. The regional specification replaces common year effects with fixed World Bank region-year effects. Each baseline uses the same observations as the specifications immediately below it. All rows cover 147 countries during 1991–2016. Standard errors clustered by country are in parentheses. K-P F is the cluster-robust Kleibergen–Paap first-stage statistic.

G.5 The relationship is present for manufacturing in GGDC

WDI's industry category combines manufacturing with mining, utilities, and construction. GGDC separates manufacturing. I therefore estimate a companion panel using manufacturing's employment share as the outcome. Table 32 shows a positive relationship directly for manufacturing: the 2SLS coefficient is 0.926 for manufacturing's share of total employment and 1.266 for its share within non-agriculture. Both estimates remain statistically significant under country-clustered, wild-bootstrap, and Anderson–Rubin inference.

Table 32: Young-adult cohort size and manufacturing employment shares (GGDC)

	Manufacturing / total employment	Manufacturing / non-ag. employment
OLS	0.783 (0.204)	0.661 (0.179)
2SLS: lagged birth rate	0.926 (0.331)	1.266 (0.370)
Kleibergen–Paap F	45.4	45.4
Wild-bootstrap p	0.012	0.003
Anderson–Rubin p	0.015	0.008
Observations	1,364	1,364
Countries	56	56

Notes: The dependent variables are GGDC manufacturing employment as a share of total employment and of non-agricultural employment. The key regressor is the share of the total population aged 25–34. The 2SLS specifications instrument it with the mean crude birth rate 25–34 years earlier. All specifications include country and year fixed effects, log GDP per capita, female labor-force participation, mean years of schooling, imports relative to GDP, the log capital–output ratio, the investment share, and log TFP. The sample spans 1991–2016. Standard errors clustered by country are in parentheses. Wild-bootstrap and Anderson–Rubin p -values use 9,999 replications clustered by country.

G.6 The manufacturing gradient is also visible within countries

The cross-country panel is complemented by a comparison across subnational regions using IPUMS International census and labor-force-survey samples. For region r in country c and year t , I estimate

$$s_{rct}^M = \beta \text{young}_{rct} + \alpha_r + \delta_{ct} + \varepsilon_{rct}, \quad (77)$$

where s_{rct}^M is manufacturing employment as a share of employment records with a valid industry classification, α_r are region fixed effects, and δ_{ct} are country–year fixed effects. The latter compare regions observed in the same country and year. The sample contains 8,316 region–year observations from 1,463 regions in 72 countries.

Table 33 reports the estimates. A one-percentage-point larger regional young-adult share is associated with a 0.381-point higher manufacturing share of industry-coded employment. The estimate is 0.331 when manufacturing is measured within non-agricultural employment. The corresponding coefficients for construction and for utilities plus mining are small and statistically indistinguishable from zero. Region and country–year fixed effects show that the pattern is also present across regions observed in the same national economy at the same time.

G.7 Trade openness and demographic comparative advantage

The model counterfactuals show that limiting the responsiveness of foreign demand for manufactures attenuates the manufacturing response to demographic change. Trade gives an economy scope to specialize in sectors that use the abilities supplied relatively abundantly by its population. Since manufacturing relies more heavily on abilities that peak earlier in life, access to foreign demand gives a younger population greater scope to expand manufacturing. I therefore examine whether the employment response associated with a larger young-adult share is stronger in more open economies.

Table 33: Young-adult shares and sectoral employment across subnational regions

	Manufacturing	Manufacturing / non-ag.	Broad industry	Construction	Utilities + mining
Young-adult share (25–34)	0.381 (0.087)	0.331 (0.113)	0.509 (0.151)	0.089 (0.071)	0.045 (0.046)
Observations	8,316	8,316	8,316	8,316	8,316
Regions	1,463	1,463	1,463	1,463	1,463
Countries	72	72	72	72	72
Region FE	Yes	Yes	Yes	Yes	Yes
Country–year FE	Yes	Yes	Yes	Yes	Yes

Notes: The unit of observation is the region–year. The IPUMS sample spans 1960–2020 and includes censuses and labor-force surveys. Regressions are unweighted. The regressor is the share of the total population aged 25–34, measured in percentage points. Except in the second column, outcomes are shares of employment with a valid industry classification; broad industry combines manufacturing, construction, utilities, mining, and residual industrial activities. Outcomes are also measured in percentage points. Standard errors clustered by country are in parentheses. For the manufacturing outcome in column 1, the null-imposed wild-cluster-bootstrap p -value is below 0.001.

Let W_{it} denote exports plus imports of goods and services divided by GDP. I extend equation (3) to

$$s_{it} = \beta \text{Young}_{it} + \delta \text{Young}_{it} W_{it}^c + \eta W_{it}^c + \lambda_1 \ln(\text{GDPpc}_{it}) + \lambda_2 [\ln(\text{GDPpc}_{it})]^2 + \mathbf{X}_{it}' \theta + \alpha_i + \gamma_t + \varepsilon_{it}, \quad (78)$$

where $W_{it}^c = W_{it} - \widetilde{W}$ and $\widetilde{W}$ is the median trade/GDP ratio in the main WDI sample, 67.74 percent. The same reference is used in every column. Hence β is the demographic coefficient at this openness level, while δ measures how that coefficient changes with openness. The controls $\mathbf{X}_{it}$ are those in the main specification, including imports relative to GDP. I instrument both Young_{it} and $\text{Young}_{it} W_{it}^c$ with the cohort instrument Z_{it} and its product $Z_{it} W_{it}^c$. Trade/GDP and the remaining controls enter as exogenous regressors. The identifying restriction requires both excluded instruments to be unrelated to the remaining determinants of sectoral employment, conditional on the included controls and fixed effects.

Table 34 reports four specifications. The first two use current trade/GDP and, respectively, the birth-rate and surviving-child-stock instruments. The third uses trade/GDP five years earlier with the birth-rate instrument. The fourth uses manufacturing’s share of total employment in the GGDC data with current trade/GDP and the birth-rate instrument. All four include the quadratic in log GDP per capita, the remaining baseline controls, and country and year fixed effects.

The interaction is positive and statistically significant at the one-percent level in every column. To express its magnitude in employment-share units, the table reports $0.20\widehat{\delta}$: the additional response to a one-percentage-point larger young-adult share associated with a twenty-percentage-point higher trade/GDP ratio. This difference is 0.479 percentage points with the birth-rate instrument and 0.375 with the child-stock instrument. Using trade five years earlier gives 0.439 points, and isolating manufacturing gives 0.357 points. These estimates show that the demographic employment relationship is stronger in more open economies, consistent with trade allowing greater specialization according to demographic comparative advantage.

Table 34: Trade openness and the demographic employment relationship

	Industry			Manufacturing
	(1) Birth rate Current	(2) Child stock Current	(3) Birth rate Five-year lag	(4) Birth rate Current
Cohort instrument				
Trade/GDP				
Young-adult share	0.531 (0.325)	0.514*** (0.190)	0.509 (0.319)	0.375 (0.426)
Young-adult share $\times$ openness	2.397*** (0.644)	1.876*** (0.385)	2.193*** (0.673)	1.787*** (0.489)
Openness	-0.367*** (0.097)	-0.289*** (0.059)	-0.329*** (0.102)	-0.218*** (0.076)
Log GDP per capita	0.166*** (0.060)	0.162*** (0.061)	0.215*** (0.058)	0.105*** (0.033)
(Log GDP per capita) ²	-0.008** (0.003)	-0.007** (0.003)	-0.010*** (0.003)	-0.006** (0.002)
Female labor-force participation	-0.199*** (0.057)	-0.198*** (0.056)	-0.202*** (0.057)	-0.181*** (0.053)
Mean years of schooling, ages 25+	-0.004 (0.004)	-0.004 (0.004)	-0.005 (0.004)	-0.005 (0.004)
Imports / GDP	0.024 (0.031)	0.017 (0.030)	0.022 (0.020)	-0.057 (0.056)
Log capital-output ratio	0.017 (0.014)	0.016 (0.014)	0.007 (0.015)	-0.026 (0.016)
Investment share	0.071*** (0.022)	0.076*** (0.022)	0.065*** (0.025)	0.027 (0.039)
Log TFP	0.013 (0.020)	0.005 (0.019)	-0.012 (0.021)	-0.041* (0.021)
Additional response at 20pp higher openness	0.479 (0.129)	0.375 (0.077)	0.439 (0.135)	0.357 (0.098)
95% confidence interval	[0.225, 0.734]	[0.223, 0.527]	[0.172, 0.705]	[0.161, 0.553]
Country and year fixed effects	Yes	Yes	Yes	Yes
Observations	3,650	3,650	3,474	1,364
Countries	147	147	144	56
Conditional first-stage F : young share	38.3	109.4	33.3	32.6
Conditional first-stage F : interaction	23.0	24.9	18.5	30.1

Notes: The dependent variable is the sector's share of total employment. Industry in columns 1–3 includes manufacturing, mining, utilities, and construction; column 4 uses GGDC manufacturing employment. The regressor of interest is the share of the total population aged 25–34. Openness is exports plus imports of goods and services divided by GDP, centered at the main WDI sample median of 67.74 percent in every column. In column 3 it is measured five calendar years earlier. Young-adult share and its interaction with openness are instrumented with the indicated cohort instrument and its product with centered openness. The young-adult coefficient is evaluated at the reference openness level; the openness coefficient is evaluated at a young-adult share of zero. Employment, population, participation, and trade shares enter as fractions. The additional response is 0.20 δ , in employment percentage points per one-percentage-point larger young-adult share. All specifications are unweighted, span 1991–2016, and include the controls shown. Standard errors clustered by country are in parentheses. Confidence intervals use the full coefficient covariance and country-cluster degrees of freedom. Conditional first-stage statistics are the Sanderson–Windmeijer F statistics for the two endogenous regressors. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

G.7.1 Control sensitivity and reduced forms

In the trade-interaction specification, I remove the contemporaneous economic controls listed in Appendix G.3 but retain openness itself, together with the young-adult share and its interaction with openness. The instruments remain the cohort instrument and its product with centered openness, as in equation (78). These comparisons use identical observations with and without the controls: 3,650 country–years from 147 countries for WDI industry and 1,364 country–years from 56 countries for GGDC manufacturing, over 1991–2016. Both outcomes are shares of total employment, and all regressions retain country and year fixed effects.

Table 35 reports the estimates. The openness interaction also remains positive in both WDI in-

dustry and GGDC manufacturing. Without the additional controls, twenty percentage points higher trade/GDP is associated with an additional industry-employment-share response of 0.691 points using the birth-rate instrument and 0.578 points using the child-stock instrument; the corresponding manufacturing estimates are 0.452 and 0.401 points.

Table 35: Trade-interaction IV estimates with and without contemporaneous economic controls

Outcome	Instrument	Full controls	No additional controls
<i>Panel A. Trade specification: coefficient at reference openness</i>			
WDI industry	Birth rate	0.531 (0.325)	0.888 (0.410)
	Child stock	0.514 (0.190)	0.773 (0.252)
GGDC manufacturing	Birth rate	0.375 (0.426)	1.159 (0.404)
	Child stock	0.326 (0.207)	0.941 (0.286)
<i>Panel B. Additional response at 20pp higher trade/GDP</i>			
WDI industry	Birth rate	0.479 (0.129)	0.691 (0.196)
	Child stock	0.375 (0.077)	0.578 (0.142)
GGDC manufacturing	Birth rate	0.357 (0.098)	0.452 (0.106)
	Child stock	0.269 (0.093)	0.401 (0.069)

Notes: Entries are employment-share percentage points per one-percentage-point larger young-adult share. Panels A and B report β and 0.20δ from equation (78); reference openness is 67.74 percent trade/GDP. All regressions retain country and year fixed effects; the trade specification also retains openness. Full controls are those in equation (3). Samples cover 1991–2016 and are identical across columns: WDI, 3,650 observations and 147 countries; GGDC, 1,364 observations and 56 countries. Country-clustered standard errors are in parentheses. In the trade specification without additional controls, Kleibergen–Paap rk Wald F statistics are 9.5 and 7.3 for WDI and 14.7 and 5.2 for GGDC, respectively using birth rates and child stocks; inference here is conventional rather than weak-instrument-robust.

Reduced forms with trade interactions. For the specifications in Appendix G.7, the reduced form includes both excluded instruments:

$$s_{it}^j = r_j Z_{it} + q_j Z_{it} W_{it}^c + h_j W_{it}^c + \mathbf{C}_{it}' \psi_j + \alpha_i + \gamma_t + u_{it}, \quad (79)$$

where $W_{it}^c = W_{it} - \widetilde{W}$, with $\widetilde{W} \simeq 0.6774$, retaining the common openness reference used in the IV regressions. The coefficient r_j describes the instrument–employment relationship at this reference; $0.20q_j$ measures how much larger that relationship is at twenty percentage points higher trade/GDP. As above, $\mathbf{C}_{it}$ contains all baseline controls in the full specification and is omitted in the alternative. Openness remains in both.

Table 36 uses the same instrument increments as Table 29. For example, in GGDC without additional controls, five additional historical births per 1,000 are associated with a 1.027-percentage-point higher manufacturing employment share at the reference openness level. At twenty percentage points higher trade/GDP, this association is 0.252 points larger, or approximately 1.28 points in total. The reduced-form interactions are positive in every specification; the GGDC child-stock interaction with full controls is imprecisely estimated.

These reduced forms describe how the employment relationship with the historical instruments

varies with openness. They do not measure the employment response to a one-point increase in the current young-adult share. The trade IV specification has two endogenous regressors, and recovering its coefficients requires both first-stage equations. In particular, the reduced-form interaction cannot be converted into the IV interaction by dividing by 1.483 or 0.582.

Table 36: Reduced-form employment relationships and trade openness

Outcome	Instrument	Controls	At reference openness	Additional response at 20pp higher openness
WDI industry	Birth rate	Full	0.934 (0.449)	0.248 (0.044)
		None	0.692 (0.406)	0.310 (0.042)
	Child stock	Full	0.332 (0.107)	0.060 (0.013)
		None	0.309 (0.111)	0.080 (0.016)
GGDC manufacturing	Birth rate	Full	0.320 (0.399)	0.181 (0.080)
		None	1.027 (0.440)	0.252 (0.071)
	Child stock	Full	0.123 (0.094)	0.042 (0.028)
		None	0.350 (0.122)	0.059 (0.028)

Notes: Entries are employment-share percentage points for the instrument increments in Table 29. The columns report scaled estimates of r_j and $0.20q_j$ in equation (79), evaluated at the common reference of 67.74 percent trade/GDP. “Full” retains all baseline controls; “None” omits these controls but retains openness and country and year fixed effects. Samples match Table 35. Standard errors clustered by country are in parentheses. These are reduced-form associations, not effects of changes in trade policy.

H Robustness of the industrial-employment peak relationships

This appendix tests the peak relationships across sector definitions, young-adult age windows, pre-peak capital and productivity controls, demographic measures constructed without reference to country-specific peak dates, peak definitions, employment-share denominators, and inference procedures.

H.1 Full peak-regression estimates

Table 37 reports every non-fixed-effect coefficient from the two main peak regressions.

H.2 Smoothed WDI peaks

Raw annual maxima can reflect transitory measurement error. I therefore reconstruct the peak as the maximum of a centered five-year moving average of the WDI industry employment share and re-estimate both peak equations using the strictly pre-peak windows and controls from the main text.

The association is preserved under smoothing. Across the three windows, a one-percentage-point larger pre-peak young-adult share is associated with a 1.2–1.6 percentage point higher smoothed industry peak and 0.17–0.23 log points higher GDP per capita at that peak.

Table 37: Full estimates for pre-peak age structure and industrial-employment peaks

	5-year window	10-year window	15-year window
<i>Panel A. Peak industry employment share</i>			
Young-adult share (25–34)	1.106 (0.288)	1.438 (0.410)	1.650 (0.485)
Pre-peak GDP per capita growth	0.045 (0.095)	0.019 (0.059)	−0.012 (0.025)
Average log GDP per capita	0.080 (0.164)	0.144 (0.170)	0.279 (0.135)
(Average log GDP per capita) ²	−0.0028 (0.0091)	−0.0068 (0.0095)	−0.0150 (0.0075)
Average imports / GDP	0.022 (0.046)	0.003 (0.053)	0.021 (0.042)
Peak year	−0.0056 (0.0031)	−0.0054 (0.0030)	−0.0059 (0.0029)
Adjusted R^2	0.525	0.513	0.517
Observations	70	68	62
<i>Panel B. Log GDP per capita at the peak</i>			
Young-adult share (25–34)	0.177 (0.025)	0.220 (0.029)	0.261 (0.041)
Pre-peak GDP per capita growth	−0.536 (0.997)	0.387 (0.476)	0.436 (0.327)
Average imports / GDP	0.407 (0.592)	0.230 (0.508)	−0.147 (0.439)
Peak year	0.009 (0.030)	−0.010 (0.031)	−0.055 (0.036)
Adjusted R^2	0.417	0.535	0.610
Observations	70	68	62
Peak-decade fixed effects	Yes	Yes	Yes

Notes: The table reports all non-fixed-effect coefficients from equations (4) and (5). Within each cell, the coefficient is followed by its heteroskedasticity-robust standard error in parentheses. The young-adult coefficients and standard errors are expressed for a one-percentage-point increase; all other coefficients are in the native units of their regressors. The specifications include peak-decade fixed effects and are estimated without a constant.

H.3 Peak-date-independent demographic measures from GDP-band averages

A remaining concern is that pre-peak demographic measures are dated relative to each country’s peak year. To address this, I construct demographic measures that do not condition on t_i^{peak} . Specifically, I average each country’s young-cohort share across all years in which real GDP per capita falls within a common income band, and I use this country-level average as the demographic regressor. These averages can include years before or after the peak, but their construction makes no reference to the country-specific peak date.

Table 39 shows that the relationship between demographics and both peak industry employment and log GDP per capita at the peak remains positive and statistically significant under these peak-date-independent measures.

Table 38: Pre-peak young-adult shares and smoothed WDI industry peaks

	5 years	10 years	15 years
	(1)	(2)	(3)
<i>Panel A: Peak industry employment share</i>			
Young-adult share (average)	1.191*** (0.230)	1.298*** (0.368)	1.562*** (0.474)
Observations	54	50	48
<i>Panel B: Log GDP per capita at the industry peak</i>			
Young-adult share (average)	0.171*** (0.025)	0.206*** (0.025)	0.229*** (0.030)
Observations	54	50	48
Peak year and peak-decade FE	Yes	Yes	Yes
GDP and import controls	Yes	Yes	Yes

Notes: The peak is the interior maximum of a centered five-year moving average of WDI industry employment. The demographic and control windows run from $t_i^{\text{peak}} - L$ through $t_i^{\text{peak}} - 1$. Controls and estimation otherwise match Tables 4 and 4. Young-adult coefficients and standard errors are reported per one-percentage-point increase in the young-adult share, with peak employment shares in percentage points and income in log points. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

H.4 Alternative young-adult window and pre-peak aggregate controls

Table 40 changes the age window and then adds two pre-peak aggregate controls. Panel A defines young adults as ages 20–34 while leaving the peak definitions, controls, and country samples unchanged. Panel B returns to ages 25–34 and, on the common PWT-covered sample, adds trailing measures of capital deepening and TFP growth that end in the year before the peak.

The age-window result preserves the sign, precision, and increasing horizon profile of the main estimates. On the PWT-covered sample, adding pre-peak capital deepening and TFP growth does not attenuate either focal relationship: the peak-height estimate rises from 1.680 to 1.979 percentage points per percentage point, while the corresponding peak-income estimate rises from 0.184 to 0.224 log points. This establishes conditional stability with respect to the two aggregate controls rather than decomposing their contribution to the relationship.

Table 39: Peak-date-independent demographic measures from GDP-band averages

	(1) GDP band 1	(2) GDP band 2
<i>Panel A: Peak industrial employment share (raw peak)</i>		
Young-adult (25–34) share (GDP-band average)	2.198***	2.060***
	(0.666)	(0.650)
GDP per capita growth (10-year)	-0.040	-0.047
	(0.033)	(0.034)
Log GDP per capita (10-year average)	0.184	0.218
	(0.140)	(0.179)
(Average log GDP per capita) ²	-0.008	-0.010
	(0.008)	(0.010)
Peak year	-0.005*	-0.006**
	(0.003)	(0.003)
Decade FE	Yes	Yes
Observations	84	80
Adj. R^2	0.506	0.496
<i>Panel B: Peak industrial employment share / non-agricultural employment (raw peak)</i>		
Young-adult (25–34) share (GDP-band average)	2.083***	2.046***
	(0.632)	(0.575)
GDP per capita growth (10-year)	-0.010	-0.017
	(0.034)	(0.037)
Log GDP per capita (10-year average)	0.114	0.066
	(0.113)	(0.153)
(Average log GDP per capita) ²	-0.007	-0.004
	(0.007)	(0.009)
Peak year	0.003	0.002
	(0.004)	(0.004)
Decade FE	Yes	Yes
Observations	78	71
Adj. R^2	0.083	0.104
<i>Panel C: Log GDP per capita at peak (raw peak year)</i>		
Young-adult (25–34) share (GDP-band average)	0.299***	0.206***
	(0.091)	(0.062)
GDP per capita growth (10-year)	0.260	0.208
	(0.429)	(0.443)
Peak year	0.025	0.015
	(0.035)	(0.037)
Decade FE	Yes	Yes
Observations	84	80
Adj. R^2	0.320	0.182

Notes: “Young-adult (25–34) share (GDP-band average)” averages each country’s young-adult share over all years in which real GDP per capita falls within a common income band (band 1: 1,000–20,000; band 2: 2,000–15,000, in the units of real GDP per capita used in the data). These measures are computed without reference to the country-specific peak year. All specifications include decade fixed effects (absorbed) and are estimated without a constant. Young-adult coefficients and standard errors are reported per one-percentage-point increase in the young-adult share, with peak employment shares in percentage points and income in log points. All other coefficients are in the native units of their regressors. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 40: Alternative age window and pre-peak capital and productivity controls

	Peak industry share	Log GDPpc at peak	N
<i>Panel A: Young-adult share defined as ages 20–34</i>			
5-year pre-peak average	0.652** (0.245)	0.109*** (0.027)	70
10-year pre-peak average	0.809*** (0.293)	0.144*** (0.023)	68
15-year pre-peak average	1.166*** (0.357)	0.193*** (0.029)	62
<i>Panel B: Ages 25–34 with capital and TFP controls, ten-year window</i>			
Common-sample baseline	1.680*** (0.456)	0.184*** (0.028)	47
Add capital deepening and TFP growth	1.979*** (0.577)	0.224*** (0.053)	47
Peak-year control; peak-decade FE	Yes	Yes	
Standard pre-peak controls	Yes	Yes	

Notes: Panel A reproduces the main WDI peak equations after replacing the population share aged 25–34 with the share aged 20–34; samples are identical to the corresponding main-text specifications. Panel B uses the focal ten-year specification and restricts both rows to countries with non-missing Penn World Table measures. Capital deepening is the trailing ten-year change in log capital per worker, and TFP growth is the trailing ten-year cumulative Solow residual; both measures end in the year before the peak. The common-sample baseline includes all standard controls but omits these two additional variables. Both equations include pre-peak GDP growth, average import exposure, a linear peak-year control, and peak-decade fixed effects; the peak-height equation additionally includes average log GDP per capita and its square. Young-adult coefficients and standard errors are reported per one-percentage-point increase in the young-adult share, with peak employment shares in percentage points and income in log points. Regressions are estimated without a constant. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

H.5 Manufacturing-specific peaks in GGDC

The WDI outcome in the main text is broad industry. Table 41 repeats the exercise using GGDC manufacturing employment. To align the timing and controls with the main specification, all averages use years $t_i^{\text{peak}} - L$ through $t_i^{\text{peak}} - 1$, excluding the peak year, and the regressions include the peak year, peak-decade fixed effects, and GDP growth. The rows with an import control also include average import exposure. The peak-height regressions additionally include average log GDP per capita and the square of that average.

Table 41: Pre-peak young-adult shares and manufacturing peaks in GGDC

	5 years	10 years	15 years
	(1)	(2)	(3)
<i>Panel A: Peak manufacturing employment share</i>			
Young-adult share (average)	1.847*** (0.606)	2.126** (0.791)	2.804** (1.008)
Young-adult share (with import control)	2.639*** (0.706)	3.214*** (1.036)	3.632** (1.296)
<i>Panel B: Log GDP per capita at the manufacturing peak</i>			
Young-adult share (average)	0.281** (0.108)	0.256*** (0.092)	0.354*** (0.100)
Young-adult share (with import control)	0.261* (0.133)	0.271** (0.114)	0.409*** (0.107)
Observations	38	36	32
Peak-year control; peak-decade FE	Yes	Yes	Yes
GDP controls	Yes	Yes	Yes

Notes: The unit of observation is a country. Panel A uses manufacturing employment as a share of total employment; Panel B uses log GDP per capita in the same peak year. The key regressor is the population share aged 25–34 averaged over the strictly pre-peak window. GDP growth is measured between the first and last years of that window. Rows labeled “with import control” add average import exposure over the same years. Panel A also controls for average log GDP per capita and the square of that average. Regressions absorb peak-decade fixed effects, include the peak year, and are estimated without a constant. Young-adult coefficients and standard errors are reported per one-percentage-point increase in the young-adult share, with peak employment shares in percentage points and income in log points. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

The manufacturing-specific estimates reinforce the WDI result. Without the import control, a one-percentage-point larger strictly pre-peak young-adult share is associated with a 1.8–2.8 percentage point higher manufacturing employment peak; adding it yields 2.6–3.6 points. The corresponding

income-at-peak coefficients are positive under both specifications.

As an intermediate GGDC aggregate—broader than manufacturing but not an exact match to WDI industry because mining is excluded—I construct M^+ as the sum of manufacturing, construction, and utilities employment. Table 42 dates each country’s peak using this measure and repeats the strictly pre-peak specifications. The young-adult coefficient is positive and precise for peak height at every horizon and positive for income at the peak under both control sets.

Table 42: Pre-peak young-adult shares and broader-sector peaks in GGDC

	5 years	10 years	15 years
	(1)	(2)	(3)
<i>Panel A: Peak M^+ employment share</i>			
Young-adult share (average)	2.009*** (0.701)	2.518*** (0.836)	3.402*** (1.133)
Young-adult share (with import control)	2.216** (0.822)	3.103*** (0.884)	3.897*** (1.168)
<i>Panel B: Log GDP per capita at the M^+ peak</i>			
Young-adult share (average)	0.261** (0.120)	0.267*** (0.092)	0.394*** (0.100)
Young-adult share (with import control)	0.256* (0.140)	0.294*** (0.104)	0.441*** (0.102)
Observations	41	36	32
Peak-year control; peak-decade FE	Yes	Yes	Yes
Standard pre-peak controls	Yes	Yes	Yes

Notes: M^+ is manufacturing plus construction and utilities employment as a share of total employment. Each country’s peak is dated using M^+ , and component shares discussed in the text are evaluated in that year. The GGDC panel spans 1960–2018; countries whose maximum occurs in 2018 are excluded, and ties are assigned to the earliest maximum. Demographic and control windows are strictly pre-peak. Both equations include pre-peak GDP growth, a linear peak-year control, and peak-decade fixed effects; the peak-height equation additionally includes average log GDP per capita and its square. Rows labeled “with import control” add average import exposure. Young-adult coefficients and standard errors are reported per one-percentage-point increase in the young-adult share, with peak employment shares in percentage points and income in log points. Regressions are estimated without a constant. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

The broader-sector result is driven overwhelmingly by manufacturing. In the focal ten-year specification with import exposure, the 3.103-point M^+ coefficient decomposes exactly into 2.839 points for manufacturing (s.e. 0.597), 0.187 for construction (s.e. 0.497), and 0.077 for utilities (s.e. 0.086), all

evaluated at the M^+ peak. Manufacturing therefore accounts for 91.5 percent of the coefficient; construction and utilities are individually imprecise. This is an accounting decomposition of the regression coefficient, not a comparison of the components at their own sector-specific peaks.

H.6 Alternative denominator and influence-robust inference

Table 43: Peak height measured relative to non-agricultural employment

	(1)	(2)
	5-year pre-peak average	10-year pre-peak average
Young-adult (25–34) share (average)	1.228*** (0.389)	1.393** (0.548)
Peak year and peak-decade FE	Yes	Yes
GDP and import controls	Yes	Yes
Observations	65	64

Notes: The dependent variable is the raw peak WDI industry employment share measured relative to non-agricultural employment. Demographic and control averages exclude the peak year and otherwise match Table 4. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 43 replaces the baseline denominator with non-agricultural employment. The coefficient remains positive and statistically precise in both windows.

The focal raw-peak estimate is also robust to more conservative inference and to individual-country influence. For the ten-year specification, HC3 standard errors give a coefficient of 1.438 with a standard error of 0.523 ($p = 0.008$). Across the 68 exact leave-one-country-out refits, the coefficient ranges from 1.215 to 1.693, and the smallest absolute t -statistic is 2.96.

I Labor force per capita and the working-age share

Labor supply per capita moves almost one-for-one with the working-age share: regressing labor force per capita on the share of the population of working age (20–64), with country fixed effects, gives a slope between 0.76 and 0.83 (Figure 35). Because sectoral employment shares are defined over employment while the paper’s outcomes are per capita, shifts in the labor-force-to-population ratio affect output per capita even holding sectoral productivities fixed. This motivates the model’s partition of cohorts into workers, who supply a unit of labor inelastically, and retirees.




Figure 35: Labor force per capita and the working-age share. Notes: scatter observations are net of country fixed effects.

J Model: derivations and proofs

This appendix provides the derivations and equilibrium conditions supporting Section 3. Appendix J.1 explains sectoral choice and the calculation of employment and effective labor, including costly agricultural exit. The proofs of Propositions 3.1 and 3.2 are in Appendices J.1.2 and J.1.3, respectively. Appendix J.2 completes the baseline trade and equilibrium conditions. Appendix J.3 develops durable accumulation and the dynasty's joint consumption and stock choices.

J.1 Worker choices and sectoral labor supply

This section derives how permanent abilities and schooling determine sectoral choices, employment, and effective labor. I first consider costless sectoral moves, the building block for the model with costly exit from agriculture in Appendix J.1.4. Throughout, sums over ability indices run over $\{P, F, C\}$ unless otherwise indicated.

J.1.1 Permanent abilities, schooling, and cohort distributions

The permanent component of log ability k is $\nu_k + \alpha_k s$. The residual draws are jointly Gaussian, with means $\mu_{0,k}$ and covariances $\Sigma_{0,kr}$. Within birth cohort c , schooling satisfies $s \sim \mathcal{N}(\bar{s}_c, \sigma_{s,c}^2)$ and is independent of the residual draws, as in (9). Because a linear combination of jointly Gaussian variables is Gaussian, the three permanent log abilities are jointly Gaussian, with moments

$$\mu_{c,k} = \mu_{0,k} + \alpha_k \bar{s}_c, \quad \Sigma_{c,kr} = \Sigma_{0,kr} + \alpha_k \alpha_r \sigma_{s,c}^2. \quad (80)$$

Here $\mu_{c,k}$ is the cohort mean of permanent log ability k , and $\Sigma_{c,kr}$ is its covariance with permanent log ability r . For $k = r$, the second expression gives the variance. Independence between schooling

and the residual draws eliminates the cross-covariance terms.

A higher cohort mean of schooling shifts mean log ability k by α_k times the increase in schooling. Greater dispersion in schooling also changes the variances and covariances of permanent abilities, because the same schooling draw enters all three components. Schooling can therefore affect both overall productivity and relative strengths; its effect on relative strengths depends on differences in the α_k s.

These permanent components are fixed at labor-market entry. Within a given cohort, their distribution does not change as workers age. Aging changes the common profiles $\rho_k(a)$ and hence the sectoral efficiency profiles $\tilde{z}_j(a)$, not workers' permanent ability or schooling draws. Different cohorts can nevertheless have different permanent-ability distributions because their schooling distributions differ.

J.1.2 Sectoral choice and relative abilities

When sectoral moves are costless, workers choose the sector offering the highest earnings. Since ability weights sum to one in every sector, a proportional increase in all three abilities raises earnings in every sector by the same proportion. Sectoral choice therefore depends on relative strengths, rather than overall ability. This reduction does not require the Gaussian distributional assumptions.

As in the main text, define the worker's average permanent log ability and the deviations from that average:

$$\bar{m} = \frac{1}{3} \sum_k (\nu_k + \alpha_k s), \quad \hat{h}_k = \nu_k + \alpha_k s - \bar{m}.$$

The component $\hat{h}_k$ measures strength in ability k relative to the worker's own average, not relative to other workers. The three deviations sum to zero, so $\hat{h}_C = -\hat{h}_P - \hat{h}_F$. A worker's permanent ability mix can therefore be represented by the fixed point $(\hat{h}_P, \hat{h}_F)$.

Proof of Proposition 3.1. Substituting $\nu_k + \alpha_k s = \bar{m} + \hat{h}_k$ into log earnings and using $\sum_k \theta_{j,k} = 1$ gives

$$\log(w_{j,t} z_j(a)) = \log w_{j,t} + \log \tilde{z}_j(a) + \bar{m} + \sum_k \theta_{j,k} \hat{h}_k. \quad (81)$$

The term $\bar{m}$ is common across sectors and cancels from every comparison. Sector j pays at least as much as sector ℓ if and only if

$$\sum_k (\theta_{j,k} - \theta_{\ell,k}) \hat{h}_k \geq \log \frac{w_{\ell,t} \tilde{z}_\ell(a)}{w_{j,t} \tilde{z}_j(a)}. \quad (82)$$

Since $\hat{h}_C = -\hat{h}_P - \hat{h}_F$, this comparison depends only on $(\hat{h}_P, \hat{h}_F)$, proving part (i).

For distinct sectoral ability weights, each inequality describes one side of a straight line in these two coordinates. Sector j is optimal when all its comparisons hold, giving the region $R_j(a, t)$ in (16). This

intersection of at most three closed half-planes is convex and may be empty or unbounded, proving part (ii). If two sectors have identical ability weights, their comparison holds for every ability mix or for none, depending on which has the higher product $w_{j,t}\tilde{z}_j(a)$. $\square$

The worker's point remains fixed over their career. The pairwise indifference lines have directions determined by differences in sectoral ability weights. Age profiles and relative wages enter the right-hand side of (82): they move each nontrivial line parallel to itself, rather than changing its slope. A worker switches sectors when these changes make another sector the highest-paying choice.

J.1.3 Employment and productive contributions

Cohort employment shares. A cohort's employment share in each sector is the fraction of its workers whose ability mixes lie in that sector's choice region. The Gaussian assumptions on endowments and schooling imply that $\hat{h}_P$ and $\hat{h}_F$ are jointly Gaussian. Let $\mu_{\hat{h},c}$ and $\Sigma_{\hat{h},c}$ denote the mean vector and covariance matrix of this pair for cohort c .

With a nonsingular ability-mix distribution, as in the quantitative model, the employment share of birth cohort $c = t - a$ can be written as

$$\begin{aligned} \varsigma_j(a, t) &= \Pr((\hat{h}_P, \hat{h}_F) \in R_j(a, t)) \\ &= \int_{R_j(a, t)} \phi_2(h_P, h_F; \mu_{\hat{h},c}, \Sigma_{\hat{h},c}) dh_P dh_F, \end{aligned} \tag{83}$$

where ϕ_2 is the bivariate normal density. Employment thus depends on how many workers lie inside the region, not on its area. With distinct sectoral ability weights, ties lie on lines with probability zero, so the sector shares sum to one.

The moments of the ability-mix distribution follow directly from subtracting each worker's average permanent log ability. For $k, r \in \{P, F\}$, their entries are

$$\begin{aligned} \mu_{\hat{h},c,k} &= \mu_{c,k} - \frac{1}{3} \sum_u \mu_{c,u}, \\ \Sigma_{\hat{h},c,kr} &= \Sigma_{c,kr} - \frac{1}{3} \sum_u \Sigma_{c,ku} - \frac{1}{3} \sum_u \Sigma_{c,ur} + \frac{1}{9} \sum_u \sum_v \Sigma_{c,uv}. \end{aligned} \tag{84}$$

Only the physical and fluid-cognitive coordinates enter the bivariate density; the crystallized-and-interpersonal coordinate is determined by their sum.

Effective labor. Employment shares describe the allocation of a cohort across sectors. To determine its contribution to production, one must also account for the efficiency of the workers choosing each sector. The overall ability component $\bar{m}$, which was irrelevant in sectoral comparisons, becomes important: workers with the same ability mix choose the same sector, but those with higher overall ability supply more effective labor.

Each cohort member contributes their efficiency in sector j if they choose it, and zero otherwise. The average contribution per cohort member is

$$\mathbb{E}\left[z_j(a) \mathbf{1}\{(\hat{h}_P, \hat{h}_F) \in R_j(a, t)\}\right].$$

This is an average over the entire cohort, not only over workers employed in j . To calculate it, first consider the efficiency the cohort would supply if all its members worked in that sector. Its average permanent component is

$$\begin{aligned} J_j(c) &= \mathbb{E}\left[e^{\sum_k \theta_{j,k}(\nu_k + \alpha_k s)}\right] \\ &= \exp\left(\sum_k \theta_{j,k} \mu_{c,k} + \frac{1}{2} \sum_k \sum_r \theta_{j,k} \theta_{j,r} \Sigma_{c,kr}\right). \end{aligned} \quad (85)$$

Multiplying by $\tilde{z}_j(a)$ gives average potential efficiency at age a . The remaining question is what fraction of this potential efficiency is supplied by the workers who actually choose j .

To answer it, weight each worker by their efficiency in sector j rather than counting everyone equally. For any measurable region $\mathcal{R}$ of the ability-mix plane, define

$$\Pr((\tilde{h}_P^j, \tilde{h}_F^j) \in \mathcal{R}) = \frac{\mathbb{E}\left[e^{\sum_k \theta_{j,k}(\nu_k + \alpha_k s)} \mathbf{1}\{(\hat{h}_P, \hat{h}_F) \in \mathcal{R}\}\right]}{J_j(c)}. \quad (86)$$

The common age factor $\tilde{z}_j(a)$ cancels from this ratio. This is the main text's efficiency-weighted distribution. It depends on the sector and cohort; it does not describe a change in workers' abilities or choices.

The Gaussian assumption makes this weighting particularly simple. The pair $(\tilde{h}_P^j, \tilde{h}_F^j)$ is jointly Gaussian with the same covariance matrix $\Sigma_{\hat{h},c}$ as the original pair. Only its means change. For $k \in \{P, F\}$, the adjusted mean is

$$\begin{aligned} \tilde{\mu}_{j,c,k} &= \mu_{\hat{h},c,k} + \text{Cov}\left(\hat{h}_k, \sum_r \theta_{j,r}(\nu_r + \alpha_r s)\right) \\ &= \mu_{\hat{h},c,k} + \sum_r \theta_{j,r} \left(\Sigma_{c,kr} - \frac{1}{3} \sum_u \Sigma_{c,ur}\right). \end{aligned} \quad (87)$$

The mean of a relative-ability component rises under this weighting when it is positively associated with permanent log efficiency in sector j , and falls when the association is negative. The sectoral choice region itself does not change.

It follows that the contribution to sector j per cohort member is

$$\begin{aligned} &\mathbb{E}\left[z_j(a) \mathbf{1}\{(\hat{h}_P, \hat{h}_F) \in R_j(a, t)\}\right] \\ &= \tilde{z}_j(a) J_j(c) \Pr((\tilde{h}_P^j, \tilde{h}_F^j) \in R_j(a, t)). \end{aligned} \quad (88)$$

This is (18): average potential efficiency multiplied by the fraction supplied by the workers who select into the sector. Employment and effective labor use the same region and the same covariance matrix, but different means in the Gaussian probability calculation.

Proof of Proposition 3.2. For this proof only, collect the three permanent log abilities in the vector

$$\mathbf{m} = (\nu_P + \alpha_P s, \nu_F + \alpha_F s, \nu_C + \alpha_C s)^\top$$

and the sectoral weights in $\theta_j = (\theta_{j,P}, \theta_{j,F}, \theta_{j,C})^\top$. Then $\mathbf{m} \sim \mathcal{N}(\mu_c, \Sigma_c)$ and permanent efficiency in sector j is $e^{\theta_j^\top \mathbf{m}}$.

For positive-definite Σ_c , completing the square gives

$$\begin{aligned} \theta_j^\top \mathbf{m} - \frac{1}{2}(\mathbf{m} - \mu_c)^\top \Sigma_c^{-1}(\mathbf{m} - \mu_c) \\ = \theta_j^\top \mu_c + \frac{1}{2}\theta_j^\top \Sigma_c \theta_j \\ - \frac{1}{2}(\mathbf{m} - \mu_c - \Sigma_c \theta_j)^\top \Sigma_c^{-1}(\mathbf{m} - \mu_c - \Sigma_c \theta_j). \end{aligned}$$

Consequently,

$$e^{\theta_j^\top \mathbf{m}} \phi_3(\mathbf{m}; \mu_c, \Sigma_c) = J_j(c) \phi_3(\mathbf{m}; \mu_c + \Sigma_c \theta_j, \Sigma_c), \quad (89)$$

where ϕ_3 is the trivariate normal density. The multiplicative constant is $J_j(c)$ in (85). After dividing by this constant, weighting the original Gaussian density by sectoral efficiency therefore gives another Gaussian density, with mean $\mu_c + \Sigma_c \theta_j$ and unchanged covariance.

To recover relative abilities, subtract the average of these three components from each one and retain the physical and fluid-cognitive coordinates. Applying this same linear transformation to both distributions gives the mean adjustment in (87) and leaves the covariance of the pair unchanged.

Integrating (89) over permanent abilities whose centered pair lies in $R_j(a, t)$, and multiplying by $\tilde{z}_j(a)$, proves (88). The argument works for any measurable region, not only a polygon.²⁷ $\square$

Why overall ability is retained. The change in means accounts for the association between relative strengths and overall ability; it does not assume they are independent. Since $\sum_r \theta_{j,r}(\nu_r + \alpha_r s) = \bar{m} + \sum_r \theta_{j,r} \hat{h}_r$, the mean shift can also be written as

$$\tilde{\mu}_{j,c,k} - \mu_{\hat{h},c,k} = \text{Cov}(\hat{h}_k, \bar{m}) + \sum_r \theta_{j,r} \text{Cov}(\hat{h}_k, \hat{h}_r).$$

The first term matters when workers' relative strengths are associated with their overall ability. In particular,

$$\text{Cov}(\hat{h}_k, \bar{m}) = \frac{1}{3} \sum_r \Sigma_{c,kr} - \frac{1}{9} \sum_u \sum_r \Sigma_{c,ur}.$$

Under joint Gaussianity, overall ability and the relative-ability pair are independent only when the row sums of Σ_c are equal. Unequal residual variances, correlations between abilities, or schooling can generate dependence. The aggregation formula retains this dependence, as well as the level of overall ability through $J_j(c)$.

²⁷The density proof uses positive-definite Σ_c , as in the quantitative model. The corresponding expectation identity also holds for a singular Gaussian distribution: express it as an affine transformation of independent standard normal variables and apply the argument in those coordinates.

Aggregation across cohorts. For the costless-mobility case considered here, weighting each cohort by its population share and summing over working ages gives

$$L_{j,t} = \sum_{a \in \mathcal{A}^W} \pi_t(a) \Pr((\hat{h}_P, \hat{h}_F) \in R_j(a, t)),$$

$$N_{j,t} = \sum_{a \in \mathcal{A}^W} \pi_t(a) \tilde{z}_j(a) J_j(t-a) \Pr((\tilde{h}_P^j, \tilde{h}_F^j) \in R_j(a, t)).$$

Each probability uses the distribution of cohort $c = t - a$. Headcount employment counts workers equally; effective labor weights them by their productive efficiency. Since $\pi_t(a)$ is a share of the model population, the sectoral employment share is $L_{j,t}/L_t$, with $L_t = \sum_{a \in \mathcal{A}^W} \pi_t(a)$. Effective labor enters production through $Y_{j,t} = A_{j,t} N_{j,t}$.

The next subsection introduces costly agricultural exit. Its aggregation result uses the same efficiency-weighted distribution, with the employment regions implied by workers' career choices; the derivation is given in Appendix J.1.4.

J.1.4 The cost of leaving agriculture

The previous subsections considered workers who can change sectors freely. I now introduce the cost of leaving agriculture in (19). This makes agricultural employment depend on past choices and future opportunities. The following derivation shows how the same relative-ability coordinates and aggregation formulas extend to this setting.

The cost and the worker's choices. A worker who was employed in agriculture in the previous year and chooses another sector incurs the nonmonetary cost $\kappa_A(a)w_{A,t}z_A(a)$. The coefficient $\kappa_A(a) \geq 0$ measures this cost relative to what the worker would earn in agriculture in the year of departure. The cost is independent of the destination sector. Entering agriculture and moving among manufacturing and services are costless; a worker who returns to agriculture can face another cost upon leaving again. The cost affects sectoral choices but is not deducted from output or the dynasty's income. A worker who leaves supplies their full efficiency in the destination sector.

Workers value current earnings at the dynasty's marginal value of resources λ_t , discount the future by β , and weight the following year by the cohort-accounting ratio $s_t(a)$ defined in (6). Throughout this subsection, wage, λ_t , and cohort-ratio paths are taken as given. Every expression refers to a particular worker; their permanent abilities and schooling are suppressed from the notation.

Only one feature of the previous year's employment matters for today's decision: whether it was in agriculture. Let $V_t^A(a)$ denote the value of current and future choices for a worker previously employed in agriculture, and $V_t^O(a)$ the value for the same worker previously employed outside agriculture. These superscripts describe last year's employment, not the sector chosen today.

Conditional on choosing non-agricultural work, the best destination is the sector with the highest

current earnings. Any exit cost is the same across the three destinations, and all lead to the same value of future choices. Write these highest earnings as

$$\bar{e}_t(a) = \max_{j \in \{M, L, H\}} w_{j,t} z_j(a).$$

The worker's problem can therefore be written using only two value functions:

$$\begin{aligned} V_t^O(a) &= \max \left\{ \begin{aligned} &\lambda_t w_{A,t} z_A(a) + \beta s_t(a) V_{t+1}^A(a+1), \\ &\lambda_t \bar{e}_t(a) + \beta s_t(a) V_{t+1}^O(a+1) \end{aligned} \right\}, \\ V_t^A(a) &= \max \left\{ \begin{aligned} &\lambda_t w_{A,t} z_A(a) + \beta s_t(a) V_{t+1}^A(a+1), \\ &\lambda_t [\bar{e}_t(a) - \kappa_A(a) w_{A,t} z_A(a)] + \beta s_t(a) V_{t+1}^O(a+1) \end{aligned} \right\}. \end{aligned} \quad (90)$$

In each maximum, the first term is the value of working in agriculture and the second is the value of working outside it. The only difference between the two problems is the current exit cost. The terminal conditions are

$$V_t^A(\bar{a} + 1) = V_t^O(\bar{a} + 1) = 0.$$

Workers entering at age 20 have no previous sector and face no current exit cost, so their initial decision uses $V_t^O(20)$. They still anticipate the costs of leaving agriculture in later years.

Entering and remaining in agriculture. Define the value of starting the year outside rather than inside agriculture as

$$\Delta_t(a) = V_t^O(a) - V_t^A(a).$$

This difference is nonnegative: a worker previously outside agriculture has all the same choices without having to pay today's exit cost. It cannot exceed that cost, valued at λ_t :

$$0 \leq \Delta_t(a) \leq \kappa_A(a) \lambda_t w_{A,t} z_A(a). \quad (91)$$

Comparing the two alternatives in (90), a worker previously outside agriculture, or entering the labor market, chooses agriculture if and only if

$$\lambda_t [w_{A,t} z_A(a) - \bar{e}_t(a)] \geq \beta s_t(a) \Delta_{t+1}(a+1). \quad (92)$$

A worker previously in agriculture remains if and only if

$$\lambda_t [(1 + \kappa_A(a)) w_{A,t} z_A(a) - \bar{e}_t(a)] \geq \beta s_t(a) \Delta_{t+1}(a+1). \quad (93)$$

Ties are resolved in favor of agriculture; the proof below shows that they occur with probability zero under the model's assumptions.

The right-hand side is the value of reaching next year with a non-agricultural employment history. Someone considering entry gives up this future advantage by choosing agriculture today. Agriculture's

current earnings advantage must compensate for that loss. Consequently, a worker enters agriculture only if it currently pays the most, and not always then. Someone already in agriculture also avoids the current exit cost by staying. This can keep them there even when another sector currently pays more.

Anticipating future costs can also induce a worker to leave while agriculture still offers the highest current earnings. From (93) and the bound (91), such a departure requires

$$\beta s_t(a) \kappa_A(a+1) \lambda_{t+1} w_{A,t+1} \tilde{z}_A(a+1) > \kappa_A(a) \lambda_t w_{A,t} \tilde{z}_A(a).$$

The worker's permanent agricultural efficiency cancels from this comparison. The left-hand side values next year's exit cost using the same discounting and cohort weights as the worker's problem. It must exceed today's cost. This condition is necessary, not sufficient: the worker's earnings opportunities outside agriculture also matter.

Relative abilities and careers. The following proposition establishes that the two-coordinate representation survives the exit cost. The career statement follows a cohort from the common costless entry decision at age 20.

Proposition J.1 (Costly exit from agriculture). Suppose the ability-weight vectors of different sectors are distinct and each sums to one, wages and λ_t are strictly positive, $s_t(a) \geq 0$, and $\kappa_A(a) \geq 0$. Suppose also that each cohort's pair $(\hat{h}_P, \hat{h}_F)$ has a nonsingular Gaussian distribution. Then:

1. Multiplying all three abilities by the same factor multiplies both value functions by that factor and leaves choices unchanged. Given whether the worker was in agriculture in the previous year, the choice depends on permanent abilities and schooling only through $(\hat{h}_P, \hat{h}_F)$.
2. Conditional on choosing non-agricultural work, the worker chooses the sector with the highest current earnings. Agricultural entry and staying obey (92) and (93), with the bound (91).
3. Starting from the common entry rule, a worker's ability mix determines their sector at every age. The resulting employment regions have well-defined probabilities, cover the ability-mix plane, and overlap only where workers are indifferent. These ties have probability zero.
4. If exit costs are zero at every working age, the employment regions, shares, and effective labor reduce to the costless case in Propositions 3.1 and 3.2.

Proof. Overall ability. Earnings can be written as

$$\lambda_t w_{j,t} z_j(a) = e^{\bar{m}} \lambda_t w_{j,t} \tilde{z}_j(a) e^{\sum_k \theta_{j,k} \hat{h}_k}. \quad (94)$$

The cost is proportional to agricultural earnings and therefore carries the same factor $e^{\bar{m}}$. At retirement both value functions are zero. Working backward from retirement, suppose that multiplying a

worker's three abilities by a common positive factor multiplies next year's values by that factor. It also multiplies every current payoff by the same factor. Both maxima in (90) therefore scale by that factor, leaving their maximizing choices unchanged. This proves part 1; it requires no upper bound on $\kappa_A(a)$.

The value of starting outside agriculture. Subtracting the two maxima in (90) gives

$$\Delta_t(a) = \min \left\{ \max \left[\lambda_t(\bar{e}_t(a) - w_{A,t}z_A(a)) + \beta s_t(a)\Delta_{t+1}(a+1), 0 \right], \right. \\ \left. \kappa_A(a)\lambda_t w_{A,t}z_A(a) \right\}, \quad (95)$$

with $\Delta_t(\bar{a}+1) = 0$. This establishes the bound. The difference is zero when both inherited employment states lead to agriculture, and equals the current exit cost when both lead outside. It lies between these values when the inherited state changes the chosen sector.

Comparing the two alternatives in each maximum gives (92) and (93). All non-agricultural destinations have the same exit cost and continuation value, so their ranking depends only on current earnings. This proves part 2.

Employment histories and regions. Let $R_A^{\text{in}}(a, t)$ contain the ability mixes satisfying (92), and let $R_A^{\text{stay}}(a, t)$ contain those satisfying (93). Both are conditions on $(\hat{h}_P, \hat{h}_F)$: by part 1, all terms share the factor $e^{\bar{m}}$. Since the exit cost is nonnegative,

$$R_A^{\text{in}}(a, t) \subseteq R_A^{\text{stay}}(a, t).$$

A worker willing to enter agriculture would also be willing to remain there if they were already employed in it.

For $j \in \{M, L, H\}$, let

$$R_j^N(a, t) = \left\{ (\hat{h}_P, \hat{h}_F) : \sum_k (\theta_{j,k} - \theta_{\ell,k}) \hat{h}_k \geq \log \frac{w_{\ell,t} \tilde{z}_\ell(a)}{w_{j,t} \tilde{z}_j(a)} \right. \\ \left. \text{for every } \ell \in \{M, L, H\} \setminus \{j\} \right\}. \quad (96)$$

This region identifies the highest-paying non-agricultural sector; it does not yet determine whether the worker leaves or avoids agriculture. Each such region is the intersection of two half-planes and hence is convex, with straight boundaries.

Follow a cohort as both age and date advance by one year. At entry, workers in R_A^{in} choose agriculture. Subsequently, a worker is in agriculture if they enter it or were already there and remain. Thus

$$R_A^\kappa(20, t) = R_A^{\text{in}}(20, t), \\ R_A^\kappa(a+1, t+1) = R_A^{\text{in}}(a+1, t+1) \\ \cup [R_A^\kappa(a, t) \cap R_A^{\text{stay}}(a+1, t+1)]. \quad (97)$$

Including all of R_A^{in} is valid because it is contained in R_A^{stay} : those workers choose agriculture whether or not they worked there previously. Among the remaining workers, the highest-paying non-agricultural sector determines employment, so

$$R_j^{\kappa}(a, t) = R_j^N(a, t) \setminus R_A^{\kappa}(a, t), \quad j \in \{M, L, H\}. \quad (98)$$

This establishes the career representation in part 3. It uses the common entry rule and common paths faced by workers within a cohort. For a cohort already alive at the initial date of a reported period, the representation applies when its inherited employment is generated by that history.

Probabilities and indifference. Dividing (95) by $e^{\bar{m}}$ gives a continuous function of the two relative abilities, because earnings are continuous and the recursion uses only sums, maxima, and minima. The entry and staying sets are therefore closed. The employment regions are formed from these sets by finite unions, intersections, and differences, so their probabilities are well defined. Agriculture and the non-agricultural employment regions are disjoint by (98). The latter overlap only on equal-earnings lines, which have probability zero under a nonsingular Gaussian distribution.

Ties between agriculture and non-agricultural work also have probability zero. To see this, consider the finitely many branch combinations of the finite-horizon recursion. On each branch, $\Delta_t(a)/e^{\bar{m}}$ is a linear combination of the sectoral terms $\exp(\sum_k \theta_{j,k} \hat{h}_k)$, with nonnegative coefficients on the non-agricultural terms. This follows backward from zero at retirement: a branch either takes zero, takes the agricultural cost bound, or adds current non-agricultural earnings and a nonnegative multiple of next year's difference while subtracting agricultural earnings.

Write each entry or staying comparison as the value of non-agricultural work minus the value of agricultural work. After dividing by $e^{\bar{m}}$, it contains a strictly positive coefficient on today's best non-agricultural term, which future terms cannot cancel. Distinct sectoral ability weights imply distinct exponential terms in the two coordinates. Each tie equation is therefore a nonzero real-analytic function; its zero set has area zero. A finite union of these sets also has zero Gaussian probability. The regions cover the plane and account for every worker, completing part 3.

Zero exit costs. If $\kappa_A(a) = 0$ at every working age, the bound gives $\Delta_t(a) = 0$. Both entry and staying then reduce to the comparison $w_{A,t} z_A(a) \geq \bar{e}_t(a)$, so $R_A^{\text{in}} = R_A^{\text{stay}} = R_A$. Equations (97) and (98) recover the costless regions, apart from zero-probability ties. This proves part 4. $\square$

What changes in the employment regions. With costly exit, the comparisons among manufacturing and services retain their straight-line boundaries. Agricultural employment is different: its boundaries depend on future choices and past employment, and need not be straight or enclose a convex region. This can change employment in all four sectors, even though the current-earnings ranking among the three non-agricultural sectors is unchanged at given wages. At the last working age, the

continuation term vanishes; entry compares current earnings, while staying also reflects the current exit cost.

Figure 36 illustrates these regions for the cohort entering the labor market in 1976 in the calibrated baseline economy.




Figure 36: Sectoral allocation with costly exit from agriculture. *Notes:* Dots show the same simulated workers from the cohort entering in 1976, at ages 20 (1976) and 64 (2020). Colors identify their chosen sectors in the calibrated baseline, including agricultural exit costs; percentages report full-cohort employment shares. Workers anticipate future earnings and exit costs, and their choices are followed from labor-market entry. Unlike Figure 4, this figure uses the baseline allocation without an illustrative adjustment to agricultural returns.

Employment and effective labor. The employment regions change, but the aggregation calculation does not. For cohort $c = t - a$,

$$\varsigma_j(a, t) = \Pr((\hat{h}_P, \hat{h}_F) \in R_j^\kappa(a, t)).$$

The proof in Appendix J.1.3 applies to any measurable region, not only one bounded by straight lines. Applying it to $R_j^\kappa(a, t)$ gives

$$\begin{aligned} \mathbb{E}[z_j(a) \mathbf{1}\{(\hat{h}_P, \hat{h}_F) \in R_j^\kappa(a, t)\}] \\ = \tilde{z}_j(a) J_j(c) \Pr((\tilde{h}_P^j, \tilde{h}_F^j) \in R_j^\kappa(a, t)). \end{aligned} \quad (99)$$

Both $J_j(c)$ and the efficiency-weighted distribution are unchanged from Appendix J.1.3, including the contribution of overall ability and its association with relative strengths. No exit cost is subtracted: the cost changes the worker's decision, not their productive contribution after choosing a sector.

Aggregating across cohorts gives

$$\begin{aligned} L_{j,t} &= \sum_{a \in \mathcal{A}^W} \pi_t(a) \varsigma_j(a, t), \\ N_{j,t} &= \sum_{a \in \mathcal{A}^W} \pi_t(a) \tilde{z}_j(a) J_j(t - a) \Pr((\tilde{h}_P^j, \tilde{h}_F^j) \in R_j^\kappa(a, t)). \end{aligned}$$

These are (21) and (22). Sectoral employment shares are $L_{j,t}/L_t$, and output is $Y_{j,t} = A_{j,t}N_{j,t}$.

Why the cost is proportional to the worker's own earnings. The proportionality of the cost is what preserves the two-coordinate choice problem. If the cost were instead an amount $q_t \kappa_A(a)$ common to all workers, with $q_t > 0$, multiplying a worker's three abilities by a common factor would raise earnings but leave the cost unchanged. Workers with the same relative strengths could then make different choices because the cost would represent a different burden relative to their earnings. Overall ability would generally matter for sectoral choice as well as production, so the two-coordinate representation would no longer suffice.

J.2 Baseline demand, trade, and equilibrium

This section completes the equilibrium description for the baseline without explicit durable accumulation. Non-durable expenditure, sectoral spending shares, and consumption across cohorts are given by (36), (37), and (38), subject to the pooled budget (35).

J.2.1 Trade and resource accounting

Households pay the composite price $p_{j,t}^Q$, while domestic producers receive $p_{j,t}^d$. For a given level of absorption $Q_{j,t}$, cost minimization under (27) yields

$$Y_{j,t}^{\text{abs}} = \omega_j \left(\frac{p_{j,t}^d}{p_{j,t}^Q} \right)^{-\eta} Q_{j,t}, \quad m_{j,t} = (1 - \omega_j) \left(\frac{p_{j,t}^m}{p_{j,t}^Q} \right)^{-\eta} Q_{j,t}. \quad (100)$$

Spending on the composite therefore equals spending on its domestic and imported components:

$$p_{j,t}^Q Q_{j,t} = p_{j,t}^d Y_{j,t}^{\text{abs}} + p_{j,t}^m m_{j,t}.$$

For $j \in \{A, M\}$, the export conditions are $X_{j,t} \geq 0$ and $p_{j,t}^d \geq e_t p_{j,t}^*/\tau_j^X$, with equality in the price condition whenever exports are positive. Exports and domestic sales exhaust output through (26). While manufacturing exports, $p_{M,t}^d = 1$ fixes $e_t = \tau_M^X/p_{M,t}^*$; otherwise, $e_t \leq \tau_M^X/p_{M,t}^*$ and all manufacturing output is sold at home. Services satisfy the foreign-demand schedules (30).

When goods markets clear, summing the spending identity across sectors and using zero profits gives

$$\begin{aligned} \sum_{a \in \mathcal{A}} \pi_t(a) E_t(a) &= \sum_{j \in \mathcal{J}} p_{j,t}^Q Q_{j,t} \\ &= \sum_{j \in \mathcal{J}} w_{j,t} N_{j,t} - \sum_{j \in \mathcal{J}} (p_{j,t}^d X_{j,t} - p_{j,t}^m m_{j,t}). \end{aligned} \quad (101)$$

Thus household spending equals labor income less net exports. Substituting the external-balance condition (31) gives the pooled budget with equality. Conditional on goods-market clearing and zero profits, the budget and external-balance conditions are therefore equivalent, not two independent resource restrictions. Agricultural exit costs do not enter this accounting because they are nonmonetary.

J.2.2 Equilibrium

Definition J.1 (Baseline equilibrium). Given the exogenous paths and parameters in Section 3, and the inherited sector assignments of workers already alive at the initial date, consistent with the common-entry histories in Appendix J.1.4, an equilibrium consists of price paths $\{p_{j,t}^d, p_{j,t}^m, p_{j,t}^Q, w_{j,t}, e_t\}$, associated budget multipliers $\{\lambda_t\}$, and allocations of consumption, labor, production, and trade satisfying:

1. *Firms.* Production and wages satisfy (23) and (24), with the manufacturing producer price normalized to one.
2. *Workers.* Sectoral careers solve (90). Cohort employment shares, sectoral employment, and effective labor satisfy (20), (21), and (22).
3. *The dynasty.* Consumption maximizes (34) subject to the pooled budget (35), which holds with equality. Expenditure, sectoral shares, and consumption indices satisfy (36), (37), and (38).
4. *Trade.* Composite prices and domestic and import demands satisfy (28) and (100). Exports satisfy the agriculture and manufacturing conditions in Appendix J.2.1 and (30) for services. External balance satisfies (31).
5. *Market clearing.* Firms employ the effective labor supplied by workers. Domestic output is allocated between domestic use and exports according to (26), and composite absorption equals household purchases as in (39).

J.3 Manufactured durable goods

This appendix develops the extension summarized in Section 3.9. I extend the non-homothetic CES demand system to include manufactured durable goods, linking their accumulation to income and relative prices. Unlike non-durable goods, which are consumed in the period of purchase, durables add to a cohort's stock, contributing to current and future consumption.

J.3.1 Stocks, purchases, and the household problem

Let $D_t(a)$ denote the durable stock held by each individual aged a at date t . The consumption index includes this stock alongside the four non-durable goods:

$$\sum_{j \in \mathcal{J}} \Omega(a, j)^{1/\sigma} \left(\frac{c_{j,t}(a)}{\mathcal{C}_t(a)^{\iota_j}} \right)^{\frac{\sigma-1}{\sigma}} + \Omega_D(a)^{1/\sigma} \left(\frac{\bar{d} + D_t(a)}{\mathcal{C}_t(a)^{\iota_D}} \right)^{\frac{\sigma-1}{\sigma}} = 1. \quad (102)$$

Here, $\iota_D > 0$ and $\Omega_D(a) > 0$ are the income-effect parameter and age-specific taste weight for durables, with $\sum_{j \in \mathcal{J}} \Omega(a, j) + \Omega_D(a) = 1$. As in the baseline, $0 < \sigma < 1$ and all income-effect parameters and

taste weights are positive. The translation parameter $\bar{d} \geq 0$, when positive, allows zero durable holdings, as observed for some households.²⁸

The dynasty chooses each cohort's durable stock. With depreciation rate $\delta_D \in (0, 1)$, new purchases per cohort member $x_t(a)$ equal the chosen stock minus undepreciated holdings carried forward:

$$x_t(a) = D_t(a) - (1 - \delta_D)D_{t-1}(a - 1), \quad a = 21, \dots, a_{\max}. \quad (103)$$

Stocks and purchases are measured in units of the manufacturing absorption composite. Entering cohorts have no inherited stock, so $x_t(20) = D_t(20)$. Installed durables cannot be sold or reassigned across cohorts:

$$x_t(a) \geq 0 \quad \Longleftrightarrow \quad D_t(a) \geq (1 - \delta_D)D_{t-1}(a - 1), \quad (104)$$

where inherited holdings are zero at age 20 and specified as initial conditions for cohorts already alive at the start of the model.

The per-member stock law also requires an accounting convention when cohort membership changes. Holdings of departing members are lost, while members joining an existing cohort are endowed with its per-member stock. The cohort ratio $s_t(a)$ in (6) therefore also enters the valuation of future durable holdings. At the terminal age $a_{\max}$, remaining stocks have no resale or scrap value.

Entering cohorts must purchase any durable stock they hold; older cohorts carry undepreciated holdings. This stock-building motive can contribute to the higher durable expenditure shares of younger households and links population age composition to manufactured-goods demand.

Let $E_t(a) \equiv \sum_{j \in \mathcal{J}} p_{j,t}^Q c_{j,t}(a)$ denote non-durable expenditure. With durable purchases, the dynasty's date- t budget is

$$\sum_{a \in \mathcal{A}} \pi_t(a) E_t(a) + p_{M,t}^Q \sum_{a \in \mathcal{A}} \pi_t(a) x_t(a) \leq \sum_{j \in \mathcal{J}} w_{j,t} N_{j,t} - \mathcal{W}_t \sum_{j \in \mathcal{J}} p_{j,t}^d Y_{j,t}. \quad (105)$$

The objective remains (34). For each cohort, the dynasty jointly chooses the composition of the non-durable bundle, the consumption index, and the durable stock. I first derive the least costly non-durable bundle for a given consumption index and stock, then the conditions determining the index and stock. These conditions must hold jointly.

J.3.2 Non-durable expenditure given durable holdings

Durable holdings contribute to consumption, so a larger stock allows the same consumption index to be attained with fewer non-durable goods. Conditional non-durable expenditure therefore equals the four-good expression in (36), evaluated with the extension's preference parameters, multiplied by an adjustment for durable holdings.

²⁸This plays a role analogous to the committed quantities of Pollak and Wales (1981) and the service terms of Kongsamut, Rebelo, and Xie (2001) and Herrendorf, Rogerson, and Valentinyi (2013). With $\bar{d} = 0$, durables are essential and entering cohorts must purchase a positive stock. Appendix K.1.3 describes the calibration.

A feasible pair of consumption and holdings requires $C_t(a) > 0$, $\bar{d} + D_t(a) > 0$, and a durable term strictly below one in (102). The remaining positive terms must be supplied by non-durable goods. At a given stock, this places an upper bound on the consumption index: non-durable spending alone cannot raise consumption indefinitely.

The expenditure adjustment depends on $\xi_t(a)$, the ratio of the durable term to the sum of the four non-durable terms in (102):

$$\xi_t(a) \equiv \frac{\Omega_D(a)^{1/\sigma} [(\bar{d} + D_t(a))C_t(a)^{-\iota_D}]^{\frac{\sigma-1}{\sigma}}}{1 - \Omega_D(a)^{1/\sigma} [(\bar{d} + D_t(a))C_t(a)^{-\iota_D}]^{\frac{\sigma-1}{\sigma}}}. \quad (106)$$

Required non-durable expenditure is

$$E_t(a) = (1 + \xi_t(a))^{\frac{\sigma}{1-\sigma}} \left[\sum_{j \in \mathcal{J}} \Omega(a, j) C_t(a)^{(1-\sigma)\iota_j} (p_{j,t}^Q)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \quad (107)$$

To obtain this expression, fix the index and durable stock. Multiplying every quantity in the four-good expenditure-minimization problem by $(1 + \xi_t(a))^{\sigma/(1-\sigma)}$ rescales the sum of its consumption-index terms from one to $1/(1 + \xi_t(a))$, exactly the amount left after the durable term. Expenditure is multiplied by the same factor. This common rescaling leaves the non-durable expenditure shares unchanged, so they remain those in (37), and quantities are $c_{j,t}(a) = \zeta_{j,t}(a)E_t(a)/p_{j,t}^Q$.

At a given consumption index, $\xi_t(a)$ decreases as durable holdings rise, reducing the additional spending required. The four-good expression is the non-durable expenditure required in the limit of abundant durable holdings, $\xi_t(a) \rightarrow 0$. Finite holdings require additional spending to attain the same index. For a fixed durable stock, expenditure diverges as the consumption index approaches its upper bound.

At a given consumption index and prices, the durable adjustment changes spending on all four non-durable goods proportionally. Income effects change their shares only when the consumption index changes. Differentiating (107) gives the two expressions needed for the household's choices:

$$\begin{aligned} \frac{\partial E_t(a)}{\partial C_t(a)} &= \frac{E_t(a)}{C_t(a)} [\bar{\iota}_t(a) + \xi_t(a)\iota_D], \\ -\frac{\partial E_t(a)}{\partial D_t(a)} &= \frac{E_t(a)\xi_t(a)}{\bar{d} + D_t(a)}, \end{aligned} \quad (108)$$

where $\bar{\iota}_t(a) = \sum_{j \in \mathcal{J}} \zeta_{j,t}(a)\iota_j$, as in the baseline. The first expression measures the extra non-durable expenditure required to raise the index at a given stock. The second measures the expenditure saved by additional holdings at a given index.

J.3.3 Consumption and durable-stock choices

The consumption index. Using (107) in place of the baseline expenditure function, the consumption-index condition for a given durable stock is

$$\frac{1}{\mathcal{C}_t(a)} = \lambda_t \frac{\partial E_t(a)}{\partial \mathcal{C}_t(a)} = \lambda_t \frac{E_t(a)}{\mathcal{C}_t(a)} [\bar{\iota}_t(a) + \xi_t(a)\iota_D], \quad (109)$$

where λ_t is the common current-value multiplier on the pooled budget (105).

Relative to the baseline condition (38), the additional term $\xi_t(a)\iota_D$ reflects how the expenditure adjustment increases when consumption rises but durable holdings remain fixed. For a given $\xi_t(a)$, a larger durable income-effect parameter ι_D makes the adjustment rise more sharply with consumption.

Durable purchases. The durable-stock condition must hold jointly with the consumption-index condition. Additional purchases cost resources today, but increase both current durable holdings and the undepreciated stock available for future consumption. Substituting the accumulation equation (103) into the budget gives, when the cohort purchases durables both today and next period,

$$\underbrace{\lambda_t E_t(a) \frac{\xi_t(a)}{\bar{d} + D_t(a)}}_{\text{marginal current benefit}} = \underbrace{\lambda_t p_{M,t}^Q - \beta s_t(a)(1 - \delta_D) \lambda_{t+1} p_{M,t+1}^Q}_{\text{user cost}}, \quad a < a_{\max}. \quad (110)$$

The left-hand side measures the current consumption benefit through the non-durable spending that additional holdings replace. At given prices and consumption index, an additional unit saves $E_t(a)\xi_t(a)/[\bar{d} + D_t(a)]$ in expenditure. At the optimal consumption choice, multiplying this saving by λ_t expresses the benefit in utility terms. Keeping next period's holdings unchanged, each additional unit bought today saves $1 - \delta_D$ units of future purchases. The right-hand side is today's purchase cost net of the discounted value of this saving. At the terminal age, the user cost is simply $\lambda_t p_{M,t}^Q$.²⁹

Allowing zero purchases. When a cohort already owns enough durables, it may choose no new purchases. Let $\chi_t(a) \geq 0$ denote the multiplier on $x_t(a) \geq 0$, measured in current utility units per cohort member. It is zero whenever purchases are positive. Omitting terms independent of consumption and durable-stock choices, the household Lagrangian is

$$\mathcal{L} = \sum_{t=0}^{\infty} \beta^t \sum_{a \in \mathcal{A}} \pi_t(a) \left\{ \log \mathcal{C}_t(a) - \lambda_t E_t(a) - \lambda_t p_{M,t}^Q x_t(a) + \chi_t(a) x_t(a) \right\}.$$

Here $E_t(a)$ is given by (107), and $x_t(a)$ is replaced using the stock law (103), with zero inherited holdings for entrants. Differentiating with respect to $D_t(a)$ affects both current expenditure and next

²⁹The annual-cost conversion used in the holdings-based calibration need not equal this model-implied user cost in expenditure units.

period's purchases; dividing by $\beta^t \pi_t(a)$ gives the relative weight $\beta \pi_{t+1}(a+1)/\pi_t(a) = \beta s_t(a)$ on the future terms, yielding the full stock condition

$$0 = \lambda_t E_t(a) \frac{\xi_t(a)}{\bar{d} + D_t(a)} - \lambda_t p_{M,t}^Q + \chi_t(a) + \beta s_t(a)(1 - \delta_D) [\lambda_{t+1} p_{M,t+1}^Q - \chi_{t+1}(a+1)], \quad a < a_{\max}, \quad (111)$$

together with

$$x_t(a) \geq 0, \quad \chi_t(a) \geq 0, \quad \chi_t(a)x_t(a) = 0. \quad (112)$$

The future multiplier matters because, when next period's purchases are constrained at zero, extra holdings today cannot generate the purchase saving described above without changing next period's holdings. Equation (110) follows when both multipliers are zero. With $\bar{d} > 0$, entering cohorts can also choose zero holdings; cohorts with inherited stocks cannot reduce them below their undepreciated holdings.

At the terminal age there is no future purchase saving, so

$$\lambda_t E_t(a_{\max}) \frac{\xi_t(a_{\max})}{\bar{d} + D_t(a_{\max})} - \lambda_t p_{M,t}^Q + \chi_t(a_{\max}) = 0. \quad (113)$$

These conditions cover both positive and zero purchases, without imposing a positive user cost at every date.

Why these conditions characterize the optimum. Given finite positive price and λ_t paths, finite fixed inherited holdings, and positive discounted cohort weights over its remaining lifetime, each cohort's conditional problem has a unique optimum. The consumption and stock conditions above, together with feasibility, characterize that allocation. This is a statement about household choices given those paths, not about uniqueness of general equilibrium.

Proof. For this proof only, let $q_j = c_{j,t}(a)$ and $q_D = \bar{d} + D_t(a)$ denote the five positive quantities entering the consumption index, and let T_n denote their corresponding terms in (102), so $\sum_n T_n = 1$. Consider any straight-line change in these quantities, and use primes for derivatives along that change. Differentiating the defining equation twice gives

$$(\log \mathcal{C})'' = \frac{\frac{\sigma-1}{\sigma} \sum_n T_n \left[\frac{q'_n}{q_n} - \iota_n (\log \mathcal{C})' \right]^2 - \sum_n T_n \left(\frac{q'_n}{q_n} \right)^2}{\sum_n T_n \iota_n} < 0$$

for every nonzero change, since $0 < \sigma < 1$ and all weights and income-effect parameters are positive. Thus utility is strictly concave in non-durable quantities and durable holdings: an average of two distinct allocations gives strictly more utility than the corresponding average of their utilities.

For given λ_t , the cohort's objective subtracts the resource value of spending from its discounted utility. Non-durable spending and durable purchases are linear in these quantities, as are the accumulation and no-resale constraints. The objective is therefore strictly concave on a convex feasible

set. Since the constraints are linear and an allocation with strictly positive purchases is feasible, the first-order conditions are necessary as well as sufficient. Strict concavity makes the optimum unique. A finite optimum exists because utility grows at most logarithmically, whereas expenditure costs grow linearly in non-durable quantities and nonnegative purchases over the finite lifetime. If an essential quantity approaches zero, utility tends to minus infinity. $\square$

Equilibrium with durables. Manufacturing absorption includes non-durable consumption and new durable purchases:

$$Q_{M,t} = \sum_{a \in \mathcal{A}} \pi_t(a) c_{M,t}(a) + \sum_{a \in \mathcal{A}} \pi_t(a) x_t(a). \quad (114)$$

The extended budget (105) and this absorption condition replace their baseline counterparts. The conditions for non-manufacturing absorption, production, worker choices, and trade are unchanged from Appendix J.2.2. Total household spending still equals the value of domestic absorption, so the resource-accounting argument there continues to apply. Durable holdings support consumption, but only purchases use current resources.

K Calibration details and additional evidence on model fit

This appendix provides the additional calibration details underlying Section 4, distinguishing external inputs, estimated demand parameters, and parameters fitted to Indonesian outcomes. It then supplements Section 4.6 with figures for annual employment shares and expenditure by age, explains the comparisons in Table 5, and reports manufacturing persistence.

K.1 Calibration details

K.1.1 Demography, schooling, and worker parameters

Demographic and schooling inputs. The solution grid runs annually from 1906 to 2094, while reported outcomes cover 1976–2025. Population weights are normalized over ages 20–89, with entry at 20, retirement after 64, and $\beta = 0.96$. Outside the reporting window, the age distributions are held at their first and last available values. Appendix L.1 describes the demographic data and alignment of the counterfactual paths.

I construct cohort schooling means and variances from weighted individual records. For IFLS respondents, I retain one observation per person, preferring observations at age 25 or above and then the latest available wave; otherwise I use the latest observation at ages 20–24. The sample comprises employed people in the four model sectors with recorded schooling and positive person weights. Indonesian IPUMS records supplement cohorts born in 1912–1923; the observed IFLS worker cohorts cover 1924–1994. For cohorts born in 1995–2005, I project completed schooling from the 2014 IFLS

household roster using schooling already attained, current enrollment, and enrollment rates at subsequent ages through 24. The projection is capped at 21 years, and people no longer enrolled receive no additional projected schooling. These youngest-cohort inputs are therefore projected attainment, not observed completed schooling. Cohort means and standard deviations are held at their endpoint values outside the available birth-year range. The model approximates each cohort’s schooling distribution by an untruncated Gaussian, as in Appendix J.1.1, and normalizes mean residual log abilities to zero.

External and fitted worker inputs. The three age–ability profiles and the sector ability weights are fixed at the externally constructed values in Appendix E. The schooling loadings, the covariance matrix of permanent abilities, and the agricultural exit-cost schedule are calibrated jointly. Tables 44 and 45 report the fitted values. Exit costs are expressed as multiples of the worker’s own potential annual agricultural earnings; they are not subtracted from sectoral production. The schedule uses shape-preserving cubic interpolation through twelve age-specific values, with no cost of entering agriculture.

Schooling–earnings targets. I estimate weighted regressions of log hourly earnings on years of schooling, five-year age indicators, and survey-wave indicators, separately by sector. The sample contains full-time workers aged 20–64 with positive earnings and person weights in the 1993, 1997, 2000, 2007, and 2014 IFLS. Earnings are trimmed at the first and ninety-ninth percentiles within each wave. The manufacturing estimate, 0.0995, anchors $\sum_k \theta_{M,k} \alpha_k$; the agricultural estimate, 0.0691, enters as a penalized target. The implied crystallized-ability loading is restricted to $[0.03, 0.20]$. Its lower bound binds in the durable extension, so the manufacturing anchor is not matched exactly there. The resulting effects on potential log earnings in agriculture and manufacturing are (0.0805, 0.0995) in the baseline and (0.0846, 0.1064) in the extension. These empirical earnings relationships are calibration targets, not causal estimates of schooling’s effects on abilities.

Employment targets and calibration procedure. The worker parameters minimize a weighted sum of squared discrepancies in employment profiles and changes in sectoral shares across approximately matched cohorts. The active employment targets are: four-sector shares in nine five-year age groups; single-age agriculture and manufacturing profiles; three-year-smoothed single-age low- and high-skill-service profiles at ages 20–34; manufacturing and low-skill-service shares within non-agricultural employment; and changes in agricultural and manufacturing shares along cohort diagonals between IFLS waves. Most age-profile targets are measured relative to each sector’s mean across ages, separating profile shape from average sector size. The age-group moments pool 1993, 1997, and 2000 in both model and data. The single-age comparisons pool the model over 1976–2025 and the IFLS over those three waves; the cohort comparisons also use 2007 and 2014.

The objective uses bootstrap-precision weights for the single-age moments and binomial-precision weights for the non-agricultural shares and cohort changes. The age-group block places greater weight on manufacturing and on the two oldest high-skill-service groups. In the baseline, manufacturing’s age-group and single-age terms receive an additional factor of three relative to the durable extension. The size of manufacturing’s employment decline across working ages, in total employment and within non-agricultural employment, is also a targeted moment, using model employment pooled over 1976–2025. Wage-age profiles and individual retention probabilities receive no weight. The objective also penalizes departures from the agricultural schooling–earnings target and abrupt changes in the exit-cost schedule. Worker parameters are fitted at given wage and marginal-utility paths, followed by the general-equilibrium adjustment of relative productivity paths described below.

Table 44: Calibrated worker parameters

Parameter	Baseline	Durable extension
α_P	0.00660	0.01539
α_F	0.30084	0.30938
α_C	0.03375	0.03000
$\text{Var}(\nu_P)$	2.30405	2.37019
$\text{Var}(\nu_F)$	0.52399	0.50210
$\text{Var}(\nu_C)$	3.85952	3.76037
$\text{Cov}(\nu_P, \nu_F)$	0.40639	0.34315
$\text{Cov}(\nu_P, \nu_C)$	-0.15904	-0.15840
$\text{Cov}(\nu_F, \nu_C)$	-0.23152	-0.16162

Notes: P , F , and C denote physical, fluid-cognitive, and crystallized-and-interpersonal abilities. The α_k are schooling loadings, and ν_k are permanent residual log abilities. The variances and covariances form Σ_0 . Sector ability weights remain fixed at the external values in Table 2.

K.1.2 Household expenditure and non-durable preferences

Expenditure data and Engel curves. I construct sectoral household expenditure from the IFLS consumption modules. The four-sector total excludes survey items left outside the sectoral concordance. In the baseline, manufacturing includes durable and semi-durable purchases; these purchases are removed for the non-durable demand block of the extension. Households are grouped into province-by-expenditure-decile cells within each survey wave, after trimming the bottom and top percentiles of household expenditure. Cell sectoral shares are ratios of survey-weighted expenditure totals. For agriculture and each service sector, I regress the log expenditure ratio relative to manufacturing on the log of the cell’s survey-weighted mean household expenditure per adult aged 20 and above, with province–wave fixed effects and weights equal to the cell’s household-expenditure mass. The income-

Table 45: Agricultural exit-cost schedule

Age	Baseline	Durable extension
20	0.15584	0.09016
22	0.20000	0.10505
24	0.20000	0.13536
26	0.20000	0.18580
28	0.21434	0.22952
31	0.20791	0.24735
34	0.25995	0.31965
38	0.36721	0.39677
43	0.50794	0.78876
49	0.74440	1.01141
56	0.87784	0.97522
64	1.01129	0.89778

Notes: A value of one equals one year’s potential agricultural earnings. Values between knots use shape-preserving cubic interpolation. Knots are nonnegative, with upper bounds of 0.2 through age 26 and 5 thereafter; no monotonicity restriction is imposed.

effect parameters use the 1993 and 1997 waves, comprising 260 cells. The expenditure regressor includes total household spending, whereas the sector shares use the mapped four-sector basket.

Substitution elasticity. The common σ is estimated from changes in BPS household-consumption benchmarks and IFLS sectoral shares. GGDC sectoral value-added deflators proxy relative consumption prices, and Penn World Table real consumption per person measures expenditure growth. I subtract the component of sectoral expenditure-ratio growth predicted by the all-wave IFLS Engel slopes, then regress the remainder on relative-price changes. The three sector equations are pooled without an intercept. The selected sample contains eleven within-source time differences over 1975–2016, excluding differences spanning the 1997–1999 crisis and breaks in source definitions. This yields $\sigma = 0.7208$, retained in both models.³⁰

Mapping Engel slopes into income effects. Let b_j^E be the estimated slope on log expenditure, with $b_M^E = 0$. Conditional on $\iota_M = 1$, the baseline uses the mean-share conversion

$$\iota_j = 1 + \frac{b_j^E \bar{\iota}}{1 - \sigma}, \quad \bar{\iota} = \left[1 - \frac{\sum_j \bar{\zeta}_j b_j^E}{1 - \sigma} \right]^{-1}, \quad (115)$$

³⁰The selected macro estimator combines BPS sectoral consumption including manufactured durables with IFLS shares net of the broader durable basket. It uses sectoral value-added deflators rather than matched household consumption-price indices. The retained elasticity is therefore a calibration estimate based on these proxies; it is not re-estimated when the model’s manufacturing basket changes.

where $\bar{\zeta}_j$ is sector j 's share of pooled expenditure in the estimation sample. This gives $\bar{\iota} = 0.5995$ and $(\iota_A, \iota_M, \iota_L, \iota_H) = (0.2364, 1, 0.9331, 1.4673)$. The durable extension retains the earlier expenditure-based conversion with the factor $\bar{\iota}$ set to one, giving $(0.4225, 1, 1.5848, 2.4759)$ for its non-durable goods. Thus the two installed vectors differ both in the manufacturing basket and in the slope-to-index conversion.

Age-specific taste weights. The 1993 IFLS supplies expenditure profiles for twelve groups defined by mean adult household age: under 25, five-year groups from 25–29 to 70–74, and 75 and above. Before recovering taste weights, I add sector-specific constants to align these profiles with aggregate expenditure anchors. The anchors combine sectoral value-added shares from a reference model equilibrium with observed import and export ratios. I infer sectoral absorption from these objects and reconcile its aggregate level with the external-balance wedge, assigning the remaining discrepancy to manufacturing. The shifts preserve within-sector age differences in percentage points; the targets are not the unadjusted level of every survey share. In the baseline, the shifts for (A, M, L, H) , in share units, are $(-0.1740, -0.0220, 0.1191, 0.0769)$, with aggregation weights proportional to population times reference expenditure. In the extension's non-durable block, the shifts are $(-0.2053, 0.0065, 0.1460, 0.0528)$, using population weights after separating durable purchases from the reference manufacturing total.

I recover the four-good taste weights at reference prices, consumption indices, and expenditure levels using

$$\Omega^{\text{four}}(a, j) = \frac{\zeta_j^R(a) [E^R(a)]^{1-\sigma}}{[p_j^{Q,R}]^{1-\sigma} [\mathcal{C}^R(a)]^{(1-\sigma)\iota_j}}. \quad (116)$$

Here, R denotes the reference allocation and $\zeta_j^R(a)$ the expenditure-share target. Retaining the row's scale reproduces reference expenditure as well as shares. In the durable extension, these four-good rows are additionally multiplied by $(1 + \xi^R(a))^{-\sigma}$ as the durable taste weights are constructed jointly, accounting for the contribution of durable holdings to the consumption index. All taste weights are held fixed across demographic counterfactuals.

K.1.3 Durable-goods parameters and targets

Holdings data and the translation parameter. Durable holdings are the reported current values of vehicles and household appliances in the 1993 and 1997 IFLS asset modules. I adjust for partial ownership when reported and express holdings per adult aged 20 and above. Households with zero holdings remain in the sample. To choose the translation parameter, I use a survey-weighted censored regression for 1993 households with province fixed effects. For positive holdings, it relates the log ratio of shifted holdings to manufactured non-durable expenditure to log annual non-durable expenditure per adult; zero holdings enter as censored observations. The selected translation is equivalent to 11,346 rupiah per adult, or 1.092 percent of household-weighted mean annual non-durable expenditure per adult. A fixed conversion based on an earlier reference calibration expresses this as $\bar{d} = 0.02918$ in

model units.³¹

Income sensitivity of durable holdings. I estimate ι_D from the 1993 and 1997 holdings data, aggregated into province-by-expenditure-decile cells. The translation is carried across waves as a fraction of each wave’s mean expenditure, rather than as a fixed nominal rupiah amount. The regression relates shifted holdings relative to manufactured non-durable spending to an iteratively constructed consumption index, controlling for province–wave effects and eleven household-age-composition shares. The index is recovered jointly from the four expenditure-ratio equations for agriculture, services, and durables and their adding-up relation. Cells are weighted by non-durable expenditure mass. The estimated index slope is 0.4160, implying $\iota_D = 1 + 0.4160/(1 - \sigma) = 2.4901$. For this estimation step, a unit of holdings is assigned the constant annual cost $1 - \beta(1 - 0.10) = 0.136$. This auxiliary conversion is retained when the quantitative model uses $\delta_D = 0.168$; it is not the model’s time- and age-specific equilibrium user cost.

Depreciation and purchase targets. I set annual depreciation to $\delta_D = 0.168$ using the ratio of survey-weighted aggregate purchases in the preceding twelve months to aggregate current holdings in the 1993 asset data. This is a calibration proxy: purchases also reflect net accumulation, so the ratio is not a direct estimate from the stock accumulation law. The durable-purchase targets combine the age pattern of the broader durable and semi-durable expenditure basket with the aggregate level of vehicle and household-appliance purchases. I subtract a common amount of 4.413 percentage points from the broad age-group shares, subject to a 0.1 percent floor, so that their survey-expenditure-weighted mean equals 5.208 percent. The denominator for this level anchor is the original mapped four-sector non-durable expenditure plus broad durable spending. Non-durable manufacturing is adjusted so that total manufactured spending remains unchanged within each group. These are constructed IFLS-based targets, not the unadjusted age-specific spending shares of vehicles and appliances. The common shift preserves age differences between groups for which the floor does not bind.

Fitting durable tastes. I choose $\Omega_D(a)$ to fit the twelve age-group purchase shares and their aggregate level, conditional on the remaining demand parameters. To keep the profile smooth, I parameterize a positive desired-stock profile using seven log-stock values at ages 20, 30, 40, 50, 60, 72, and 89, interpolating between them. At each trial, the consumption-index and durable optimality conditions at a reference price and marginal-utility path map that profile into durable taste weights and the corresponding adjustment to non-durable taste weights. I then solve the full joint household problem, including the restriction that purchases cannot be negative, and compare the implied

³¹The conversion divides 11,345.83 rupiah by adult-weighted mean annual non-durable expenditure of 952,435.38 rupiah and multiplies by reference-model mean expenditure of 2.44928. This reference conversion is held fixed; it is not recomputed using expenditure in the final equilibrium. The expenditure fraction 0.01092 and the model-unit parameter 0.02918 are therefore different objects.

1993 purchases with the targets. The objective weights squared discrepancies by inverse variances from a 2,000-draw household bootstrap. The desired-stock profile is a way to parameterize tastes, not a separate observed stock target. The resulting taste weights are fixed before the demographic counterfactuals; they are not recalibrated under each alternative population path.

Initial stocks. Initial durable stocks are constructed rather than taken directly from observed household holdings. I expand the constructed 1993 purchase-share profile to single ages and accumulate the implied purchases recursively, using depreciation of 0.168 and a reference expenditure scale of one. This supplies initial conditions for cohorts already present at the start of the solution horizon in 1906. Stocks at the beginning of the reporting window, 1976, have consequently already evolved under the model; they are not reset to the constructed profile. Cohorts entering subsequently have no inherited stock. Table 46 summarizes the scalar demand parameters.

Table 46: Household demand parameters

Parameter	Baseline	Durable extension
σ	0.7208	0.7208
ι_A	0.2364	0.4225
ι_M	1	1
ι_L	0.9331	1.5848
ι_H	1.4673	2.4759
ι_D	—	2.4901
δ_D	—	0.1680
$\bar{d}$ (model units)	—	0.02918

Notes: The baseline includes durable purchases in manufactured consumption without modeling stocks. In the extension, the four ι_j refer to non-durable goods. The manufacturing normalization is $\iota_M = 1$ in both models. Age-specific taste weights are calibrated separately as described in the text.

K.1.4 Sectoral productivity

Real value added per worker is constructed from the linked GGDC sectoral accounts through 2018. The manufacturing aggregate includes utilities and construction but excludes mining. I extrapolate raw manufacturing labor productivity through 2025 using its average annual log change over 2013–2018. To distinguish sectoral productivity from the efficiency of workers employed there, I divide the normalized labor-productivity series by effective labor per manufacturing worker in a reference model allocation:

$$A_{M,t} = \frac{LP_{M,t}}{LP_{M,1993}} \frac{q_{M,1993}^R}{q_{M,t}^R}, \quad q_{M,t}^R = \frac{N_{M,t}^R}{L_{M,t}^R}. \quad (117)$$

Here, $LP_{M,t}$ is observed or extrapolated real value added per employed worker, and $q_{M,t}^R$ is effective labor per employed worker in the reference allocation. This normalizes $A_{M,1993} = 1$. The adjustment uses the reference worker allocation, rather than being recomputed separately for each final model. The resulting manufacturing-productivity path is common to the baseline and durable extension.

Conditional on this path, I adjust the annual productivity paths of agriculture and the two service sectors to fit their employment shares, solving the transition equilibrium after each adjustment. The annual employment targets use GGDC through 2017, followed by the labor-force-survey and WDI extension described in Appendix K.2. These relative productivity paths are fitted separately in the two models. All productivity paths are then held fixed in the demographic counterfactuals.

K.1.5 Trade flows, world prices, and trade parameters

Sources and sectoral coverage. Merchandise flows use CEPII’s BACI database for 1995–2024 and the NBER–United Nations World Trade Flows for earlier years, using the latter’s growth rates to extend the BACI level backward from 1995. The model’s manufacturing aggregate adds utilities and construction using OECD Inter-Country Input–Output (ICIO) tables and excludes mining. Gross service flows combine WTO International Trade Statistics and balance-of-payments data for 1981–1994, OECD–WTO Balanced Trade in Services (BaTIS) under BPM5 for 1995–2004, and BaTIS under BPM6 for 2005–2024; earlier service flows are backcast. OECD ICIO sectoral accounts, UNIDO manufacturing accounts, and WDI national accounts supply the production denominators accompanying the gross-flow measures.

Gross trade and value-added trade serve different roles. Gross import penetration determines the fixed home-bias weights. Value-added trade targets use the 2025 edition of OECD ICIO over 1995–2022, extended outside that window using the gross-flow series. Let $FVA_{j,t}$ denote foreign value added originating in sector j and embodied in Indonesian final demand, $VAX_{j,t}$ Indonesian value added originating in sector j and embodied in foreign final demand, and $VA_{j,t}$ domestic sectoral value added. The import-share target is

$$s_{j,t}^m = \frac{FVA_{j,t}}{VA_{j,t} - VAX_{j,t} + FVA_{j,t}}. \quad (118)$$

This is an import share of value-added absorption, not a gross import-to-output ratio.

World-price proxies. Manufacturing’s international price is normalized to one. Agricultural prices use the World Bank agricultural commodity-price index relative to a dollar-denominated manufacturing-price benchmark, constructed from the GGDC manufacturing deflator and WDI official exchange rates. For services, Indonesian sectoral deflators relative to manufacturing provide price proxies; these are not direct observations of international service prices. All series equal one in 1993. After GGDC coverage ends, the service-price proxies and manufacturing-price conversion factor are held at their 2018 values; the agricultural commodity-price series continues to vary.

Import and export parameters. I set $\eta = 8$ and fix ω_j at one minus average gross import penetration over 1995–2024. In sector order (A, M, L, H) , the resulting weights are $(0.8505, 0.8704, 0.9400, 0.8863)$. Given these weights, I invert (100) to recover annual import-price wedges from the value-added import shares at reference producer prices and exchange rates. The reference producer prices use relative GGDC deflators, with manufacturing’s price and the exchange rate normalized to one. These are effective price wedges and need not exceed one.

Agriculture and manufacturing face price-taking export opportunities. Services face downward-sloping foreign demand with elasticity $\varepsilon = 8$. I recover each service sector’s annual export-demand shifter $\Xi_{j,t}$ from its export-to-GDP target at the reference allocation. Export costs are normalized to $\tau_j^X = 1$. The import wedges and service export shifters are held fixed when fitting the remaining model parameters and in demographic counterfactuals. Their inversion reproduces the trade targets at the reference allocation, not necessarily at the final equilibrium prices and output levels.

External balance. Within the ICIO window, the empirical external-balance ratio is net value-added exports of the four modeled sectors divided by WDI national GDP. I smooth it with a centered three-year moving average and join it outside that window to the earlier mining- and oil-adjusted external-balance series, applying level adjustments at the joins. The model applies this empirical ratio, $\mathcal{W}_t$, to its own nominal production value $\sum_j p_{j,t}^d Y_{j,t}$. The empirical denominator is therefore national GDP, whereas the model denominator is the value of production in its four sectors. Negative values imply a trade deficit.

K.2 Employment shares

Figure 37 compares sectoral employment shares over time. Both models reproduce the sustained decline in agriculture and the expansion of manufacturing and services. The annual employment-share fit statistics in Table 5 are calculated over the reporting period, 1976–2025. The data use GGDC through 2017; the post-2017 data extension uses labor force survey and WDI data, with interpolation in 2020–2022 and extrapolation of non-agricultural shares in 2024–2025.

The model’s age profiles in Figure 6 pool employment over 1976–2025. The empirical profiles pool the 1993, 1997, and 2000 IFLS waves, excluding mining and expressing employment shares over the four model sectors. Each sector’s IFLS age profile is shifted by a constant so that its average, using IFLS age weights, matches the GGDC mean for those survey years. This adjustment preserves percentage-point differences across ages. The manufacturing decline reported in Table 5 is measured from each profile’s highest share among ages 20–25 to age 64: the highest share occurs at age 20 in both models and at age 24 in the data. The change per ten years of age is ten times the least-squares slope of the manufacturing share of total employment on single age over ages 20–64, computed on the same pooled profiles as the decline.

The figure includes a retirement correction based on workers' previous sector: the high-skill-services share is raised by 0.8 percentage points at ages 55–59 and 1.3 points at ages 60–64, with the other sectors' shares reduced proportionately. The table's correlations use these adjusted profiles; its manufacturing decline and slope use the empirical profile before this correction.




Figure 37: Sectoral employment shares over time. Markers show the data; solid lines show the baseline and dashed lines the durable-goods extension. The shaded years extend the GGDC series using labor force survey and WDI data, with interpolation and extrapolation as described in the text.

K.3 Expenditure by age

Figure 38 compares model and IFLS expenditure in 1993. Households are grouped by the mean age of members aged 20 and above, using twelve age groups: under 25, five-year groups from 25–29 to 70–74, and 75 and above. Model expenditure is aggregated by individual age over the same groups, using population weights. The left panel shows manufactured goods as a share of total spending, in percentage-point deviations from each series' unweighted mean across the twelve groups. Manufactured goods comprise manufactures in the baseline, and non-durable manufactures and durable purchases in the durable-goods extension. The right panel shows durable purchases as a share of total spending in the IFLS-based calibration targets and in the extension.

The durable-purchase targets retain the broader durable basket's age pattern but shift its level to the aggregate spending share of vehicles and household appliances, with a floor of 0.1 percent in each age group. These are constructed targets, not the unadjusted IFLS shares underlying the young-old differences in Appendix F. The durable-purchase correlation in Table 5 uses this target profile. Its

young–old manufactured-spending difference compares the simple averages of the two group shares below 30 and the four group shares at ages 60 and above.




Figure 38: Spending by age, models and 1993 IFLS. Solid lines show the baseline and dashed lines the durable-goods extension. Left: manufactured spending as a share of total expenditure, in percentage-point deviations from each series’ mean across the twelve age groups. Right: durable purchases as a share of total expenditure, compared with the IFLS-based calibration targets described in the text. Age groups use mean adult household age in the data and individual age in the models.

K.4 Manufacturing persistence

Manufacturing persistence measures the share of workers initially employed in manufacturing who remain in that sector. Model one-year transitions are pooled over origin years 1976–2024, so all transitions end within the reporting window. The IFLS comparison uses the four consecutive survey intervals from 1993 to 2014. Retention over an interval of Δ years is annualized by taking its $1/\Delta$ power; this converts survey endpoints into an annual rate rather than observing annual transitions directly. Figure 39 shows that both models overpredict persistence in every reported age group. These transition probabilities are not targeted in either calibration.

L Counterfactual construction and additional results

This appendix documents the counterfactual exercises of Section 5: the demographic paths, the restricted economies and the decomposition of demographic gaps, the alternative export assumptions, comparisons with the durable-goods extension, and the schooling and welfare calculations. Unless otherwise stated, results refer to the baseline model without explicit durable accumulation.

L.1 Demographic paths

Each counterfactual replaces the Indonesian age distribution $\pi_t(a)$, $a \in \mathcal{A}$, with that of another country at a comparable stage of structural transformation, and recomputes the cohort ratios $s_t(a)$ from (6). All other primitives are those of the calibrated Indonesian economy, and the economy is




Figure 39: Annual manufacturing persistence by age at the start of the transition. Solid and dashed lines show the baseline and the durable-goods extension, respectively. Markers show annualized IFLS retention between survey waves; these are not directly observed one-year transitions.

re-solved in general equilibrium under each path.

The paths are aligned at a point when agriculture’s share of employment is about 60 percent. For Indonesia and Thailand, I use the last year with a share of at least 60 percent: 1976 and 1992, respectively. For the United States, I use 1850, the first year of the census series, when the agricultural share is already slightly below 60 percent. These years are mapped to 1976, so the reporting window 1976–2025 corresponds to 1850–1899 on the U.S. path and to 1992–2041 on the Thai path.

The Indonesian and Thai paths are single-year age distributions from UN WPP 2024, which reports estimates through 2023 and the medium-variant projection thereafter. The Thai path therefore relies on projections from model year 2008 onward, and the Indonesian path in model years 2024 and 2025. The U.S. path is built from the IPUMS USA Full Count census schedules of 1850, 1860, 1870, 1880, and 1900, with the enslaved-person schedules added in 1850 and 1860 (Appendix A). Age heaping is smoothed within five-year age groups, preserving each group’s total, and the years between censuses are interpolated linearly. Because the 1890 schedules no longer exist,³² the U.S. path in model years 2007–2025 is interpolated between the 1880 and 1900 censuses.

³²The originals were destroyed after the 1921 Commerce Building fire. See the National Archives, *Prologue Magazine*, Spring 1996, “First in the Path of the Firemen”: <https://www.archives.gov/publications/prologue/1996/spring/1890-census-1.html>.

L.2 Construction of the restricted economies

The restricted economies remove ability aging, age-dependent demand, or both, as described below. Each is re-solved in general equilibrium under every demographic path and each set of export assumptions.

Sorting restriction. I set $\tilde{z}_j(a) = \bar{z}$ at every age and in every sector, including agriculture. I first take the geometric mean of the four sectoral profiles at each age and normalize the resulting age profile to have an unweighted mean of one over working ages. The common level $\bar{z}$ is the average of this normalized profile under the 1993 Indonesian working-age distribution. Permanent ability heterogeneity, the sectoral ability weights, the schooling loadings, the cohort schooling distributions, and the coefficients $\kappa_A(a)$ of the cost of leaving agriculture are unchanged. Workers therefore continue to sort across sectors according to their abilities and equilibrium wages, and continue to face the cost of leaving agriculture. The restriction removes the age variation in sectoral efficiency generated by ability aging; it does not remove worker selection, cohort schooling differences, or mobility frictions.

Age-dependent demand restriction. In each specification, I replace the age-specific sectoral taste weights $\Omega(a, j)$ with a common vector $\bar{\Omega}_j$. I construct this vector from that specification's full Indonesian economy in 1993, rather than by averaging its age-specific sectoral weights.

Let $w_a = \pi_{1993}^{\text{IDN}}(a)$ denote the Indonesian adult age weights, normalized over ages 20–89, and let a superscript $\star$ denote the reference allocation and prices. Define

$$C^\star = \sum_a w_a C^\star(a), \quad E^\star = \sum_a w_a E^\star(a), \quad q_j^\star = \frac{\sum_a w_a p_j^{Q,\star} c_j^\star(a)}{E^\star}.$$

Here, $E^\star$ is non-durable expenditure per person, and $q_j^\star$ is sector j 's share of aggregate non-durable expenditure, rather than the population-weighted average of individual expenditure shares.

In the baseline without explicit durable accumulation, set $S^\star = 1$. In the durable-goods extension, I also replace the durable taste weight with its 1993 population-weighted mean,

$$\bar{\Omega}_D = \sum_a w_a \Omega_D(a), \tag{119}$$

and define

$$D^\star = \sum_a w_a D^\star(a), \quad T_D^\star = \bar{\Omega}_D^{1/\sigma} [(\bar{d} + D^\star)(C^\star)^{-\iota_D}]^{(\sigma-1)/\sigma}, \quad S^\star = 1 - T_D^\star.$$

The common sectoral weights in either specification are

$$\bar{\Omega}_j = \frac{q_j^\star (E^\star)^{1-\sigma} (S^\star)^\sigma}{(p_j^{Q,\star})^{1-\sigma} (C^\star)^{(1-\sigma)\iota_j}}. \tag{120}$$

At the reference prices, consumption index, and durable stock, these weights reproduce $E^\star$ and the aggregate shares $q_j^\star$. All other preference parameters are unchanged.

The resulting weights are used at every age and on every demographic path. In the durable-goods extension, the dynasty additionally chooses a common durable stock per member, $\bar{D}_t$, and common purchases, $\bar{x}_t$, so $D_t(a) = \bar{D}_t$ and $x_t(a) = \bar{x}_t$ at every age. The common stock and purchases remain endogenous.

Inherited holdings are pooled across current members. Entrants bring no stock, and the holdings of exiting members are not transferred to the pool. With the population normalization $\sum_a \pi_t(a) = 1$, common inherited holdings $\bar{B}_t$ and purchases satisfy

$$\bar{B}_t = (1 - \delta_D)[1 - \pi_t(a_0)]\bar{D}_{t-1}, \quad \bar{D}_t = \bar{B}_t + \bar{x}_t, \quad \bar{x}_t \geq 0, \quad (121)$$

where $a_0 = 20$. At the first date of the solution horizon, inherited holdings are

$$\bar{B}_0 = (1 - \delta_D) \sum_{a>a_0} \pi_0(a) D_{-1}(a - 1).$$

Entry dilutes inherited holdings per member, so the age structure still affects aggregate purchases through the share of entrants; what the restriction removes is the accumulation of separate stocks over each cohort's life cycle.

For a given $\bar{D}_t$, the consumption-index condition (109) continues to apply. The marginal benefit of a larger common stock is the population-weighted average of members' marginal benefits,

$$\overline{\text{MV}}_t = \sum_{a \in \mathcal{A}} \pi_t(a) \lambda_t E_t(a) \frac{\xi_t(a)}{\bar{d} + \bar{D}_t}. \quad (122)$$

Let $v_t = \beta(1 - \delta_D)[1 - \pi_{t+1}(a_0)]$ denote the discounted share of a unit bought today that remains in the pool next period, and let $\bar{\chi}_t \geq 0$ be the current-value multiplier on $\bar{x}_t \geq 0$. The optimality conditions are

$$\begin{aligned} \overline{\text{MV}}_t - \lambda_t p_{M,t}^Q + v_t \lambda_{t+1} p_{M,t+1}^Q + \bar{\chi}_t - v_t \bar{\chi}_{t+1} &= 0, \\ \bar{x}_t \geq 0, \quad \bar{\chi}_t \geq 0, \quad \bar{\chi}_t \bar{x}_t &= 0, \end{aligned} \quad (123)$$

with all next-period terms omitted at the terminal date. When purchases are positive today and next period, (123) reduces to $\overline{\text{MV}}_t = \lambda_t p_{M,t}^Q - v_t \lambda_{t+1} p_{M,t+1}^Q$, the counterpart of (110). The common stock thus remains an intertemporal purchase decision with non-negative purchases.

The economies solved. For each specification, demographic path, and set of export assumptions, I solve four economies: *Full*, with neither restriction; *Sorting only*, with the age-dependent demand restriction; *Demand only*, with the sorting restriction; and *Neither*, with both restrictions. Apart from these changes, the calibrated primitives, including productivity paths, are unchanged, and every economy is re-solved in general equilibrium.

The restricted economies retain permanent ability heterogeneity, cohort schooling, the parameters governing agricultural exit costs, and the headcount margin. Workers continue to choose sectors, and demand continues to respond to income and prices. In the durable-goods extension, the common stock

also retains replacement demand and the dilution of inherited holdings by entrants. The demographic gap in *Neither* is therefore not a headcount-only effect.

L.3 Sectoral employment under the demographic paths

Table 47 reports the full-model allocation of employment across all four sectors under the baseline export assumptions. In 2025, the older Thai path has larger agricultural and low-skill-services shares than the U.S. path, and smaller manufacturing and high-skill-services shares, in both specifications.

Table 47: Sectoral employment under the three demographic paths

Demographic path	A	M	L	H
<i>Baseline model</i>				
<i>1976</i>				
Indonesia	67.17	8.84	16.97	7.02
Thailand	66.62	9.25	17.12	7.01
U.S.	66.45	9.69	16.94	6.92
<i>2025</i>				
Indonesia	27.65	20.90	36.94	14.51
Thailand	28.94	18.38	39.45	13.23
U.S.	26.19	23.25	36.00	14.56
<i>Mean, 1976–2025</i>				
Indonesia	45.52	16.13	27.51	10.84
Thailand	46.68	14.19	28.57	10.55
U.S.	44.94	17.06	27.18	10.83
<i>Durable-goods extension</i>				
<i>1976</i>				
Indonesia	65.00	9.28	18.45	7.26
Thailand	64.54	9.73	18.54	7.20
U.S.	64.67	10.02	18.28	7.03
<i>2025</i>				
Indonesia	27.63	21.21	37.06	14.11
Thailand	28.80	18.70	39.50	13.00
U.S.	26.20	23.52	36.24	14.04
<i>Mean, 1976–2025</i>				
Indonesia	45.05	16.49	27.92	10.55
Thailand	46.20	14.57	28.85	10.38
U.S.	44.53	17.39	27.60	10.48

Notes: Percent of total employment, full economies with price-taking exports in agriculture and manufacturing. A: agriculture; M: manufacturing, utilities and construction; L: low-skill services; H: high-skill services. Total employment excludes mining. Means give equal weight to the 50 years 1976–2025.

L.4 Decomposition of demographic gaps

The same four-economy decomposition applies to the baseline model and the durable-goods extension. For a pair of demographic paths, let G^F denote the difference in an outcome in the full model, G^{-s} the gap without ability aging, G^{-d} the gap without age-dependent demand, and $G^{-s,-d}$ the gap with both restrictions. These correspond to the *Full*, *Demand only*, *Sorting only*, and *Neither* economies. Every comparison uses the same pair of demographic paths, outcome definition, and reporting window within a given specification and set of export assumptions.

The contribution of each mechanism can depend on whether the other operates. I therefore average its contribution with the other mechanism present and absent, using a two-factor Shapley decomposition (Shorrocks, 2013):

$$\begin{aligned}\mathcal{S} &= \frac{1}{2} \left[(G^F - G^{-s}) + (G^{-d} - G^{-s,-d}) \right], \\ \mathcal{D} &= \frac{1}{2} \left[(G^F - G^{-d}) + (G^{-s} - G^{-s,-d}) \right].\end{aligned}\tag{124}$$

The contribution of sorting, $\mathcal{S}$, averages the change in the demographic gap when ability aging is introduced with age-dependent demand present and absent. The contribution of demand, $\mathcal{D}$, is constructed analogously. In the baseline it captures age-specific sectoral taste weights; in the extension it also captures age-specific durable taste weights and cohort-specific durable accumulation. The decomposition satisfies

$$G^F = \mathcal{S} + \mathcal{D} + G^{-s,-d},$$

so the full-model gap equals the two contributions plus the gap in the *Neither* economy in both specifications. The demographic gap in *Neither* is therefore not a headcount-only effect.

Tables 48 and 49 report the four gaps entering the decomposition in the baseline and the durable-goods extension, respectively, under both sets of export assumptions. Tables 7, 8, and 56 report the resulting contributions. Removing ability aging substantially narrows the manufacturing gaps in both specifications. The demand contribution to manufacturing employment is small under price-taking goods exports and larger under downward-sloping export demand.

Table 48: Demographic gaps in the full and restricted economies of the baseline model

Pair	Price-taking exports in agriculture and manufacturing				Downward-sloping export demand in all sectors			
	Full	Sorting	Demand	Neither	Full	Sorting	Demand	Neither
		only	only			only	only	
<i>Panel A. Peak manufacturing employment share (percentage points)</i>								
U.S.–Thailand	4.84	4.66	1.48	1.32	3.11	2.51	0.98	0.35
Indonesia–Thailand	2.61	2.47	1.03	0.89	1.64	1.19	0.60	0.13
<i>Panel B. Mean manufacturing employment share (percentage points)</i>								
U.S.–Thailand	2.86	2.82	0.68	0.60	1.99	1.72	0.55	0.27
Indonesia–Thailand	1.94	1.90	0.49	0.44	1.32	1.13	0.36	0.16
<i>Panel C. Cumulative real output-per-worker growth (percentage points)</i>								
U.S.–Thailand	19.5	20.7	8.6	9.6	21.0	22.6	9.1	10.4
Indonesia–Thailand	19.9	21.1	10.3	11.7	20.5	22.0	10.5	12.1
<i>Panel D. Cumulative real output-per-capita growth (percentage points)</i>								
U.S.–Thailand	99.6	100.7	93.5	94.4	100.9	102.3	93.6	94.8
Indonesia–Thailand	93.1	94.3	88.2	89.5	93.6	94.9	88.1	89.4

Notes: Each entry is the difference between two demographic paths in the baseline model without explicit durable accumulation. The four economies are defined in Appendix L.2: “Sorting only” removes age-dependent demand, “Demand only” removes ability aging, and “Neither” imposes both restrictions. The contributions in Tables 7 and 8 are obtained from these gaps by (124). Outcomes are defined as in Table 6.

Table 49: Demographic gaps in the full and restricted economies of the durable-goods extension

Pair	Price-taking exports in agriculture and manufacturing				Downward-sloping export demand in all sectors			
	Full	Sorting	Demand	Neither	Full	Sorting	Demand	Neither
		only	only			only	only	
<i>Panel A. Peak manufacturing employment share (percentage points)</i>								
U.S.–Thailand	4.76	4.65	1.47	1.39	2.97	2.54	0.88	0.49
Indonesia–Thailand	2.62	2.51	1.00	0.89	1.65	1.29	0.55	0.19
<i>Panel B. Mean manufacturing employment share (percentage points)</i>								
U.S.–Thailand	2.82	2.83	0.69	0.62	1.88	1.79	0.46	0.36
Indonesia–Thailand	1.92	1.92	0.52	0.45	1.26	1.17	0.32	0.22
<i>Panel C. Cumulative real output-per-worker growth (percentage points)</i>								
U.S.–Thailand	22.0	23.8	9.7	11.3	23.8	26.0	10.2	12.2
Indonesia–Thailand	22.3	24.2	12.2	14.0	22.9	24.9	12.4	14.4
<i>Panel D. Cumulative real output-per-capita growth (percentage points)</i>								
U.S.–Thailand	105.4	107.0	98.8	100.3	107.0	108.9	99.1	100.8
Indonesia–Thailand	98.6	100.1	93.9	95.5	99.0	100.6	93.8	95.5

Notes: Each entry is the difference between two demographic paths in the model with durable goods. The four economies are defined in Appendix L.2: “Sorting only” removes age-dependent demand, “Demand only” removes ability aging, and “Neither” imposes both restrictions. The contributions in Table 56 are obtained from these gaps by (124). Outcomes are defined as in Table 6.

Sectoral decomposition. Tables 50–53 apply (124) to each sector’s employment-share gap in 2025 and to its mean gap over 1976–2025. The gaps are the first demographic path minus the Thai path. Summing across sectors gives zero for each column.

In the baseline model with price-taking goods exports, ability aging contributes 2.06 percentage points to the Thai path’s 3.45-point greater low-skill-services share in 2025. The U.S. path’s high-skill-services advantage is 2.15 points in the *Neither* economy; sorting and age-dependent demand reduce it by 0.32 and 0.50 points, respectively, leaving a full-model gap of 1.33 points.

Table 50: Sectoral decomposition: baseline model, price-taking goods exports

Pair	Sector	Full	Sorting	Demand	Neither
<i>2025</i>					
U.S.–Thailand	A	-2.75	-1.00	0.66	-2.41
	M	4.86	3.38	0.16	1.32
	L	-3.45	-2.06	-0.32	-1.06
	H	1.33	-0.32	-0.50	2.15
Indonesia–Thailand	A	-1.28	-0.50	0.29	-1.07
	M	2.51	1.49	0.14	0.89
	L	-2.51	-0.86	-0.48	-1.17
	H	1.28	-0.12	0.05	1.35
<i>Mean, 1976–2025</i>					
U.S.–Thailand	A	-1.74	-0.68	0.55	-1.62
	M	2.86	2.20	0.06	0.60
	L	-1.40	-1.26	-0.12	-0.03
	H	0.28	-0.27	-0.49	1.04
Indonesia–Thailand	A	-1.16	-0.41	0.33	-1.07
	M	1.94	1.46	0.04	0.44
	L	-1.07	-0.87	-0.11	-0.09
	H	0.29	-0.17	-0.25	0.71

Notes: Percentage points. Sorting and Demand are Shapley contributions; Neither is the gap with both restrictions imposed. Components add to Full, up to rounding. Sector definitions and annual averaging follow Table 47.

Table 51: Sectoral decomposition: baseline model, downward-sloping export demand

Pair	Sector	Full	Sorting	Demand	Neither
<i>2025</i>					
U.S.–Thailand	A	-1.89	-0.41	0.44	-1.91
	M	3.11	2.14	0.62	0.35
	L	-2.56	-1.41	-0.56	-0.60
	H	1.34	-0.32	-0.50	2.16
Indonesia–Thailand	A	-0.80	-0.25	0.14	-0.68
	M	1.52	0.94	0.45	0.13
	L	-2.02	-0.58	-0.63	-0.81
	H	1.29	-0.12	0.05	1.36
<i>Mean, 1976–2025</i>					
U.S.–Thailand	A	-1.38	-0.37	0.52	-1.53
	M	1.99	1.44	0.27	0.27
	L	-0.91	-0.82	-0.29	0.20
	H	0.31	-0.25	-0.51	1.06
Indonesia–Thailand	A	-0.87	-0.24	0.31	-0.94
	M	1.32	0.96	0.20	0.16
	L	-0.75	-0.57	-0.24	0.06
	H	0.30	-0.16	-0.26	0.72

Notes: Percentage points. Sorting and Demand are Shapley contributions; Neither is the gap with both restrictions imposed. Components add to Full, up to rounding. Sector definitions and annual averaging follow Table 47.

Table 52: Sectoral decomposition: durable-goods extension, price-taking goods exports

Pair	Sector	Full	Sorting	Demand	Neither
<i>2025</i>					
U.S.–Thailand	A	-2.60	-1.15	1.09	-2.53
	M	4.82	3.36	0.07	1.39
	L	-3.25	-1.91	-0.52	-0.82
	H	1.03	-0.29	-0.63	1.96
Indonesia–Thailand	A	-1.17	-0.59	0.56	-1.14
	M	2.50	1.50	0.11	0.89
	L	-2.44	-0.80	-0.65	-0.99
	H	1.10	-0.11	-0.02	1.23
<i>Mean, 1976–2025</i>					
U.S.–Thailand	A	-1.67	-0.79	0.73	-1.62
	M	2.82	2.17	0.03	0.62
	L	-1.25	-1.13	-0.18	0.05
	H	0.10	-0.25	-0.59	0.94
Indonesia–Thailand	A	-1.15	-0.54	0.46	-1.07
	M	1.92	1.43	0.03	0.45
	L	-0.94	-0.74	-0.18	-0.02
	H	0.17	-0.15	-0.31	0.64

Notes: Percentage points. Sorting and Demand are Shapley contributions; Neither is the gap with both restrictions imposed. Components add to Full, up to rounding. Sector definitions and annual averaging follow Table 47.

Age composition and manufacturing employment. The aggregate manufacturing share can differ because the age composition of employment differs or because workers of the same age have different sectoral employment shares. Let $L_t^r(a)$ denote employment headcounts at age a on demographic path r , and $L_{M,t}^r(a)$ the corresponding manufacturing headcount. Define

$$p_t^r(a) = \frac{L_t^r(a)}{\sum_{b \in \mathcal{A}^W} L_t^r(b)}, \quad m_t^r(a) = \frac{L_{M,t}^r(a)}{L_t^r(a)}, \quad s_{M,t}^r = \sum_{a \in \mathcal{A}^W} p_t^r(a) m_t^r(a),$$

Table 53: Sectoral decomposition: durable-goods extension, downward-sloping export demand

Pair	Sector	Full	Sorting	Demand	Neither
<i>2025</i>					
U.S.–Thailand	A	-1.73	-0.59	0.95	-2.09
	M	2.98	2.10	0.38	0.49
	L	-2.30	-1.22	-0.70	-0.38
	H	1.05	-0.29	-0.63	1.98
Indonesia–Thailand	A	-0.70	-0.35	0.45	-0.80
	M	1.49	0.94	0.35	0.19
	L	-1.90	-0.48	-0.79	-0.63
	H	1.12	-0.10	-0.02	1.24
<i>Mean, 1976–2025</i>					
U.S.–Thailand	A	-1.26	-0.47	0.85	-1.64
	M	1.88	1.42	0.10	0.36
	L	-0.76	-0.72	-0.35	0.31
	H	0.13	-0.23	-0.60	0.96
Indonesia–Thailand	A	-0.83	-0.34	0.53	-1.02
	M	1.26	0.95	0.09	0.22
	L	-0.62	-0.46	-0.30	0.14
	H	0.19	-0.14	-0.32	0.65

Notes: Percentage points. Sorting and Demand are Shapley contributions; Neither is the gap with both restrictions imposed. Components add to Full, up to rounding. Sector definitions and annual averaging follow Table 47.

where $\mathcal{A}^W$ covers ages 20–64. Thus, $p_t^r(a)$ is an age group’s share of total employment and $m_t^r(a)$ is manufacturing’s share within that age group. Let Δ denote the U.S.-path value minus the Thai-path value, and let a bar denote the arithmetic average of the two paths’ values. Then

$$\Delta s_{M,t} = \underbrace{\sum_{a \in \mathcal{A}^W} \Delta p_t(a) \bar{m}_t(a)}_{\text{different workforce age composition}} + \underbrace{\sum_{a \in \mathcal{A}^W} \bar{p}_t(a) \Delta m_t(a)}_{\text{different manufacturing shares within ages}}. \quad (125)$$

The first term applies the same manufacturing employment profile, averaged across the two paths, to their different workforce age distributions. The second applies the same age distribution to differences in manufacturing employment shares at each age.

In the full baseline model with price-taking goods exports, workforce age composition accounts for 4.15 percentage points of the 4.86-point U.S.–Thailand manufacturing employment-share difference in 2025; differences within ages account for 0.72 points. This is an accounting comparison of the saved equilibrium allocations, not a further counterfactual or a Shapley contribution.

L.5 Downward-sloping export demand in all sectors

Section 5.4 re-solves the counterfactual economies with downward-sloping export demand in every sector. Exports follow the foreign demand schedule (30),

$$X_{j,t} = \Xi_{j,t} \left(\frac{(p_{j,t}^d/e_t) \tau_j^X}{p_{j,t}^*} \right)^{-\varepsilon}, \quad j \in \mathcal{J},$$

with $\varepsilon = 8$. For the two service sectors the shifters $\Xi_{j,t}$ are those of the baseline calibration. For agriculture and manufacturing, $\Xi_{j,t}$ is chosen separately in each specification to reproduce exports

at the equilibrium prices of that specification’s Indonesian full model in every year; in years with no exports of a good in that reference economy, foreign demand for it is effectively zero. Manufacturing remains the numeraire and the exchange rate adjusts to satisfy external balance (31). Within each specification, the same shifters are used on every demographic path and in every restricted economy. All other primitives are unchanged.

By construction, re-solving the Indonesian full model with these export-demand schedules reproduces its price-taking allocation within each specification. In a demographic counterfactual, the prices faced by foreign buyers of agricultural and manufactured goods respond to the quantities exported, and the exchange rate is determined by external balance rather than by the world price of manufactures. Foreign demand remains effectively zero in years when the reference economy does not export the good. In the baseline without explicit durable accumulation, the full-model manufacturing peak is 19.37 percent on the Thai path and 22.48 percent on the U.S. path, compared with 18.41 and 23.25 percent under price-taking exports. Table 8 reports the baseline decomposition; Table 56 reports the corresponding results for the durable-goods extension.

Sectoral allocation and manufacturing exports. Table 54 reports the change in each sector’s employment share when goods exports face downward-sloping demand. In 2025, manufacturing’s share falls on the U.S. path and rises on the Thai path. Agriculture and low-skill services rise on the U.S. path and fall on the Thai path, while high-skill-services shares change little. The full Indonesian allocation is preserved by construction.

Table 55 reports the corresponding manufacturing output and export allocation. Manufacturing output per adult falls on the U.S. path and rises on the Thai path in 2025. The share of manufacturing output exported also falls on the U.S. path and rises on the Thai path.

L.6 Comparison with the durable-goods extension

The durable-goods extension of Section 3.9 is calibrated separately to the same employment targets and household expenditure data as the baseline. In the baseline without explicit durable accumulation, purchases of durable and semi-durable goods are included in manufactured consumption, so the age-specific taste weights carry the age profile of total manufactured spending. The income-effect parameters, age-specific taste weights, worker parameters, and productivity paths of agriculture and the two service sectors are calibrated separately in the two specifications. The substitution elasticity, manufacturing productivity, and price-taking trade-calibration parameters are common to both. The mean absolute gap between model and data in Indonesia’s annual sectoral employment shares is 0.38 percentage points in the baseline and 0.17 in the extension.

Table 56 reports the extension’s demographic gaps and mechanism contributions under both export assumptions. With price-taking agricultural and manufacturing exports, sorting accounts for 3.27 of

Table 54: Sectoral employment changes under downward-sloping export demand

Demographic path	A	M	L	H
<i>Baseline model</i>				
<i>2025</i>				
Indonesia	0.00	0.00	0.00	0.00
Thailand	-0.49	0.99	-0.50	0.00
U.S.	0.38	-0.77	0.39	0.00
<i>Mean, 1976–2025</i>				
Indonesia	0.00	0.00	0.00	0.00
Thailand	-0.29	0.62	-0.32	-0.01
U.S.	0.07	-0.26	0.17	0.01
<i>Durable-goods extension</i>				
<i>2025</i>				
Indonesia	0.00	0.00	0.00	0.00
Thailand	-0.47	1.02	-0.53	-0.01
U.S.	0.39	-0.82	0.42	0.01
<i>Mean, 1976–2025</i>				
Indonesia	0.00	0.00	0.00	0.00
Thailand	-0.33	0.66	-0.32	-0.02
U.S.	0.09	-0.28	0.18	0.01

Notes: Percentage-point change from price-taking agricultural and manufacturing exports to downward-sloping export demand, within each full-model demographic path. Sector definitions and annual averaging follow Table 47.

Table 55: Manufacturing production and exports under the two export assumptions

Demographic path	Output index	Export/output (%)	
	Downward-sloping demand	Price-taking	Downward-sloping demand
<i>Baseline model</i>			
<i>2025</i>			
Indonesia	100.00	41.12	41.12
Thailand	104.28	41.79	47.46
U.S.	97.32	44.22	40.40
<i>Mean, 1976–2025</i>			
Indonesia	100.00	25.98	25.98
Thailand	103.59	23.10	28.51
U.S.	98.58	27.56	25.62
<i>Durable-goods extension</i>			
<i>2025</i>			
Indonesia	100.00	33.83	33.83
Thailand	104.29	31.65	38.50
U.S.	97.20	38.46	33.90
<i>Mean, 1976–2025</i>			
Indonesia	100.00	21.44	21.44
Thailand	103.68	16.88	23.04
U.S.	98.48	23.44	21.43

Notes: Full economies. Manufacturing output per adult under price-taking exports is normalized to 100 for each path and period. The mean-period index is the ratio of mean outputs. Export/output is the percentage of manufacturing output exported; its mean is the arithmetic average of annual ratios.

the 4.76 points of the U.S.–Thailand peak gap, and age-dependent demand contributes 0.09 points. With downward-sloping export demand in all sectors, the gap is 2.97 points, of which sorting accounts for 2.07 and demand for 0.41. The full-model manufacturing peaks under downward-sloping export demand are 19.73 percent on the Thai path and 22.70 percent on the U.S. path, compared with 18.76 and 23.52 percent under price-taking exports. Sorting accounts for 69 and 70 percent of the peak gap under the two export assumptions, respectively, compared with 69 percent under either set of

assumptions in the baseline.

Table 56: Decomposition of demographic gaps in the extended model with durable goods under different export assumptions

Pair	Price-taking exports in agriculture and manufacturing				Downward-sloping export demand in all sectors			
	Full	Sorting	Demand	Neither	Full	Sorting	Demand	Neither
<i>Panel A. Peak manufacturing employment share (percentage points)</i>								
U.S.–Thailand	4.76	3.27	0.09	1.39	2.97	2.07	0.41	0.49
Indonesia–Thailand	2.62	1.62	0.11	0.89	1.65	1.10	0.36	0.19
<i>Panel B. Mean manufacturing employment share (percentage points)</i>								
U.S.–Thailand	2.82	2.17	0.03	0.62	1.88	1.42	0.10	0.36
Indonesia–Thailand	1.92	1.43	0.03	0.45	1.26	0.95	0.09	0.22
<i>Panel C. Cumulative real output-per-worker growth (percentage points)</i>								
U.S.–Thailand	22.0	12.4	-1.7	11.3	23.8	13.7	-2.1	12.2
Indonesia–Thailand	22.3	10.1	-1.9	14.0	22.9	10.5	-2.0	14.4
<i>Panel D. Cumulative real output-per-capita growth (percentage points)</i>								
U.S.–Thailand	105.4	6.7	-1.5	100.3	107.0	8.0	-1.8	100.8
Indonesia–Thailand	98.6	4.7	-1.6	95.5	99.0	5.2	-1.7	95.5

Notes: Durable-goods extension, 1976–2025. “Sorting” and “Demand” are Shapley contributions; “Neither” is the gap with ability aging and age-dependent demand removed. Components add to “Full”, up to rounding. Outcome definitions follow Table 6.

L.7 Schooling and family size

The counterfactuals assign to every path the mean and dispersion of years of schooling of Indonesian birth cohorts (Section 4.2). Section 5.5 adjusts the cohort means for differences in family size. Let $\bar{s}_c$ denote mean years of schooling of Indonesian birth cohort c , let f_c denote the total fertility rate of the path’s source country in the calendar year that corresponds to the birth year of cohort c , and let f_c^{IDN} denote the Indonesian rate in the cohort’s birth year. A unit difference in total fertility is treated as one additional sibling, so adjusted schooling on the path is

$$\bar{s}_c^{\text{adj}} = \bar{s}_c - 0.107 (f_c - f_c^{\text{IDN}}),$$

with the within-cohort dispersion of schooling unchanged. The coefficient is the effect of one additional co-resident sibling on the years of schooling completed by children aged 6–18, estimated by Feng (2021) from twin births in the 2010 Indonesian census. Total fertility rates are taken from the Gapminder series; for U.S.-path cohorts born before the series begins in 1800, its first value is used. On the Thai path the adjustment lowers the schooling of cohorts born before 1955, by at most 0.12 years, and raises that of later cohorts, by at most 0.34 years; on the U.S. path it lowers the schooling of all cohorts, by between 0.05 and 0.26 years. With cohorts weighted by their employment in 2025, the adjustment raises the schooling of the workforce by 0.22 years on the Thai path and lowers it by 0.19 years on the U.S. path.

The second exercise raises mean schooling on the Thai path by a fixed number of years for every cohort born from 1957 onward, the cohorts that enter the workforce at age 20 from 1977 onward. Earlier cohorts, the within-cohort dispersion of schooling, and the U.S. path are unchanged. In the

baseline without explicit durable accumulation, 3.2 additional years approximately equalize output per capita on the Thai and U.S. paths in 2025. With five additional years, the peak manufacturing employment share on the Thai path rises from 18.41 to 19.59 percent, and the gap to the U.S. path falls from 4.84 to 3.66 points. The durable-goods extension uses a separate completed experiment with 3.3 additional years, which puts Thai-path output per capita 1.3 percent above the U.S. path; five additional years raise the Thai manufacturing peak from 18.76 to 20.40 percent and reduce the gap from 4.76 to 3.12 points. For reference, mean schooling of Indonesian cohorts rises from 6.44 years for the cohort born in 1956 to 12.19 years for the cohort born in 2005. Each exercise is evaluated in general equilibrium within its separately calibrated specification, with price-taking agricultural and manufacturing exports. These schooling increases are imposed without modeling the resources needed to provide the additional education. Indeed, the policy intervention required for this would be very large and costly and that is precisely what the exercise is trying to illustrate. Table 58, Panel A, reports the family-size adjustment in both specifications.

Sectoral employment responses to schooling. Table 57 compares each schooling experiment with the full economy under the same demographic path and unchanged cohort schooling. The family-size adjustments are reported separately from the uniform increases in schooling. In 2025, the baseline model's 3.2-year increase reduces the Thai path's agricultural share by 3.92 points and raises high-skill services by 3.10 points; manufacturing rises by 0.79 points. With five additional years, the corresponding changes are -6.03 , 4.79 and 1.21 points. Low-skill services change little in either experiment. The durable-goods extension gives the same ranking of sectoral employment gains.

L.8 Welfare measures

This appendix details the consumption-welfare and child-adjusted output comparisons in Section 5.6. Both use the full-model U.S. and Thai equilibrium paths without the schooling adjustments. Agricultural exit costs are transfers within the dynasty and do not enter.

Expenditure-equivalent consumption gains. For the baseline without explicit durable accumulation, I calculate the additional expenditure needed at Thai prices to attain the U.S. path's utility from consumption at every age. Let $\mathcal{C}_t^{\text{USA}}(a)$ and $\mathcal{C}_t^{\text{THA}}(a)$ denote the consumption indices on the two paths. Preferences are common across paths, so attaining the same index at a given age delivers the same utility. From (36), the minimum expenditure needed to attain index $\mathcal{C}$ at age a , given absorption prices $\mathbf{p} = (p_j)_{j \in \mathcal{J}}$, is

$$e_a(\mathbf{p}, \mathcal{C}) = \left[\sum_{j \in \mathcal{J}} \Omega(a, j) \mathcal{C}^{(1-\sigma)\iota_j} p_j^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \quad (126)$$

Table 57: Sectoral employment responses to the schooling experiments

Experiment	A	M	L	H
<i>Baseline model</i>				
<i>2025</i>				
Family-size adjustment: Thailand	-0.35	0.06	0.07	0.22
Family-size adjustment: U.S.	0.16	-0.04	0.09	-0.21
Thailand: +3.2 years	-3.92	0.79	0.03	3.10
Thailand: +5 years	-6.03	1.21	0.03	4.79
<i>Mean, 1976–2025</i>				
Family-size adjustment: Thailand	-0.05	0.03	-0.03	0.05
Family-size adjustment: U.S.	0.21	-0.05	-0.04	-0.12
Thailand: +3.2 years	-1.46	0.42	-0.13	1.17
Thailand: +5 years	-1.96	0.59	-0.32	1.70
<i>Durable-goods extension</i>				
<i>2025</i>				
Family-size adjustment: Thailand	-0.35	0.08	0.07	0.20
Family-size adjustment: U.S.	0.16	-0.05	0.09	-0.19
Thailand: +3.3 years	-4.01	1.11	-0.05	2.96
Thailand: +5 years	-5.95	1.66	-0.16	4.44
<i>Mean, 1976–2025</i>				
Family-size adjustment: Thailand	-0.06	0.04	-0.03	0.05
Family-size adjustment: U.S.	0.21	-0.06	-0.05	-0.11
Thailand: +3.3 years	-1.75	0.63	0.04	1.07
Thailand: +5 years	-2.28	0.85	-0.05	1.48

Notes: Percentage-point changes from the same demographic path with unchanged schooling, under baseline export assumptions. Uniform increases apply to cohorts born from 1957 onward; the baseline model uses 3.19 years, reported as 3.2. Family-size adjustments follow the formula above. Sector definitions and annual averaging follow Table 47.

Using Thai absorption prices $\mathbf{p}_t^{Q, \text{THA}}$ and the Thai age distribution for both paths, I define

$$\Gamma_t = \frac{\sum_{a \in \mathcal{A}} \pi_t^{\text{THA}}(a) e_a(\mathbf{p}_t^{Q, \text{THA}}, \mathcal{C}_t^{\text{USA}}(a))}{\sum_{a \in \mathcal{A}} \pi_t^{\text{THA}}(a) e_a(\mathbf{p}_t^{Q, \text{THA}}, \mathcal{C}_t^{\text{THA}}(a))} - 1. \quad (127)$$

The numerator is the minimum total spending needed to attain the U.S.-path utility at each age, allowing the composition of spending to adjust to Thai prices. The denominator is actual expenditure under the Thai path. Thus, $100\Gamma_t$ is the percentage increase in total expenditure required, not a common percentage increase assigned to every age.

For 1976–2025, I average the fifty annual expenditure gaps, giving each year equal weight:

$$\bar{\Gamma} = \frac{1}{50} \sum_{t=1976}^{2025} \Gamma_t. \quad (128)$$

The reported period average is $100\bar{\Gamma}$ percent. The resulting expenditure equivalent is 44.4 percent in 2025 and averages 22.1 percent over 1976–2025. Age-specific expenditure increases range from 43.7 to 44.9 percent in 2025 and are positive at every age and in every year. The calculation is unchanged by a consistent renormalization of the consumption index: replacing ι_j with $k\iota_j$ and $\mathcal{C}$ with $\mathcal{C}^{1/k}$, for any $k > 0$, leaves every term in (126) unchanged.

Durable-goods extension. I apply the same expenditure comparison to the model with durable goods: what spending is needed, at Thai prices, to attain the U.S. path’s utility at every age and

date? Durable purchases provide services beyond the year of purchase, so non-durable consumption and durable purchases must be chosen jointly over time.

Let $e_a^D(\mathbf{p}, \mathcal{C}, D)$ denote the non-durable expenditure required for a given consumption index and durable stock, as defined in (107). With prices and population weights taken from the Thai equilibrium, annual spending needed to deliver the U.S. utility targets is

$$B_t = \sum_{a \in \mathcal{A}} \pi_t^{\text{THA}}(a) \left[e_a^D(\mathbf{p}_t^{Q, \text{THA}}, \mathcal{C}_t^{\text{USA}}(a), D_t(a)) + p_{M,t}^{Q, \text{THA}} x_t(a) \right]. \quad (129)$$

The consumption targets are those of the U.S. equilibrium in the durable-goods specification.

I choose stocks and purchases to minimize $\sum_{t=1976}^{2094} q_t B_t$, where $q_t = \beta^{t-1976} \lambda_t^{\text{THA}} / \lambda_{1976}^{\text{THA}}$ values expenditure at different dates using the Thai dynasty's marginal value of resources. The calculation respects the accumulation and no-resale constraints in (103)–(104). Initial holdings are the Thai stocks inherited from 1975, and entering cohorts have no inherited durables. I retain the model's continuation through 2094, allowing purchases made during the reporting period to provide services over cohorts' remaining lifetimes.

Applying this calculation to the Thai utility targets reproduces the original Thai stocks, purchases, and expenditure. Let $B_t^{\text{USA|THA}}$ denote spending along the jointly cost-minimizing path for the U.S. utility targets, and let B_t^{THA} denote actual Thai expenditure. I report

$$\Gamma_t^D = \frac{B_t^{\text{USA|THA}}}{B_t^{\text{THA}}} - 1, \quad \bar{\Gamma}^D = \frac{1}{50} \sum_{t=1976}^{2025} \Gamma_t^D. \quad (130)$$

The weights q_t determine the timing of purchases in the cost-minimization problem; the reported average gives each year equal weight.

Spending along this path is 43.7 percent higher than Thai expenditure in 2025 and 21.4 percent higher on average over 1976–2025. These results are close to the baseline expenditure equivalents of 44.4 and 22.1 percent, respectively. As in the baseline, the calculation is invariant to a consistent renormalization of the consumption index: multiplying all income-effect parameters, including ι_D , by $k > 0$ and replacing $\mathcal{C}$ with $\mathcal{C}^{1/k}$ leaves the expenditure requirements unchanged.

Output and children's consumption needs. The output comparisons are calculated separately for each specification. For demographic path s , let n_t^s be the number of persons younger than 20 per person aged 20–89, taken from the same population sources as the paths. Child-adjusted output per person, also termed output per equivalent person, is

$$y_t^{\text{eq},s} = \frac{y_t^{PC,s}}{1 + 0.3 n_t^s}, \quad (131)$$

where $y_t^{PC,s}$ is real output per person aged 20–89, valuing production at that path's 1976 producer prices. This changes only the population denominator; model output and consumption allocations

Table 58: Schooling and welfare

				Baseline model without explicit durable accumulation		Extended model with durable goods	
<i>Panel A. U.S.–Thailand gap, without and with the schooling adjustment</i>				Unadjusted	Adjusted	Unadjusted	Adjusted
Peak manufacturing employment share (percentage points)				4.84	4.75	4.76	4.62
Mean manufacturing employment share (percentage points)				2.86	2.79	2.82	2.73
Output per capita in 2025 (percent)				42.8	36.7	44.1	37.7
Cumulative growth of output per capita (percentage points)				99.6	88.2	105.4	93.2
<i>Panel B. U.S. path relative to Thailand's path (percent)</i>				2025	1976–2025	2025	1976–2025
Expenditure equivalent, common age weights				44.4	22.1	43.7	21.4
Output per capita (ages 20–89)				42.8	22.8	44.1	23.8
Output per equivalent person, children weighted 0.3				21.9	5.4	23.0	6.3

Notes: Price-taking exports in agriculture and manufacturing. Panel A adjusts cohort schooling by -0.107 years per unit difference in birth-year total fertility relative to Indonesia (Appendix L.7). Panel B uses Thai prices and age weights for the expenditure equivalent. The output rows value production at each path's own 1976 producer prices; valuing all paths at the Indonesian baseline's 1976 prices changes them by at most 0.6 percentage points (to 42.8, 22.8, 21.9 and 5.5 in the baseline model and 43.5, 23.3, 22.5 and 5.9 in the extension). The 1976–2025 columns report arithmetic averages of annual percentage differences (Appendix L.8).

remain unchanged. In 2025 there are 0.80 persons younger than 20 per person aged 20–89 on the U.S. path, 0.49 in Indonesia, and 0.20 on the Thai path.

For either output measure y_t , the annual U.S.–Thailand comparison is $100(y_t^{\text{USA}}/y_t^{\text{THA}} - 1)$. The 1976–2025 column in Table 58, Panel B, reports the arithmetic average of these fifty annual percentage differences. In the baseline, child-adjusted output per person is 21.9 percent higher under the U.S. path in 2025 and 5.4 percent higher on average over 1976–2025.