Population Aging and the Realignment of World Trade

Joseph Kopecky*

10th September 2026

Abstract

The world's workforces are aging, some faster than others. I quantify how this divergence realigns the pattern of world trade, in a multi-country, multi-sector Ricardian model in which workers of different ages supply different skills. Aging moves trade in two ways. Diverging workforce sizes move its volume: countries with more workers produce and export more. The changing skill mix moves its composition: as a workforce ages, the sectors a country is good at shift, and so does what it exports. By 2050 demographics alone lower China's exports to the United States by 9% while raising India's by 33%. For the world's largest trading partners, a decades-long demographic tailwind that promoted trade integration is now reversing, even as it keeps building for younger exporters. This paper estimates the scale of that reversal and traces the compositional shift underneath it, as trade moves out of the sectors that lean on the young and into those that use older workers. Aging also moves the external balance, turning the advanced and East Asian economies into net creditors and the young into debtors, broadly matching observed net foreign asset positions, and leaving the corridor results intact. The welfare effects of these flows are largely offsetting. The goods a fast-aging country buys grow more expensive as the services it makes grow cheaper. Much larger welfare effects arise from demographic change, but fall on domestic margins.

JEL classification codes: F11, J11, F17 *Keywords:* Population aging, Comparative advantage,

Global imbalances, Trade, Demographics

*Trinity College Dublin (jkopecky@tcd.ie).

1. Introduction

The demographic transition is a slow, predictable force reshaping the global economy. The population aged 65 and older will more than double between 2024 and 2050 (United Nations, 2024), and working-age numbers are now falling across most of the rich world and in China even as they keep rising across much of South Asia, Latin America, and Sub-Saharan Africa. What this does to labour markets, public finances, and growth is well studied (Bloom, Canning, and Fink, 2010; Maestas, Mullen, and Powell, 2023). What it does to the *pattern* of world trade is not. The link between trade and aging has been studied as a reduced-form regularity, but its expected effect on the scale, composition, and balance of world trade has never been carefully traced. This paper pins it down, with a structural model.

Aging reshapes trade through the workforce, in two ways. The first is a matter of size. As a population ages its working-age numbers eventually fall, and a country with fewer workers produces and exports less. In countries where the young keep arriving, this "demographic dividend" points the other way, toward rising trade for years to come. This size effect is the primary source of the demographic headwinds and tailwinds pushing the volume of trade between any pair of countries. The second force is age-specific skill. Workers of different ages are good at different things: verbal fluency and comprehension hold up or improve into later working life, while processing speed, memory, and physical capacity decline early (Salthouse, 2009; Schaie, 2005; de Zwart, Frings-Dresen, and Kilbom, 2001). Because sectors draw on these abilities in different proportions, an aging workforce shifts a country's comparative advantage across sectors, away from those that lean on the young and toward those that use the skills of the old. I call this life-cycle comparative advantage. It has a smaller net effect on the size of flows, but large distributional effects on which sectors move within them. Beyond the workforce, aging also moves the external balance: older workers hold more wealth, so a country aging faster than the world wants to save more than it invests at home and lends the difference abroad. Aging thus reshapes trade directly, through the size and skill of the workforce, and moves the external balance indirectly, through saving.

Two facts motivate this paper. The first is a correlation in the trade data: Brakman, Kohl, and van Marrewijk (2025) and Kopecky (2023) put demographic structure into gravity equations and find that it has real predictive power for trade, both in past data and projected out across the century. A coefficient of this kind says that demographics matter for trade, not why. The second is the mechanism behind it: Cai and Stoyanov (2016) and Gu and Stoyanov (2019) show that the skills a workforce supplies shift as it ages, the life-cycle

comparative advantage I build into the model. A structural model joins the two. Holding income, capital, and institutions fixed, it isolates the demographic component of trade: it says not just that demographics move trade but why, and prices that component corridor by corridor and sector by sector. I embed an age-skill mechanism into the multi-sector Eaton–Kortum model of Caliendo and Parro (2015), replacing homogeneous labour with three age-dependent skills whose supplies are set by a country's age structure; the trade side then turns shifts in those supplies into bilateral, sectoral flows, with counterfactuals following from exact hat algebra.

The central results are about the pattern of trade, and demographics have been pushing on it for decades. Run the model back through the age structures of the past sixty years, holding all non-demographic structure fixed, and the demographic component of China's exports to the United States climbs with its working-age boom and peaks in the decade after 2014. What had been a tailwind behind the world's largest corridors for half a century becomes a headwind: by 2050, demographics alone lower China's exports to the United States by 9% while raising India's by 33%, the same reorientation running across the young exporters generally. Within each corridor the skill channel decides which sectors move: losses concentrate where production leans on the skills of the young, sparing electronics and falling hardest on textiles and metals. Demographic change is a structural force, and it adds to the realignment of world trade already under way.

In the baseline model, each country's trade balance is held fixed, so the corridors above isolate the effects of size and composition. The paper then extends the model to let the balance respond to aging. Older workers hold more wealth, so as a country ages its desired saving rises; letting national balances adjust turns the advanced and East Asian economies into net creditors and the young into debtors, broadly matching observed net foreign asset positions. The trade corridors themselves are all but unchanged once imbalances are allowed.

The reorientation's own welfare is small and largely offsetting: as the world ages, the goods a fast-aging country buys grow dearer while the services it makes grow cheaper, a cost and a benefit of similar size. Aging moves welfare far more than this, but through the size of the workforce, not through trade: output per person tracks the working-age share, which aging pushes down in the old countries and up in the young. This is the same force a closed economy would feel, standard labour-supply accounting (Cooley and Henriksen, 2018; Cooley, Henriksen, and Nusbaum, 2024; Kitao, 2015), and I keep it apart from the gains and losses that flow through trade.

The paper bridges the quantitative trade literature (Eaton and Kortum, 2002; Caliendo and Parro, 2015) and the literature on the macroeconomics of aging (Auclert, Malmberg,

Martenet, and Rognlie, 2025; Cravino, Levchenko, and Rojas, 2022; Maestas et al., 2023; Cooley et al., 2024). It provides the structural counterpart to Cai and Stoyanov (2016) and builds on Sposi (2022), whose demographics-and-imbalances saving channel I bring into the multi-sector trade model, so that a single demographic shock drives both the sectoral composition and the external balance. It complements Brakman et al. (2025), who project aggregate demographic trade effects through gravity without the within-country compositional channel. The China–US predictions connect to the literature on trade with China (Autor, Dorn, and Hanson, 2013), adding a demographic channel alongside the geopolitical one; the skill-endowment mechanism relates to factor-endowment views of trade composition (Burstein and Vogel, 2017) and to Acemoglu and Restrepo (2022) on aging and automation, a margin I hold fixed.

Section 2 presents the model; Section 3 the calibration; Sections 4 and 5 the two realignments, of trade and of the external balance; Section 6 the aggregate effects on welfare, openness, and policy; Section 7 gathers the robustness checks; and Section 8 concludes.

2. Model

I embed age-dependent skill supply into the multi-sector Eaton–Kortum framework of Caliendo and Parro (2015). Demography enters the model in one place: a country's age structure determines its effective supply of three age-dependent task inputs, through life-cycle profiles that differ across inputs. If those profiles were identical, aging would change the size of the workforce but not the mix of inputs, and there would be no compositional trade effect at all. The rest of the model is the standard machinery that turns relative input supplies into bilateral, sectoral trade.

The economy consists of N countries indexed by $c \in \{1, \ldots, N\}$ and J tradeable manufacturing sectors indexed by $j \in \{1, \ldots, J\}$, together with a single non-tradeable sector indexed by $j = 0$. Time is discrete and indexed by t. I model a sequence of static trade equilibria linked by exogenous demographic transitions; the baseline imposes balanced trade, an assumption Section 2.5 states and defends and Section 5 relaxes.

2.1. Demographics and Skill Endowments

Each country c at time t has a population distributed across $\mathcal{A}$ age groups, with $n_{c,a,t}$ the mass of individuals in age group a.[1] Workers supply three age-dependent task inputs, indexed by $k \in \{A, D, P\}$:

[1] In the calibration I use five-year age groups from 20–24 through 60–64, giving $\mathcal{A} = 9$ working-age groups. Dependants (ages 0–19 and 65+) do not supply labour.

- **Age-appreciating cognitive** ($k = A$): oral and written communication and comprehension, maintained or improving through the mid-50s and beyond;
- **Age-depreciating cognitive** ($k = D$): memory, divided attention, perceptual and processing speed, declining from the late 20s;
- **Physical** ($k = P$): strength, stamina, coordination, and dexterity, declining from the mid-20s.

Following the neuroscience and gerontology literatures surveyed for the purposes of international trade in Cai and Stoyanov (2016), I aggregate the underlying abilities into these three composites; Section 3.1 describes the construction. The two declining inputs, D and P, follow closely related life-cycle profiles, and for that reason the baseline calibration imposes a common profile on them.

The efficiency of an age-a worker in input k is

$$e^k(a) = \bar{e}^k \cdot g^k(a), \tag{1}$$

where $\bar{e}^k > 0$ is a normalising constant and $g^k : \{1, \ldots, \mathcal{A}\} \to \mathbb{R}_+$ is the age profile for input k, satisfying

$$g^A(a') \geq g^A(a) \quad \text{for } a' > a, \tag{2}$$

$$g^m(a') < g^m(a) \quad \text{for } a' > a \geq \underline{a}, \quad m \in \{D, P\}, \tag{3}$$

the appreciating profile non-decreasing (in the calibration, rising to the mid-50s and then flat) and the two declining profiles falling beyond their early-career peak $\underline{a}$. These shapes are calibrated from the cognitive-aging literature in Section 3.3.

Workers supply their task bundles inelastically, so the aggregate effective supply of input k in country c is

$$L^k_{c,t} = \sum_{a=1}^{\mathcal{A}} n_{c,a,t} \cdot e^k(a). \tag{4}$$

Demographic change shifts the age distribution $\{n_{c,a,t}\}$ and therefore the vector of effective endowments $\mathbf{L}_{c,t} = (L^A_{c,t}, L^D_{c,t}, L^P_{c,t})$. Two countries with identical total populations but different age structures have different effective endowments, and hence different comparative advantage. This is the variation that drives all of the model's trade results. All of the demographic action enters the rest of the model through $\mathbf{L}_{c,t}$, and its compositional part (the tilt in $\mathbf{L}$ at fixed scale) exists only because the profiles in (2)–(3) differ across inputs.

2.2. Production

2.2.1 Tradeable sectors

Within each tradeable sector j there is a continuum of varieties $\omega \in [0,1]$. The unit cost of producing any variety in sector j in country c is

$$mc_{c,t}^{j} = \underbrace{\left[\frac{1}{Z_{c,t}^{j}} \prod_{k \in \{A,D,P\}} \left(w_{c,t}^{k}\right)^{\alpha_{j}^{k}}\right]^{\beta_{j}}}_{\text{value added}} \cdot \underbrace{\prod_{\ell=0}^{J} \left(P_{c,t}^{\ell}\right)^{\gamma_{\ell}^{j}}}_{\text{intermediates}}, \tag{5}$$

where $w_{c,t}^{k}$ is the price of input k, $\alpha_{j}^{k} \geq 0$ is its cost share in sector j's value added (with $\sum_{k} \alpha_{j}^{k} = 1$), $P_{c,t}^{\ell}$ is the price index of sector ℓ, $\gamma_{\ell}^{j} \geq 0$ is the share of sector ℓ intermediates in sector j, and $Z_{c,t}^{j}$ is a value-added productivity term. Value added and intermediates combine Cobb–Douglas with value-added share $\beta_{j} = 1 - \sum_{\ell} \gamma_{\ell}^{j}$.[2]

The cost shares $\{\alpha_{j}^{k}\}$ link demographics to comparative advantage on the cost side; Section 3.1 describes their construction from O*NET ability scores and Bureau of Labor Statistics occupational data, following Cai and Stoyanov (2016). Equation (5) contains the first of the two channels through which demographics reach trade: a country abundant in input k has a low price $w_{c,t}^{k}$, which lowers unit costs most in sectors that use input k intensively. This is the Heckscher–Ohlin factor-price channel, and it operates whether or not demographics also shift sectoral productivity.

2.2.2 The non-tradeable sector

The non-tradeable sector $j = 0$ uses the same Cobb–Douglas technology over the three inputs and intermediates, with its own cost shares $\{\alpha_{0}^{k}\}$ and coefficients $\{\gamma_{\ell}^{0}\}$, calibrated from the data exactly as the tradeable sectors are (Section 3.2, Appendix C). It enters household consumption and serves as an intermediate input, but it is not traded internationally: $\pi_{cp,t}^{0} = 1$ if $c = p$ and zero otherwise.

[2] A note on notation relative to Caliendo and Parro (2015): there α denotes final consumption shares, which are η_{c}^{j} here (Section 2.4), and the Fréchet location below is written T following Eaton and Kortum (2002) rather than their λ. Bilateral subscripts here run source-first, so π_{cp}^{j} is exporter c's share of importer p's expenditure, the reverse of their importer-first convention.

2.2.3 Ricardian productivity

Country c's productivity for variety ω in sector j is drawn from a Fréchet distribution,

$$\Pr\left[z_c^j(\omega) \leq z\right] = \exp\left(-T_{c,t}^j z^{-\theta_j}\right), \tag{6}$$

with location $T_{c,t}^j > 0$ (absolute advantage) and dispersion $\theta_j > 1$ (scope for comparative advantage). This Fréchet structure (Eaton and Kortum, 2002) delivers gravity, multilateral resistance, and tractable counterfactuals. I allow the location to depend on demographics,

$$\ln T_{c,t}^j = \mu_c + \mu_j + \sum_{k \in \{A,D,P\}} \rho\, \alpha_j^k \ln L_{c,t}^k + \nu_{c,t}^j, \tag{7}$$

where μ_c, μ_j are country and sector effects, $\nu_{c,t}^j$ a residual, and $\rho \geq 0$ governs the strength of the second channel: a larger effective endowment of input k raises productivity itself in sectors that use input k intensively. The economic content is an external economy of specialization. Where an input is abundant, the sectors built around it accumulate the techniques, organizational knowledge, and specialized suppliers that show up as higher measured productivity, and equation (7) is the structural counterpart of exactly this pattern in the Cai and Stoyanov (2016) reduced form. The channel works in the same direction as the factor-price channel in (5), and the direction of the mechanism does not depend on how the two divide the work quantitatively. One feature of writing $T_{c,t}^j$ in the levels of the endowments is that a proportional rise in all three inputs raises $T_{c,t}^j$ by ρ uniformly across sectors, so demographic *scale* enters productivity as a Hicks-neutral shifter alongside the labour-supply channel. This mirrors the Cai and Stoyanov (2016) reduced form the specification is built from; it leaves the paper's object, the *composition* response, untouched, because that response is a factor-price effect that survives at $\rho = 0$.

2.3. Trade

Bilateral trade follows the Eaton–Kortum/Caliendo–Parro gravity structure. Country c's share of country p's sector-j expenditure is

$$\pi_{cp,t}^j = \frac{T_{c,t}^j \left(mc_{c,t}^j\, d_{cp}^j\right)^{-\theta_j}}{\Omega_{p,t}^j}, \qquad \Omega_{p,t}^j = \sum_{c'=1}^{N} T_{c',t}^j \left(mc_{c',t}^j\, d_{c'p}^j\right)^{-\theta_j}, \tag{8}$$

where $d_{cp}^j \geq 1$ is the iceberg trade cost ($d_{cc}^j = 1$) and $\Omega_{p,t}^j$ is the multilateral-resistance term. The sectoral price index is $P_{p,t}^j = \zeta_j \left(\Omega_{p,t}^j\right)^{-1/\theta_j}$ with ζ_j a constant. For the non-tradeable

sector $d^0_{cp} \to \infty$ for $c \neq p$, so $\pi^0_{cp,t}$ is the identity and $P^0_{c,t}$ depends only on c's own costs. Bilateral sector-j exports are $X^j_{cp,t} = \pi^j_{cp,t} E^j_{p,t}$, with $E^j_{p,t}$ total expenditure of p on sector j.

2.4. Demand

A representative household in country c has Cobb–Douglas preferences over the $J+1$ sectors with expenditure shares $\{\eta^j_c\}$, $\sum_{j=0}^{J} \eta^j_c = 1$, so that final demand for sector j is $F^j_{c,t} = \eta^j_c Y_{c,t}$ with $Y_{c,t}$ total income, and total expenditure adds intermediate demand,

$$E^j_{c,t} = \eta^j_c \, Y_{c,t} + \sum_{\ell=0}^{J} \gamma^\ell_j \, Y^\ell_{c,t}, \tag{9}$$

where $Y^\ell_{c,t}$ is gross output of sector ℓ, so expenditure on sector j combines final demand $\eta^j_c Y_{c,t}$ with intermediate demand from every using sector. Demographics affect trade only through the supply side: expenditure shares are held fixed, so I abstract from any age-driven shift of consumption toward non-tradeables.[3]

2.5. Equilibrium and Balanced Trade

Given demographics, trade costs, and residual productivities, a competitive equilibrium is a set of input prices $\{w^k_{c,t}\}$, price indices $\{P^j_{c,t}\}$, trade shares $\{\pi^j_{cp,t}\}$, and expenditures $\{E^j_{c,t}\}$ such that goods markets clear ($\sum_p X^j_{cp,t} = Y^j_{c,t}$), input markets clear ($L^k_{c,t} = \sum_{j=0}^{J} \alpha^k_j \beta_j Y^j_{c,t} / w^k_{c,t}$), income equals factor payments ($Y_{c,t} = \sum_k w^k_{c,t} L^k_{c,t}$), and trade balances.

In the baseline, expenditure equals income in every country, with no financial account, so each country's total exports identically equal its total imports:

$$\sum_{j=1}^{J} \sum_{p} X^j_{cp,t} = \sum_{j=1}^{J} \sum_{p} X^j_{pc,t} \quad \forall c. \tag{10}$$

Balanced trade is a substantive assumption, and the reason for it is that observed external balances are not neutral objects to carry into a demographic counterfactual. A long literature traces current-account positions in part to the same life-cycle saving behaviour this paper studies (Backus, Cooley, and Henriksen, 2014; Krueger and Ludwig, 2007; Auclert et al., 2025; Sposi, 2022); holding the 2014 balances fixed through 2050 would import into the baseline, unmodelled, a piece of the very mechanism whose trade consequences I want

[3] Allowing η^j_c to depend on the age distribution would add a demand-side channel; calibrating it credibly requires age-specific sectoral expenditure data across many countries, which is not currently available.

to isolate. Starting from balance keeps the composition channel clean. The cost is silence on the aging–saving–current-account corollary of the same shock, and the paper does not stay silent for long: Section 5 adds a saving block in which the age distribution drives desired wealth, and with it who lends to whom. Little rides on the choice for the flow results: the standard alternative in this framework carries observed aggregate imbalances as exogenous transfers (Dekle, Eaton, and Kortum, 2008; Caliendo and Parro, 2015), and Section 7 solves the model that way, with the realized 2014 imbalances carried forward, and finds the corridor reorientation essentially unchanged.

2.6. Counterfactual Analysis: Exact Hat Algebra

Counterfactuals use the exact hat algebra of Dekle et al. (2008). Writing $\hat{x} = x'/x$, a change in demographics from $\{n_{c,a,t}\}$ to $\{n'_{c,a,t'}\}$ implies endowment changes

$$\hat{L}_c^k = \frac{\sum_a n'_{c,a,t'}\, e^k(a)}{\sum_a n_{c,a,t}\, e^k(a)}, \tag{11}$$

and the counterfactual equilibrium solves, in hat form,

$$\widehat{mc}_c^j = \prod_k \left(\hat{w}_c^k\right)^{\alpha_j^k \beta_j} \prod_\ell \left(\hat{P}_c^\ell\right)^{\gamma_\ell^j}, \tag{12}$$

$$\hat{\pi}_{cp}^j = \frac{\hat{T}_c^j \left(\widehat{mc}_c^j\right)^{-\theta_j}}{\sum_{c'} \pi_{c'p}^j\, \hat{T}_{c'}^j \left(\widehat{mc}_{c'}^j\right)^{-\theta_j}}, \qquad \ln \hat{T}_c^j = \sum_k \rho\, \alpha_j^k \ln \hat{L}_c^k, \tag{13}$$

together with goods- and input-market clearing, using only observed trade shares $\{\pi_{cp,t}^j\}$, the calibrated parameters, and the endowment changes $\{\hat{L}_c^k\}$; the level objects $T_{c,t}^j$ and d_{cp}^j are not needed. The value-added productivity term is held fixed ($\hat{Z}_c^j = 1$), so demographics act on technology only through $\hat{T}_c^j$. Equations (11)–(13) also make precise where demography acts: the profiles determine $\hat{L}_c^k$, and $\hat{L}_c^k$ reaches comparative advantage through factor prices $\hat{w}_c^k$ and, scaled by ρ, through $\hat{T}_c^j$.

The model computes a sequence of static equilibria rather than a dynamic transition path. Each is the long-run allocation consistent with its period's age structure; a dynamic treatment would add the gradual path toward it, set by the frictions that slow workers' movement across sectors. Caliendo, Dvorkin, and Parro (2019) quantify that margin for a shock of comparable size: after the China trade shock, manufacturing employment reaches

its long-run reallocation within about fifteen years.[4] My horizons run to thirty-five years and beyond, and aging spreads its reallocation across decades rather than compressing it into a few, so the demands it places on mobility in any period are smaller and the economy stays near each period's long-run allocation. What these frictions mainly reprice is the welfare cost of transition, the margin I already treat as a small residual rather than as the pattern of trade I study. Two qualifications remain, and neither disturbs the headline. Because aging is continuous, the most recent demographic increments have not fully worked through by the terminal date, so the static snapshot runs a little ahead of a frictional economy. And the corridor volumes are carried by workforce size, not by the reallocation of workers across sectors, so they are largely insulated from mobility frictions to begin with; the dynamic margin can bear only on the smaller compositional channel.[5] The remaining dynamic margin, capital accumulation, I set aside on its own terms: aging raises saving and hence the capital stock, which independently shifts comparative advantage. I hold capital fixed to isolate the demographic skill channel, so the magnitudes here are the undamped skill-channel benchmark, and the conclusion returns to what adding capital would change.

3. Calibration

The model is calibrated to the base year 2014, the last year of the WIOD 2016 release (Timmer, Dietzenbacher, Los, Stehrer, and de Vries, 2015), for 30 countries and 21 sectors (20 tradeable manufacturing sectors and the non-tradeable sector). The empirical discipline is concentrated in two places: the trade elasticities, taken from Caliendo and Parro's tariff-based estimates, and the base-year input–output and trade structure, which the calibration matches exactly. Appendix D lists the full data sources.

3.1. Skill Intensities

The skill-intensity matrix $\{\alpha_j^k\}$ maps each sector to its cost shares in the three task inputs. Following Cai and Stoyanov (2016), I construct it from O*NET version 29.1 ability importance ratings across occupations and Bureau of Labor Statistics Occupational Employment and Wage Statistics (May 2014) employment counts by occupation within each manufac-

[4] The China shock lowers the U.S. manufacturing employment share by 0.36 percentage points both at fifteen years and in the long run (Caliendo et al., 2019, Fig. 1).

[5] A faithful dynamic treatment in the manner of Caliendo et al. (2019) would require time series of gross cross-sector labour flows for every country in the sample, which exist at usable quality for only a handful of economies, so a full port to a thirty-country calibration is not currently feasible. If one were nonetheless concerned that mobility frictions dampen the compositional results, a feasible check would calibrate the adjustment speed to their estimates, apply it to the demographic path, and measure how far the frictional terminal allocation departs from the static one.

turing industry (National Center for O*NET Development, 2024; U.S. Bureau of Labor Statistics, 2014). The underlying abilities are reduced to three composites by principal components, aggregated to sectors with employment weights, and normalised so that $\alpha_j^A + \alpha_j^D + \alpha_j^P = 1$.[6] These should be read as task-intensity indices rather than structural cost shares (the same caveat that applies to Cai and Stoyanov (2016)'s construction) and they are assumed common across countries. These indices bear on the results unequally. The sectoral pattern is set by the cross-sector *ranking* of intensities, which the O*NET data discipline more tightly than their cardinal spread; the composition-welfare magnitude scales with the cardinal values themselves. The corridor volumes, which move with the size of the workforce, depend on neither. To the extent that developing-country manufacturing relies more on physical tasks than its US counterpart, the model understates cross-country intensity differences, consistent with the larger reduced-form coefficients on the OECD subsample (Appendix B).

The matrix varies substantially across sectors. Computers and electronics is the most appreciating-intensive manufacturing sector ($\alpha^A = 0.586$), reflecting its engineers, technical writers, and managers, while wood products and textiles are the most depreciating- and physical-intensive; the manufacturing mean appreciating intensity is 0.229. The non-tradeable services sector is strongly appreciating-intensive: its cost share $\alpha_0^A = 0.404$ is 1.8 times the manufacturing mean and the second highest of all twenty-one sectors, with correspondingly low physical intensity.

3.2. Trade Parameters

Bilateral trade shares $\{\pi_{cp}^j\}$, input–output coefficients $\{\gamma_\ell^j\}$, value-added shares $\{\beta_j\}$, and final consumption shares $\{\eta_c^j\}$ are extracted from the WIOD 2016 release for 2014. I aggregate the 43 WIOD countries to 30 economies: 29 named countries plus a rest-of-world composite that absorbs the remaining 14 countries and WIOD's own rest-of-world region and the 56 WIOD industries to 20 tradeable manufacturing sectors plus one non-tradeable sector (WIOD industries 24–56). The non-tradeable sector's domestic trade share is set to unity and the consumption shares are renormalised over the 21 sectors; the resulting non-tradeable consumption weight η_c^0 averages 0.79 (range 0.65 to 0.88). Appendix C documents the services build and the base-year accounting identity it satisfies.

The trade elasticities $\{\theta_j\}$ are taken from Caliendo and Parro (2015), who estimate them from tariff variation; following them I impose a lower bound for sectors whose point

[6] I use the "Importance" ratings (Scale ID = IM) and normalise each ability by its loading on Inductive Reasoning before extracting components, following Cai and Stoyanov's (2016) methodology, so that the first three components map interpretably onto the appreciating, depreciating, and physical composites.

estimates are implausibly low.

3.3. Age–Productivity Profiles

The profiles $\{g^k(a)\}$ govern how each input evolves over the working life, calibrated from the cognitive-aging and occupational-health literatures. For the appreciating input, Schaie (2005) documents verbal ability and comprehension rising or flat through age 60; I fit a log-linear profile rising about 0.5% per year to the mid-50s and then flat. For the declining inputs, Salthouse (2009) and de Zwart et al. (2001) document processing speed, working memory, and physical capacities peaking in the late 20s and declining roughly 0.5 to 1% per year thereafter; the baseline imposes a common declining profile on the depreciating-cognitive and physical inputs, and Appendix A differentiates them and shows the results are invariant. The profiles are common across countries. Lagakos, Moll, Porzio, Qian, and Schoellman (2018) find experience–wage profiles steeper in rich countries, so uniform profiles may understate life-cycle comparative advantage for advanced economies and overstate it for developing ones.

3.4. Demographics

Demographic data come from the United Nations World Population Prospects 2024 Revision (United Nations, 2024), by five-year age group, spanning 1950–2100: observed estimates through 2023 and medium-variant projections thereafter. For each country and year I apply the calibrated profiles to the age distribution to obtain effective supplies $L^k_{c,t}$.

3.5. The Productivity Elasticity

One parameter, ρ, is not set from outside data. It governs the strength of the productivity channel in (7): how much an abundant skill raises measured productivity in the sectors that use it, on top of the factor-price effect. I set $\rho = 0.5$. Choosing ρ to match the model's skill–export gradient to the one measured in the data gives a nearly flat fit across the range I search (0.1 to 3.0): the data do not pin it down, but they never reject 0.5. This changes none of the results. The composition response, which sectors a country is good at, runs through relative factor prices and is present even with the productivity channel switched off ($\rho = 0$); ρ scales only the productivity term layered on top. Section 7 confirms the counterfactuals barely move as ρ ranges across this interval.

4. The Reorientation of Trade

I now turn to the paper's main results. Demographic change has large quantitative effects on both how much countries trade and which goods they trade. I take the two in turn. The first is volume: both the size of a workforce and its age composition bear on the volume of its bilateral exports, but size dominates, and the age margin barely moves the totals. The second is composition: underneath the aggregate flow, the shifting age-skill mix drives a large shuffle in the composition of what a country exports.

4.1. The Volume of Trade

I begin with a single corridor, one of the world's largest: bilateral trade between China and the United States. Figure 1 sets the demographic driver against the trade the model predicts from it. The top panel traces China's working-age population share relative to the United States. From the late 1970s a working-age boom shifted labor abundance toward China, and the model predicts a steady rise in China–US trade through the same decades (bottom panel, dashed). Demographics alone would have raised the corridor about 70% from 1980 to 2014. Actual trade grew many times over, as globalization and China's own industrialization transformed the relationship, all of which the counterfactual holds fixed at 2014; only age structure moves. So demographics cannot account for the entirety of the boom, but it was a meaningful structural tailwind behind it. Past 2014 the tailwind crests and reverses: by 2050 it leaves the corridor 9% below its 2014 level.

Table 1 places the 2050 reorientation next to the demographic tailwind of the three decades before it. Demographics lifted all eight corridors between 1980 and 2014; after 2014 that lift divides. For the aging exporters it becomes a headwind, and China, Japan, Germany, and Korea all move from gains to losses. For the young it only weakens: India's and Mexico's exports keep rising, by 33% and 27% to 2050, but well below their earlier pace, when demographics alone added over 100% to each.

Table 1 also splits each corridor's change into two parts: a scale component, from the changing size of the workforce, and a composition component, from its changing skill mix.[7] Almost all of the volume is scale. Across the eight corridors composition carries about 2% of the move, never shifts a corridor by more than a percentage point, and for India's exports to the United States runs mildly the other way. This lines up with the gravity evidence

[7]Scale moves all three inputs by the geometric mean of their changes, $\hat{L}_c^{\text{scale}} = (\hat{L}_c^A \hat{L}_c^D \hat{L}_c^P)^{1/3}$, and composition is the residual $\hat{L}_c^{k,\text{comp}} = \hat{L}_c^k / \hat{L}_c^{\text{scale}}$, which shifts the mix at fixed scale. The full shock factors into the two, $\hat{L}_c^k = \hat{L}_c^{\text{scale}} \hat{L}_c^{k,\text{comp}}$, and since trade responds to it roughly log-linearly the channels combine multiplicatively, $1 + x_{\text{full}} \approx (1 + x_{\text{scale}})(1 + x_{\text{comp}})$. In percent their contributions then sum only up to the cross-term $x_{\text{scale}}\, x_{\text{comp}}$, negligible over 2014–2050 and a few points over the longer window.

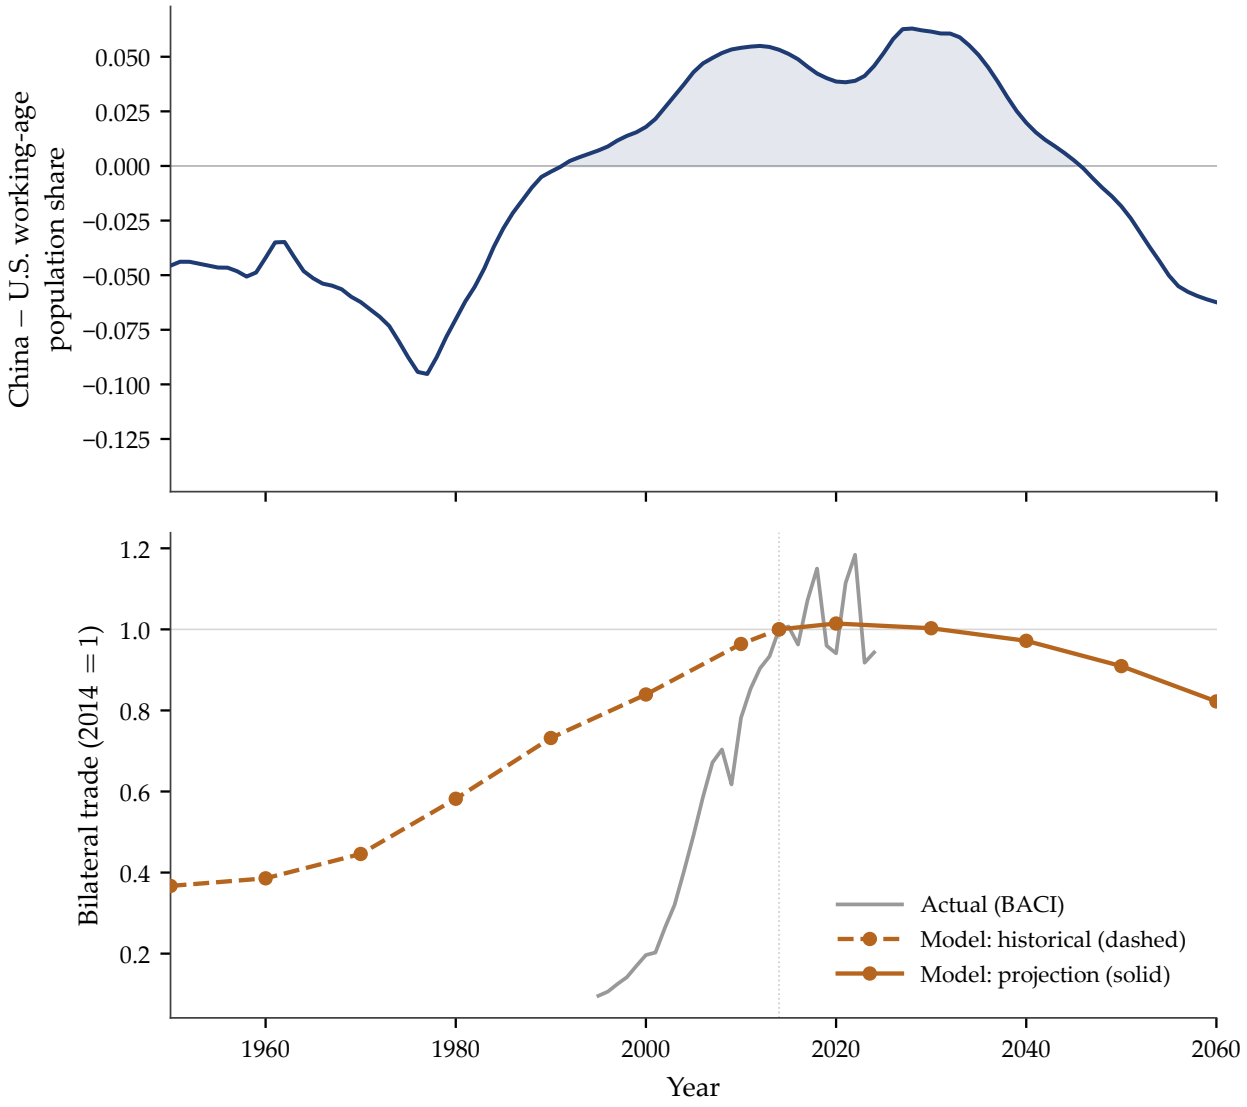


Figure 1: *China–US working-age population shares and bilateral trade. Top: the difference between China's and the US's working-age population shares (ages 20–64), UN WPP 2024. Bottom: model-predicted bilateral trade (dashed for historical counterfactuals, solid for forward projections) against actual BACI trade (grey), normalised to 2014. The gap reflects non-demographic factors; the model isolates the demographic channel.*

of Brakman et al. (2025) and Kopecky (2023): relative working-age population is strongly related to bilateral trade flows.

Composition reallocates the workforce rather than adding to it. Scale raises the number of workers of every kind, while composition only shifts the mix among them. That shift moves relative skill prices, and with them comparative advantage, raising a country's edge in the sectors that use older-age skills and lowering it in those that lean on the young. A gain in one sector comes with a loss in another, and across the corridor the two largely offset. The net effect on the total is therefore small, and second-order compared the size effect. This, however, masks a reallocation underneath that can be large. The composition of the flow across these sectors can swing sharply, as Section 4.2 will show.

Table 1: *Demographic Change in Bilateral Exports (% from a 2014 Baseline)*

	1980–2014			2014–2050		
Corridor	Total	Comp.	Scale	Total	Comp.	Scale
Headline corridors						
CHN → USA	+71.8	+3.3	+66.2	−9.1	+0.6	−9.7
IND → USA	+109.2	+3.6	+102.3	+33.4	−0.5	+34.4
MEX → USA	+111.3	+3.2	+104.8	+27.1	+0.2	+26.9
Other corridors						
JPN → USA	+12.4	+3.1	+9.1	−15.8	−0.1	−15.8
DEU → USA	+21.2	+3.2	+17.5	−6.0	−0.1	−5.9
KOR → USA	+62.7	+3.1	+57.7	−23.2	−0.1	−23.0
IDN → CHN	+146.4	+4.0	+137.3	+16.3	+0.9	+15.2
IND → CHN	+137.3	+4.2	+128.2	+16.1	+0.8	+14.8

Notes: Percentage change in total bilateral manufacturing exports (row exporter to column importer) from demographics alone, holding all non-demographic structure at its 2014 values (services-inclusive 21-sector model). 1980–2014 is a backcast under each period's age structure; 2014–2050 is the forward projection. Total is the full demographic change; its Scale (workforce size) and Comp. (skill mix) components combine multiplicatively and need not sum to it.

4.2. The Composition of Trade

The overall decline masks sharp differences across sectors, and those differences map onto the skill-intensity matrix. Table 2 lists the projected change in China's exports to the United States by sector. Every sector falls as the corridor contracts, but by amounts that track appreciating-skill intensity. Computers and electronics, the most appreciating-intensive manufacturing sector and nearly a third of the corridor, is essentially flat (−0.4%), while

the physical- and depreciating-intensive sectors fall far more, textiles and basic metals by around 17–18% and petroleum, a small-base outlier, by 69%. Figure 2 shows why electronics holds: its near-flat total is a scale-driven decline offset by a composition gain, the shift toward the older workers the sector uses intensively. Isolating that composition channel, the part of each sector's change that comes from the changing age-skill mix rather than the corridor's overall size, the gradient is sharp: appreciating-intensive sectors gain share and physical-intensive sectors lose it. The reallocation is modest next to the volume decline but systematic, concentrated in routine, physically intensive manufacturing and largely absent in high-skill sectors such as electronics.

Table 2: *Sectoral Decomposition: China's Exports to the United States by 2050*

Sector	% Change	Base share (%)	α_j^A
Computer & Electronics	−0.4	28.6	0.59
Machinery	−6.0	9.0	0.23
Furniture & Other Mfg	−9.3	5.9	0.25
Rubber & Plastics	−9.4	3.2	0.18
Non-metallic Minerals	−9.7	3.3	0.18
Electrical Equipment	−10.4	12.7	0.28
Motor Vehicles	−10.6	3.3	0.16
Fabricated Metals	−11.3	4.2	0.19
Chemicals	−13.5	5.8	0.27
Textiles & Apparel	−16.8	15.0	0.13
Basic Metals	−17.7	3.3	0.13
Wood	−17.9	1.0	0.13
Food & Beverages	−20.0	1.4	0.18
Petroleum	−69.4	0.6	0.21

Notes: Percentage change in China's goods exports to the US by sector, 2050 vs. 2014 (services-inclusive 21-sector model), sorted by exposure. α_j^A is the sector's appreciating-skill intensity.

4.3. The Young, the Aging, and the Old

The China–United States corridor is one instance of a pattern that holds across countries, and across time. Figure 3 plots, for each of three 36-year windows between 1950 and 2050, every country's composition tilt, the slope of its sector export changes on appreciating-skill intensity, against how much its workforce ages over the window. Each window feeds its demographic change through the fixed 2014 trade structure, so the panels vary only in demographics and isolate the composition channel rather than trace historical trade. The relationship is the same in all three: a workforce aging faster than the world's tilts

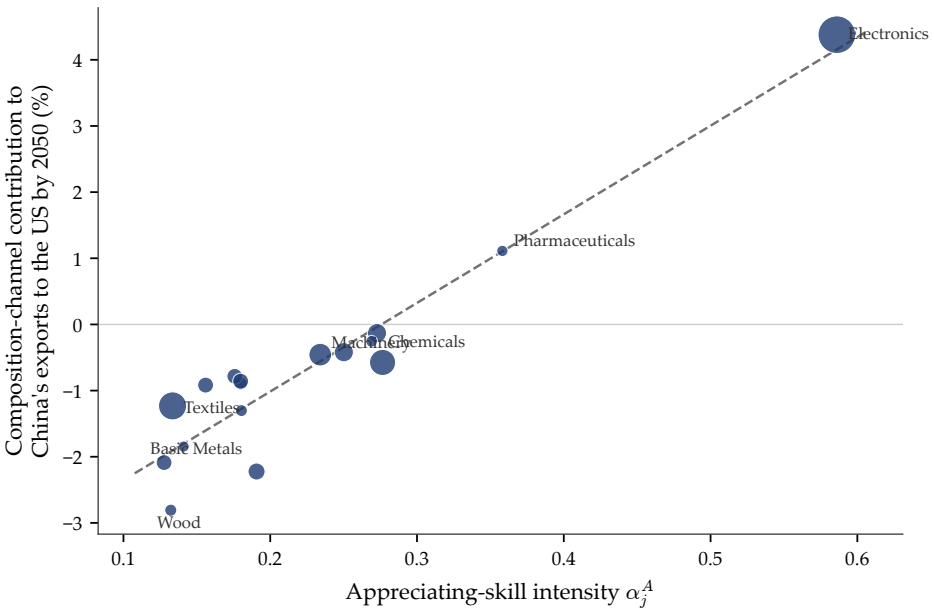


Figure 2: *China's composition-channel sector exports to the United States by 2050 against appreciating-skill intensity α_j^A (the composition channel is defined in the text). Points are tradeable sectors sized by corridor share; the dashed line is the OLS fit ($R^2 = 0.90$), petroleum excluded as a small-base outlier.*

toward the sectors that use age-appreciating skills, one aging slower tilts the other way. Comparative advantage of age is thus relative, and it moves. Japan held it when it led the world's aging in the late twentieth century, tilting sharply toward appreciating sectors; as its own aging stalled and the emerging giants aged past it, the advantage passed to China and its peers. The two cross midway through the century.

The young exporters, still moving into the appreciating-skill ages, reorient their export mix the most, and their volumes keep rising: India, Mexico, and Indonesia. The aging are where the arc turns. China is the clearest: its own aging alone would cut its US exports by 13.4%, more than the corridor's actual 9% fall, because the rest of the world, chiefly the still-growing US workforce, pushes 4.5% back the other way. The old show the two margins pulling apart: Japan's workforce is shrinking fast, so its export volume falls hard. Its own skill supply barely moves, since it has already aged, yet its comparative advantage still rotates toward physical sectors. This highlights the relative nature of life-cycle comparative advantage: the rest of the world's workforce keeps aging while Japan's stands still.

These are also the corridors of the "China-plus-one" reallocation that firms and governments have pursued for geopolitical reasons. The model has none of that machinery: no relocation decisions, no risk motive, and gross flows rather than the value-added rerouting that literature tracks. What it adds is narrower and independent: a demographic force

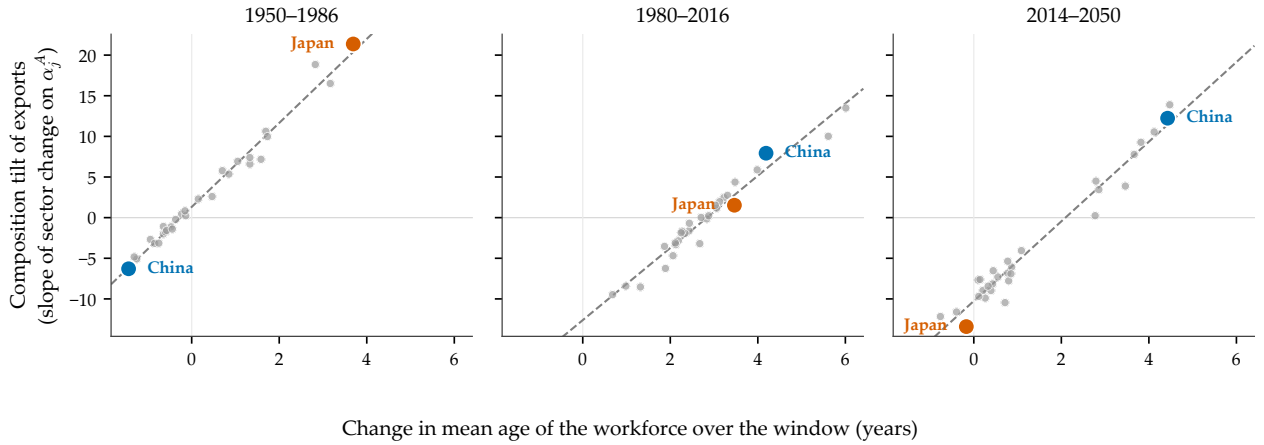


Figure 3: *Composition tilt of exports against the change in workforce mean age, for three 36-year windows between 1950 and 2050. Each point is a country, and the relationship holds in every window (corr $\geq$ 0.98). Japan (orange) and China (blue) cross as the advantage of age passes between them. The panels are comparative-statics experiments on the composition channel, constructed in the text.*

pushing the same corridors in the same direction, through each country's own aging rather than through diversion away from China. The US market share China's aging releases does not reroute to India or Mexico; over half is absorbed by domestic US production, and the rest spreads across suppliers in rough proportion to their size. Estimates of the policy-driven realignment that leave the demographic trend out will fold part of it into the treatment effect.

4.4. The Distributional Costs of Reorientation

The corridors of Section 4.1 and the sector shifts of Section 4.2 are movements in trade; behind them are workers. Figure 4 maps the reorientation onto employment, placing each country's goods sectors at their projected 2050 export change and sizing them by the workers they employ. There are two takeaways from the picture. The first is the employment-weighted mean, marked by the bars: the export change facing the average goods worker. It follows the same demographic gradient as the volume result of Section 4.1, from a 20% decline for the average Korean goods worker to a 46% gain for the average Indian one. The second is the spread within each row, the way the reorientation falls across a country's own sectors. Some economies are almost uniformly exposed, Korea and Japan nearly all in falling sectors, India and Mexico nearly all in rising ones; but many straddle zero, winning and losing sectors side by side, as the comparative advantage of Section 4.2 sorts industries within a single economy.

This is the paper's one direct look at the workers involved, and the counts are large. In

Japan, 97% of the 9.7 million people employed in goods sectors[8] are in industries whose exports fall by more than ten percent; in Korea, 91% of 3.7 million; across the six fastest-aging economies, some 24 million workers in all; and in China, 35 million, nearly a fifth of goods-sector employment. The composition tilt behind these moves is small in the aggregate, and its net welfare effect largely cancels, as I will show in Section 6.1. That cancellation prices a consumption basket; it does not speak to the scale of the labour-market adjustment behind it. Moving this many workers between contracting and expanding sectors carries an adjustment cost the present model is not equipped to price, but the number exposed marks it as a question worth taking seriously.

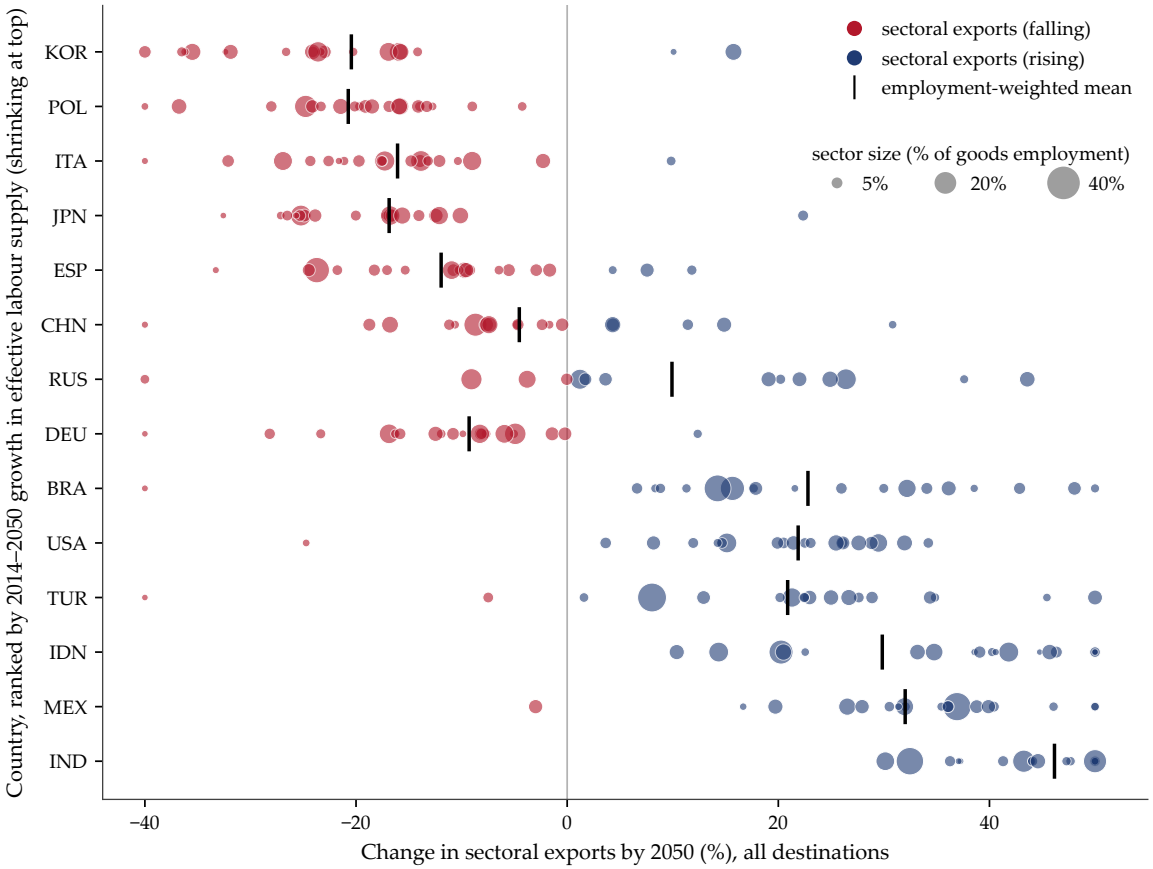


Figure 4: *The incidence of the reorientation. Each row is a model economy, ordered by the 2014–2050 change in effective labour supply, sharpest contraction at top. Employment is from the WIOD Socio-Economic Accounts (2014); petroleum and transport-equipment tails are clipped to* $[-40, +50]$% *for display.*

This reorientation is the model's claim about *what* a country trades. Aging also realigns a second dimension of world trade, the external *balance*, to which I now turn.

[8]Employment is persons engaged in 2014 from the WIOD Socio-Economic Accounts, summed across the model's tradeable goods sectors.

5. The External Balance: Who Lends to Whom

External balances move through the same compositional mechanism as the reorientation itself: the age distribution that reweights fixed skill profiles in Section 2.1 here reweights fixed age profiles of wealth and labour income. I add a minimal saving block to ask whether opening international credit markets alter my headline trade results, explore the type of creditor–debtor *pattern* implied by aging, and whether it's consistent with empirical patterns.

5.1. A compositional saving block

This saving channel has a developed literature: Backus et al. (2014) trace low-frequency capital flows to demography, Krueger and Ludwig (2007) work out its consequences for returns and welfare in an overlapping-generations world, Auclert et al. (2025) quantify the demographic wealth channel and project its implications for the world interest rate and global imbalances, and Sposi (2022) brings it to trade. Relative to that literature the block I add is minimal, a reweighting rather than a full life-cycle problem.

The block relaxes the balanced-trade assumption of Section 2.5. Households of age a hold a fixed profile of net worth $\bar{a}(a)$ and supply a fixed profile of labor income $\bar{y}(a)$; only the age distribution moves with demographics. National desired wealth is

$$A_c(r) = Y_c \cdot s \cdot \omega_c \cdot r^{-\epsilon}, \qquad \omega_c = \frac{\sum_a n_{c,a}\, \bar{a}(a)}{\sum_a n_{c,a}\, \bar{y}(a)}, \tag{14}$$

where ω_c is the demographic wealth-to-income ratio, the object aging moves: the age weights shift toward the high-$\bar{a}$, low-$\bar{y}$ older cohorts, so ω_c rises. Here ϵ is a small interest-(semi)elasticity of desired wealth and s is a scale pinned so the world market clears at the base year. A single world rate r^* equates world desired wealth to world capital demand, $\sum_c A_c(r) = \sum_c K_c(r)$ with $K_c(r) = Y_c\, \kappa_c\, r^{-\sigma}$ and κ_c the country's base capital–output ratio; net foreign assets are $\mathrm{NFA}_c = A_c(r^*) - K_c(r^*)$, positive for a creditor that lends abroad. Both schedules fall in r, so an aging-driven rise in desired wealth lowers r^* only if capital demand is the more interest-elastic, $\sigma > \epsilon$; the calibration maintains this ($\epsilon = 0.3$, $\sigma = 1.0$, both placeholders in Section 5.3). The profiles are calibrated to US data (Survey of Consumer Finances 2022 net worth; National Transfer Accounts labor income); because the base year is a fixed point, their scale is irrelevant and only their shape across ages matters.

The block determines stocks and the world rate; the mapping into trade flows is kept explicit and simple. The wedge that enters goods-market clearing is each country's realized 2014 current account plus a demographic adjustment, the latter an annuitization (at rate

0.04) of the model's demographic change in the net-foreign-asset position; it enters the trade equilibrium exactly as the exogenous imbalance in the unbalanced-trade solve of Section 7. The annuitization rate is a placeholder. The block endogenizes the financial account, not the full life cycle; Section 5.3 collects its limitations.

With the block in place, the reorientation of Section 4 is essentially unchanged. Figure 5 re-runs the 2014–2050 counterfactual with balanced trade replaced by the regime the block implies, in which aging economies run surpluses and young ones borrow. Every bilateral corridor lands almost exactly where it did under balanced trade (correlation 0.994), and the sectoral composition tilts are equally undisturbed (correlation 0.999); the demographic current accounts the block generates are small, at most 1.5% of GDP, far too little to reorganize who exports what. The volume and composition results therefore stand as the balance moves: what follows adds a dimension rather than qualifying one.

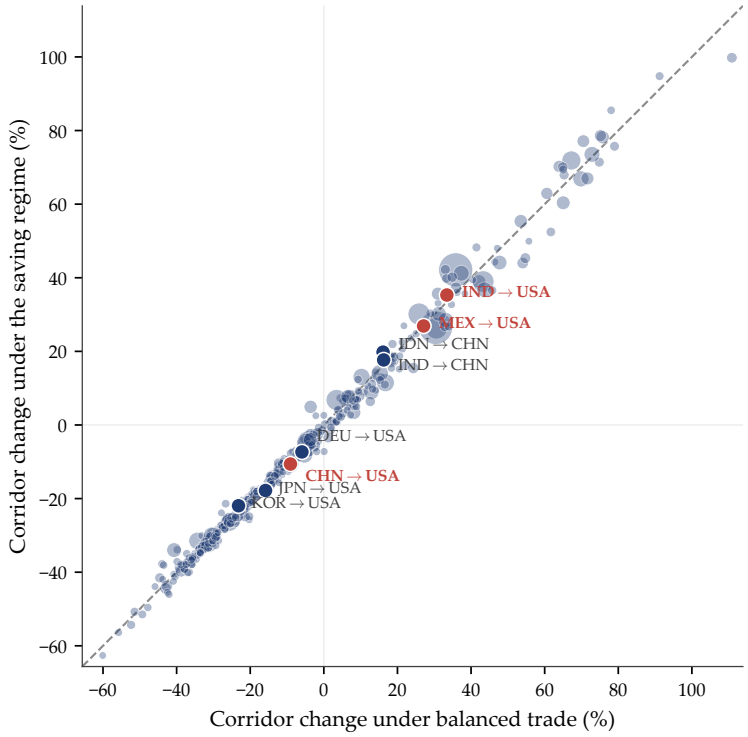


Figure 5: *The Section 4 trade results survive the saving regime. Each point is a bilateral goods corridor, sized by its base-year value, plotting its 2014–2050 percentage change under balanced trade (horizontal) against the change under the endogenous saving regime of this section (vertical). Corridors lie almost exactly on the 45-degree line (correlation* 0.994*).*

5.2. Aging realigns the balance

Forward aging lowers the world interest rate sharply: relative to the 2014 base, normalised to clear at $r^* = 1$, the baseline elasticities put the fall at 35% by 2050, the aging "savings glut." The direction of this fall is disciplined by the demographic shift in desired wealth; its size is not, since it scales with the block's placeholder interest elasticities (Section 5.3), and I quote it to fix ideas rather than as an estimate. Aging also opens a clear creditor–debtor map, and the map is the point. Table 3 reports the implied positions: the aging advanced and East Asian economies, Japan, Korea, Germany, Italy, and China, are pushed toward larger net-creditor positions as they age, while the demographically young, India, Indonesia, and Mexico, are pushed toward deficit. This is the direction the demographic-imbalance literature predicts (Higgins, 1998): aging economies, facing a falling marginal product of capital at home, export capital.

The saving block reproduces the broad cross-country *pattern* of these positions. Figure 6 plots its demographic net foreign assets against observed 2014 positions from the External Wealth of Nations database (Lane and Milesi-Ferretti, 2007, 2018): across the twenty-nine model economies the two correlate $+0.60$. The model's age gradient, $+3.1$ percentage points of NFA-to-GDP per year of mean age in 2014, matches the observed slope of $+3.6$, and the model projects it to steepen to $+5.5$ by 2050 as aging deepens (Table 3).

Two checks qualify what the correlation validates. First, the model is informative beyond the obvious predictor: mean age on its own correlates only $+0.27$ with the observed positions, indistinguishable from zero at this sample size, yet adding the model's position raises the explained variance from 0.07 to 0.53. Second, that content comes mostly from the model's capital input, not its demographic one. A country's model position is its demographic wealth-to-income ratio ω_c minus its capital-output ratio κ_c, a number taken from the data, and almost all of the cross-country agreement with observed positions comes from κ_c; the demographic ω_c, little more than mean age relabeled, adds little on its own. The cross-section therefore validates the model's *positions*, not the demographic saving mechanism specifically. The mechanism's own contribution is the forward *shift* in positions as populations age and ω_c rises, which a single cross-section cannot test.

What the model does not capture is the *level* for economies whose positions are set by non-demographic forces. The clearest case is the United States: a mild creditor on demographics alone, it is in reality the world's largest net debtor, because global demand for safe dollar assets lets it borrow persistently (Caballero, Farhi, and Gourinchas, 2017; Gourinchas and Rey, 2007); financial centres such as Ireland deviate similarly. The demographic channel governs the *supply* of global saving, who is driven to lend, more tightly than it pins down who absorbs it, so I take the cross-country pattern of positions, not the

level for any one country, as what the exercise establishes.

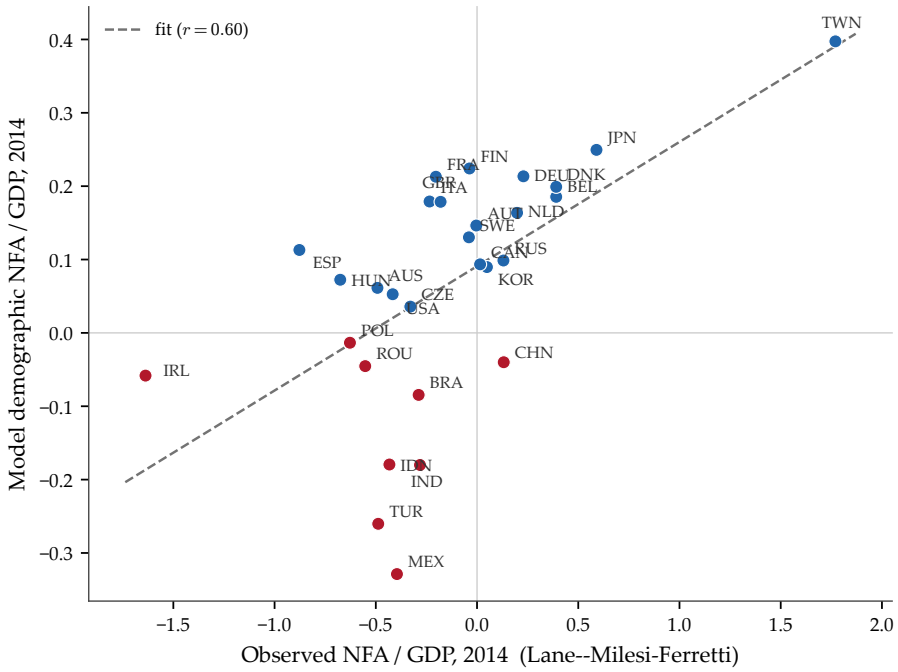


Figure 6: *The model's demographic net foreign assets against observed positions, 2014. Model NFA-to-GDP (the saving block evaluated at 2014 demographics) against observed NFA-to-GDP from the External Wealth of Nations database (Lane and Milesi-Ferretti, 2007, 2018), for twenty-nine model economies. The two correlate +0.60 (Spearman +0.67): aging creditors upper right, young debtors lower left. The United States, Ireland, and Spain sit off the line, their positions set by non-demographic forces the model omits (reserve-currency demand for safe assets, financial-centre activity, convergence-era deficits). The scatter validates the block's combined positions rather than the demographic wealth term in isolation (see text).*

The model's implied current-account *flows* are right-signed but conservative relative to the reduced-form estimates of Sposi (2022); the exercise rests on the positions, not the flows. The pattern is not an artifact of the US saving profiles: replacing them with stylized hump-shaped profiles moves the 2050 world rate from 0.65 to 0.73 and the 2050 position gradient from +5.5 to +4.3 percentage points per year of mean age, the same direction and a similar magnitude (not tabulated).

The two dimensions compose cleanly: they are one decomposition of the single demographic shock. The same moving age distribution delivers the composition result through skill profiles and the balance result through wealth profiles, and because a net imbalance shifts the level of absorption rather than relative comparative advantage, the two can be read additively. Aging reshapes *what* a country trades and *with whom it runs surpluses*. The one thing this composition holds fixed is capital as a stock: the large implied fall in the world interest rate is a capital-market response a full treatment would feed back into the

trade equilibrium.

Table 3: *Net Foreign Asset Positions under Demographic Aging, 2014 vs. 2050*

Country	Mean age '50	NFA/GDP '14	NFA/GDP '50	Change	Aging shifts
Korea	54.4	0.08	0.49	+0.41	creditor
China	50.1	−0.05	0.18	+0.22	creditor
Italy	51.1	0.17	0.35	+0.17	creditor
Japan	51.3	0.24	0.35	+0.11	creditor
Germany	47.4	0.21	0.27	+0.06	creditor
United States	43.1	0.05	0.03	−0.02	debtor
India	38.8	−0.18	−0.21	−0.03	debtor
Indonesia	38.1	−0.18	−0.23	−0.04	debtor
Mexico	39.8	−0.33	−0.43	−0.10	debtor

Notes: The model's demographic component of net foreign assets relative to GDP, base-year vs. 2050 demographics, from the saving block clearing world saving against investment at the endogenous world rate ($r^*(2050) = 0.65$, a 35% fall from the normalized base). "Change" is the demographic shift in the position and "Aging shifts" its sign. The cross-section regression of model NFA/GDP on mean age gives +5.5 percentage points per year in 2050 ($R^2 = 0.71$); Figure 6 validates these positions against observed data. Levels for individual countries also reflect non-demographic factors (Section 5.3).

5.3. Caveats

The saving block is a deliberately minimal model of external flows. Its wealth and labor-income age-profiles are US-only, and its interest elasticities and stock-to-flow annuitization are placeholders; as Section 5.2 noted, the cross-sectional fit also leans on the capital-output ratios more than on the demographic term. None of this is meant as a theory of the current account. The block is here to establish two narrower things: that once demographics are free to push on a saving margin they move positions in the direction the data show, and that letting them do so leaves the trade flows that are this paper's central object almost untouched (Figure 5). A complete mapping of demographic pressure into capital flows is a well-developed literature in its own right. Grafting that machinery onto the Caliendo and Parro (2015) trade model would be costly and, the evidence here suggests, would not overturn the trade results; the one channel that could still bite, the general-equilibrium feedback from the falling world interest rate, is an extension worth exploring that I discuss further in the conclusion.

6. Aggregate Effects: Welfare, Openness, and Policy

Beyond these two positive realignments, the same demographic shock has aggregate counterparts: its consequences for welfare, for the overall scale of trade, and for retirement policy. Each is smaller in magnitude than the realignments themselves, and I present them as decompositions rather than as headline numbers.

6.1. Welfare: A Two-Channel Decomposition

I present the welfare effect as a decomposition, with the services-inclusive (21-sector) model as primary. The decomposition separates two things: the aggregate effect of a changing workforce *size*, and the comparative-advantage effect of a changing skill *composition*. The first is large but holds capital and technology fixed; the second is this paper's object, and it is modest in magnitude but systematic in its cross-country pattern.

The total demographic welfare effects are large: by 2050, holding technology and capital fixed, Japan's real income falls 35%, Korea's 40%, and China's 29%, while India's rises 48% and Mexico's 36%. These totals are almost entirely the workforce-size channel, countries with shrinking working-age populations simply producing less; the skill-composition shift contributes only -0.03% of Japan's -35% (Table 4). Although computed in the trade model, the effect is close to a pure domestic quantity: a shrinking workforce exports less, but under balanced trade it imports less in step, so the loss is the fall in what the country produces rather than a separate penalty from trading less. Because the model holds capital, participation, and technology fixed, the totals are partial-effect labour-supply benchmarks, not welfare forecasts; they sit at the upper end of the macroeconomic aging literature (Cooley et al., 2024; Kitao, 2015), where those offsetting margins are free to move. Appendix E traces their historical counterpart, the demographic dividend in reverse. The paper's object is the composition channel, to which the rest of this section turns.

Table 4 decomposes the composition channel into two opposing forces by valuing the same real-income change through two parts of the consumption basket.[9] For a fast-aging economy the goods-price channel is a *cost*: its traded-goods consumption basket grows more expensive relative to its factor income as world aging reprices the goods it buys. China's goods-price contribution is -0.32%, comparable in scale to the welfare effect of a major trade agreement. The non-traded services-price channel is an offsetting *benefit* of

[9] Welfare is real income: the change in nominal factor income $\sum_k \hat{w}_c^k w_c^k \hat{L}_c^k L_c^k / Y_c$ deflated by the ideal price index $\prod_j (\hat{P}_c^j)^{\eta_c^j}$. Under the composition-only shock the income change is basket-invariant, so the effect splits by re-weighting the deflator: the goods-price column deflates by the goods-only (renormalized) price index, and the services-price column is the residual contribution of the non-traded price; the two sum to the full-basket net.

similar size, +0.47% for China, because the appreciating-intensive services these economies produce and consume become relatively cheaper as the appreciating input grows abundant. The two largely cancel. The goods-only model, which places all consumption weight on traded goods, reports the larger standalone cost (−0.92% for China) because it omits the services benefit entirely. The services channel here connects to the demographic-inflation literature (Juselius and Takáts, 2018; Goodhart and Pradhan, 2020; Kopecky, 2022), but runs opposite to its usual sign: that work emphasizes the demand side, where aging populations consuming more non-traded services push their relative price *up*, whereas here an aging *producer* economy's services grow *cheaper* as it becomes abundant in the skills they use. The fixed expenditure shares of Section 2.4 isolate this supply-side force; the opposing demand-side channel is an age-varying-consumption extension, a promising direction for future work that I discuss in the conclusion.

Table 4: *Composition-Welfare Decomposition, 2014–2050 (%)*

Country	Goods-only comp. (20s)	Goods-CPI cost (21s)	Services-price benefit (21s)	Net full (21s)
China	−0.92	−0.32	+0.47	+0.15
India	−0.93	−0.34	+0.45	+0.11
Mexico	−0.82	−0.30	+0.43	+0.13
Korea	−0.53	−0.13	+0.33	+0.19
Turkey	−0.96	−0.33	+0.50	+0.17
Japan	−0.07	−0.03	+0.00	−0.03
Germany	+0.05	+0.02	−0.05	−0.03
France	+0.09	+0.01	−0.07	−0.06
United States	−0.09	−0.03	+0.05	+0.02

Notes: Composition-only welfare (the skill-mix shift, removing the workforce-size change), 2050 vs. 2014. "Goods-only" is the 20-sector model. The 21-sector decomposition splits the full-basket composition effect into a traded-goods comparative-advantage contribution (through goods prices) and a non-traded services-price contribution; the two right columns sum to the net. The near-cancellation hinges on the non-tradeable consumption weight $\eta^0 \approx 0.79$ and the calibrated services intensity.

The net of the two channels is modest, within roughly ±0.2% for every country, but it is not uniform: it follows the shift in the workforce's age-skill mix over 2014–2050, not the level of a country's age structure. The economies whose mix shifts most see small net *gains*, China +0.15%, Korea +0.19%, India +0.11%, Mexico +0.13%, and Turkey +0.17%, as the services benefit outweighs the goods cost; India and Mexico are young in levels, which is why they appear among the debtors of Section 5, but their workers are moving quickly into the older prime ages, which is what matters here. Several already-old economies

whose aging is largely behind them see small net *losses*, France -0.06% and Germany and Japan -0.03%. The composition channel reshuffles welfare across countries along the same skill-mix shift that drives the sectoral reorientation, even as the aggregate magnitudes stay small.

The near-cancellation, and hence the net sign for any given country, hinges on the large non-tradeable consumption weight and on the appreciating intensity of a services sector that is calibrated, not validated; the π-diagonal sensitivity of Section 7 probes that dependence directly. I therefore do not lean on any single net welfare number. The model is also frictionless: the net prices the change in the consumption basket, not the sectoral reallocation of workers and resources that produces it, which happens here at no cost. What is robust is the structure: two economically meaningful, offsetting channels, a modest net, and a systematic cross-country gradient that mirrors the reorientation.

6.2. Aggregate Openness

A more tentative aggregate result concerns the overall volume of trade. Defined as a balanced-trade-compatible ratio of goods trade to GDP, world openness is roughly flat under the goods-only model ($+0.15\%$ by 2050) but falls by 1.07% once the non-tradeable services sector is included. The mechanism is the same services sink, operating through supply and relative prices rather than demand: expenditure shares are held fixed (Section 2.4), so as populations age it is production and relative prices that tilt toward non-traded services, and the traded share of activity declines. Like the welfare decomposition, this result is driven by the calibrated services sector and is mixed across countries, pronounced for some aging economies and weak or absent for others; I present it as a demographic contribution to the widely discussed slowdown in trade growth, not as its engine.

6.3. Retirement-Age Policy

A natural policy response to workforce shrinkage is to raise the effective retirement age. I implement a stylized version, extending the working-age population to include the 65–69 group. Because the calibration stops at 60–64, I set the 65–69 population to 90% of the 60–64 cohort and extrapolate each input's age-efficiency one five-year step beyond 60–64 by its most recent slope, floored at a small positive value, so the added cohort supplies the far end of the appreciating-skill profile. The aggregate gains are large: Japan recovers 11.3 percentage points of its welfare loss, Germany 8.7, the United States 8.6, and Korea 7.0. Almost all of this is the workforce-size channel of Section 6.1: more workers, more output.

The point of interest here is the *composition* channel, and it is small. Like the main

welfare result, its sign depends on the consumption basket: in the goods-only model it is slightly negative (from -0.17% for Korea to -0.37% for the United States), and in the services-inclusive model slightly positive (around $+0.09\%$ to $+0.14\%$); in absolute value it never exceeds 0.4% and is tiny next to the workforce-size gain, which is itself essentially unchanged across the two models. The substantive point is basket-independent: extending the retirement age recovers a great deal of workforce-size welfare but barely moves the comparative-advantage shift that is this paper's subject. Adding workers at the far end of the appreciating-skill profile does not meaningfully rebalance an aging economy's skill mix toward the global average. Retirement extension is thus an effective answer to the workforce-size problem but not to the competitiveness problem; addressing the latter would require changing the effective skill mix, either by rebalancing the age structure through younger-worker immigration or by retraining that shifts the mapping from age to task inputs, a margin the model holds fixed.

7. Robustness

I run a large battery of robustness tests on the model's headline results, collected in Table 5: each row is a different perturbation, and the columns report the four headline corridors under it. The point estimates move from row to row, by more under some tests than others, but the relationship among the corridors holds throughout. Across all 870 corridors, each test's changes correlate with the baseline at 0.97 or above, and the headline corridors keep their sign and ranking under every one.

Balanced trade. Section 5 shows the corridors are essentially unchanged under the model's endogenous saving regime, but aging in my context can only account imperfectly for large external imbalances. Model results are equally robust to the literature-standard alternative of carrying the realized 2014 current-account imbalances as fixed exogenous transfers (Dekle et al., 2008; Caliendo and Parro, 2015): the corridor changes then correlate 0.996 with the balanced baseline across all 870 corridors, every sign and ranking preserved and the result insensitive to how the world bond market is cleared (Table 5).

Differential skill profiles. The baseline collapses the depreciating-cognitive and physical inputs into a common decline, since the two are nearly collinear in the data. Appendix A re-solves the model under a battery of alternatives that differentiate them and vary their steepness. This matters for the welfare arithmetic: the composition-welfare magnitude moves from 0.55 to 1.54 times the baseline, a wide enough range that how the decline is split is a genuine source of uncertainty about that number. It does not matter for the trade.

Table 5: *Robustness of the corridor reorientation, 2014–2050 (% change in bilateral goods exports)*

Perturbation	CHN→US	IND→US	MEX→US	JPN→US
Baseline (headline)	−9.1	33.4	27.1	−15.8
External imbalances (Dekle–Eaton–Kortum)	−10.4	34.5	26.7	−17.4
Trade elasticity $\theta = 3$	−9.3	33.7	32.1	−14.7
Trade elasticity $\theta = 8$	−10.4	26.8	27.3	−16.0
Skill intensity: simple average	−9.0	33.4	27.1	−15.8
Skill intensity: unweighted	−9.3	34.0	26.9	−15.8
Skill dispersion $\lambda = 0.5$ (compress)	−9.6	33.5	26.7	−15.8
Skill dispersion $\lambda = 1.5$ (stretch)	−8.7	33.5	27.4	−15.8
Age slope: appreciating $\times \frac{1}{2}$	−9.2	33.3	26.9	−15.8
Age slope: appreciating $\times 2$	−8.7	33.7	27.5	−15.9
Services 10% tradeable	−13.6	39.6	24.9	−16.1
Zero migration	−16.5	28.7	20.7	−25.0

Notes: Each row re-solves the 2014–2050 counterfactual under the stated perturbation on the committed 21-sector model, reproducing the baseline corridors before perturbation to under 0.05 percentage points. Across all 870 bilateral corridors each perturbation's changes correlate with the baseline between 0.97 (trade elasticity) and above 0.999; magnitudes move most under services tradeability and zero migration. The composition-welfare net (not shown) is the fragile object: its sign flips under skill-dispersion compression ($\lambda = 0.5$: 1 of 12 countries) and the dropping of employment weights, and holds for 10 of 12 countries under zero migration.

Corridor magnitudes, their signs, and the cross-country ordering are left intact, and with them the scale of reallocation, because the reorientation runs on workforce scale rather than on how the declining inputs are divided.

Age-profile slopes. The size of a corridor's move depends on how steeply skills appreciate and decline with age: a steeper life-cycle profile turns the same demographic shift into a larger movement in effective skill supply, and so a larger corridor change. To see how much the magnitudes owe to the slopes I calibrate, I vary those slopes directly: the appreciating profile from 0.25 to 1.0% per year (half to double the baseline) with the declining profiles held at baseline, and the common declining slope scaled by ±50%. Every corridor keeps its sign in every variant, and the ranking is preserved up to a swap of the two near-tied corridors into China; magnitudes move within a narrow band, at most 3.9 percentage points, at India–US against its +33% baseline (Table 5). The reorientation is not an artifact of the assumed slopes; their steepness prices its magnitude within a band of a few points.

The productivity elasticity. Across the grid $\rho \in [0.1, 3.0]$ the model's composition-only appreciating-skill export-share gradient (ϕ_A run on model output, Appendix B) stays small and positive and the welfare signs are stable, because ρ scales only the Ricardian

productivity channel, while the factor-price channel that carries the response is independent of it.

Trade elasticity. Setting every sector's trade elasticity to a common value across $\{3, 4.5, 8\}$, which brackets the sector-specific baseline (median near 5), leaves the reorientation essentially unchanged: the corridor changes correlate 0.97 with the baseline at every value, their ranking is preserved, and every corridor keeps its sign (Table 5). The composition-welfare sign is stable for all twelve sample countries at every elasticity. The one visible sensitivity is that the South–South corridors into China shrink at a common high elasticity, since they sit on high-elasticity sectors in the baseline, while keeping their sign and rank.

Skill-intensity construction and dispersion. How the skill-intensity matrix is built barely touches the corridors, and largely cannot. Replacing the principal-components aggregation with a simple average, or dropping employment weights, re-weights the same O*NET indices without reordering the sectors, and the corridor estimates barely move (Table 5). The intensities' cardinal spread is the same story: compressing or stretching each input's dispersion around its cross-sector mean by up to a factor of two (λ from 0.5 to 1.5, the base year re-bootstrapped so each variant clears) rescales the appreciating-skill gradient of Figure 2 but not its order, moving its slope while its fit holds (R^2 between 0.50 and 0.62, against 0.58 at $\lambda = 1$). The one object that genuinely responds is the composition-welfare sign: robust to the aggregation choice, but flipping if employment weights are dropped or the dispersion is compressed ($\lambda < 1$), though never exceeding 0.6% in magnitude anywhere on the grid.

Services tradeability. The welfare decomposition does not require services to be strictly non-tradeable. Making the services sector ten percent tradeable (each destination's services π-diagonal set to 0.90, the remaining 0.10 spread over foreign suppliers in proportion to their shares of that destination's goods imports, and the base re-bootstrapped so the perturbed model still satisfies the accounting identity) preserves the two-channel structure: the services-price benefit remains positive for eleven or twelve of the twelve sample countries and the net composition effect stays small and heterogeneous. What moves is corridor *magnitudes*, by up to ten to twenty percentage points, with every sign preserved (Table 5). I read this as sensitivity to one stark perturbation rather than as a bound: the baseline home shares already embed home bias, so there is no clean zero-cost benchmark against which an upper bound could be defined. Direction is robust to the corner; magnitude is not.

Migration. Migration is the demographic margin with the largest trade and welfare consequences in the aging literature (di Giovanni, Levchenko, and Ortega, 2015), and here it enters only through the population projection. Replacing the medium-variant demographics with the UN zero-migration variant for 2050 leaves the corridor changes correlated 0.98 with the baseline and every sign preserved (Table 5). The economically interesting movement is in the corridors into the United States, which become more negative (China–US from -9% to -16%, Japan–US from -16% to -25%): without immigration the US working-age population stops growing, so it imports less from every source. The composition-welfare sign is stable for ten of the twelve countries, the exceptions again being the near-zero Japan and Germany.

Taken together these checks justify the paper's emphasis. The *directional* reorientation, the corridor signs and ranking and the sectoral pattern, is robust to every perturbation tried: dropping the non-tradeable sector, realistic external imbalances, the age-profile slopes, the trade elasticity, the skill-intensity construction and its cardinal dispersion, services tradeability, and migration. Corridor *magnitudes* are more sensitive to the most aggressive of these, services tradeability and zero migration, and the sign of the small composition-welfare net is the most sensitive object of all.

8. Conclusion

Demographic change is realigning world trade. Two forces drive the realignment: divergent workforce growth moves the volume of trade along the major corridors, while the shifting age-skill mix moves its sectoral composition, a comparative advantage of age that tilts each country toward the industries its changing workforce favours. One sets how much crosses each border, the other what crosses it, and together they realign both the size and the shape of world trade. The paper's primary contribution is to quantify the two together in a single general-equilibrium model, and to trace them across the demographic transition from 1950 to 2050. Allowing wealth to operate through the same life-cycle mechanism provides a pattern of external balances, the aging advanced and East Asian economies turning into creditors and the young into debtors, broadly in line with observed positions, while reinforcing the central trade findings. The net welfare implications of the reorientation are small, the product of two forces that nearly cancel: a cost from traded-goods prices set against a benefit from non-traded services prices of similar size.

Several limitations mark the natural next steps. The external-balance result rests on a deliberately minimal saving block: US-only life-cycle profiles, placeholder elasticities, and a stock-based capital demand whose cross-country fit leans more on capital-output ratios

than on demographics. A fuller model would give each country its own profiles and let saving and capital accumulate endogenously, and two channels I held outside the trade equilibrium would then reopen. The first is the general-equilibrium feedback from the sharply falling world interest rate, which a complete treatment would feed back into trade. The second is capital as a factor of production, which would open a capital–skill reallocation channel. The welfare of these imbalance and capital channels, taking the realignments here as given and asking what they are worth, is the natural sequel to this paper.

The model's equilibrium is frictionless. This reorientation reallocates workers across sectors at no cost, tens of millions of them in the fastest-aging economies. Pricing that labour-market adjustment, and its incidence, is a question these results raise but cannot answer, and the answer could be a large one: the China shock is a reminder that trade-driven reallocation can be slow and costly for the workers it displaces (Autor et al., 2013). Two remaining extensions are more mechanical, but could add to the findings rather than merely qualify them. First, demand is independent of age; an age-varying shift of consumption toward non-tradeable services would add a demand-side channel to the supply-side mechanism. Second, the thirty-economy WIOD sample collapses the developing world into a rest-of-world composite. The larger OECD inter-country tables could break that composite into named economies, sharpening the in-sample test and enriching the trade map, with the rapidly growing Sub-Saharan Africa the most important omission.

The demographic transition is not only a domestic challenge of pensions and healthcare. It is a force already realigning the global pattern of trade, in what is made where and in who lends to whom. Understanding how far demographics drives that realignment will only grow more important across the rest of the century.

References

Acemoglu, D. and P. Restrepo (2022). Demographics and automation. *Review of Economic Studies 89*(1), 1–44.

Auclert, A., H. Malmberg, F. Martenet, and M. Rognlie (2025). Demographics, wealth, and global imbalances in the twenty-first century. *Review of Economic Studies*.

Autor, D. H., D. Dorn, and G. H. Hanson (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review 103*(6), 2121–2168.

Backus, D., T. Cooley, and E. Henriksen (2014). Demography and low-frequency capital flows. *Journal of International Economics 92*, S94–S102.

Bloom, D. E., D. Canning, and G. Fink (2010). Implications of population ageing for economic growth. *Oxford Review of Economic Policy 26*(4), 583–612.

Brakman, S., T. Kohl, and C. van Marrewijk (2025). DemoGravity: World population and trade in the 21st century. *Review of International Economics 33*(2), 486–501.

Burstein, A. and J. Vogel (2017). International trade, technology, and the skill premium. *Journal of Political Economy 125*(5), 1356–1412.

Caballero, R. J., E. Farhi, and P.-O. Gourinchas (2017). The safe assets shortage conundrum. *Journal of Economic Perspectives 31*(3), 29–46.

Cai, J. and A. Stoyanov (2016). Population aging and comparative advantage. *Journal of International Economics 102*, 1–21.

Caliendo, L., M. Dvorkin, and F. Parro (2019). Trade and labor market dynamics: General equilibrium analysis of the China trade shock. *Econometrica 87*(3), 741–835.

Caliendo, L. and F. Parro (2015). Estimates of the trade and welfare effects of NAFTA. *The Review of Economic Studies 82*(1), 1–44.

Cooley, T. and E. Henriksen (2018). The demographic deficit. *Journal of Monetary Economics 93*, 45–62.

Cooley, T. F., E. Henriksen, and C. Nusbaum (2024). Demographic obstacles to European growth. *European Economic Review 169*, 104829.

Cravino, J., A. A. Levchenko, and M. Rojas (2022). Population aging and structural transformation. *American Economic Journal: Macroeconomics 14*(4), 479–498.

de Zwart, B. C. H., M. H. W. Frings-Dresen, and Å. Kilbom (2001). Gender differences in upper extremity musculoskeletal complaints in the working population. *International Archives of Occupational and Environmental Health 74*(1), 21–30.

Dekle, R., J. Eaton, and S. Kortum (2008). Global rebalancing with gravity: Measuring the burden of adjustment. *IMF Staff Papers 55*(3), 511–540.

di Giovanni, J., A. A. Levchenko, and F. Ortega (2015). A global view of cross-border migration. *Journal of the European Economic Association 13*(1), 168–202.

Eaton, J. and S. Kortum (2002). Technology, geography, and trade. *Econometrica 70*(5), 1741–1779.

Feenstra, R. C., R. Inklaar, and M. P. Timmer (2015). The next generation of the penn world table. *American Economic Review 105*(10), 3150–3182.

Gaulier, G. and S. Zignago (2010). BACI: International trade database at the product-level. the 1994–2007 version. Working Paper 2010-23, CEPII.

Goodhart, C. and M. Pradhan (2020). *The Great Demographic Reversal: Ageing Societies, Waning Inequality, and an Inflation Revival*. Cham: Palgrave Macmillan.

Gourinchas, P.-O. and H. Rey (2007). From world banker to world venture capitalist: Us external adjustment and the exorbitant privilege. In R. H. Clarida (Ed.), *G7 Current Account Imbalances: Sustainability and Adjustment*, pp. 11–66. University of Chicago Press.

Gu, K. and A. Stoyanov (2019). Skills, population aging, and the pattern of international trade. *Review of International Economics 27*(2), 499–519.

Higgins, M. (1998). Demography, national savings, and international capital flows. *International Economic Review 39*(2), 343–369.

Juselius, M. and E. Takáts (2018). The enduring link between demography and inflation. BIS Working Paper 722, Bank for International Settlements.

Kitao, S. (2015). Fiscal cost of demographic transition in Japan. *Journal of Economic Dynamics and Control 54*, 37–58.

Kopecky, J. (2022). Okay boomer... excess money growth, inflation, and population aging. *Macroeconomic Dynamics*. Volume/pages to verify.

Kopecky, J. (2023). Growing older and growing apart? Population age structure and trade. *Journal of Economic Studies 50*(8), 1694–1709.

Krueger, D. and A. Ludwig (2007). On the consequences of demographic change for rates of return to capital, and the distribution of wealth and welfare. *Journal of Monetary Economics 54*(1), 49–87.

Lagakos, D., B. Moll, T. Porzio, N. Qian, and T. Schoellman (2018). Life-cycle wage growth across countries. *Journal of Political Economy 126*(2), 797–849.

Lane, P. R. and G. M. Milesi-Ferretti (2007). The external wealth of nations mark ii: Revised and extended estimates of foreign assets and liabilities, 1970–2004. *Journal of International Economics 73*(2), 223–250.

Lane, P. R. and G. M. Milesi-Ferretti (2018). The external wealth of nations revisited: International financial integration in the aftermath of the global financial crisis. *IMF Economic Review 66*(1), 189–222.

Maestas, N., K. J. Mullen, and D. Powell (2023). The effect of population aging on economic growth, the labor force, and productivity. *American Economic Journal: Macroeconomics 15*(2), 306–332.

Mundlak, Y. (1978). On the pooling of time series and cross section data. *Econometrica 46*(1), 69–85.

National Center for O*NET Development (2024). O*net 29.1 database. U.S. Department of Labor, Employment and Training Administration. https://www.onetcenter.org/database.html.

Salthouse, T. A. (2009). When does age-related cognitive decline begin? *Neurobiology of Aging 30*(4), 507–514.

Schaie, K. W. (2005). *Developmental Influences on Adult Intelligence: The Seattle Longitudinal Study*. New York: Oxford University Press.

Sposi, M. (2022). Demographics and the evolution of global imbalances. *Journal of Monetary Economics 126*, 1–14.

Timmer, M. P., E. Dietzenbacher, B. Los, R. Stehrer, and G. J. de Vries (2015). An illustrated user guide to the world input–output database: The case of global automotive production. *Review of International Economics 23*(3), 575–605.

United Nations (2024). World population prospects: The 2024 revision. Technical report, Department of Economic and Social Affairs, Population Division, New York.

U.S. Bureau of Labor Statistics (2014). Occupational employment and wage statistics, may 2014. U.S. Department of Labor. https://www.bls.gov/oes/.

A. Differential Skill Profiles (Invariance)

The baseline imposes a common declining profile on the depreciating-cognitive and physical inputs, which are nearly collinear in the data. To check that this is innocuous, I re-solve the model under a battery of alternative profiles that differentiate the two inputs and vary their steepness, including profiles that flatten or steepen each input and that preserve physical capacity into later life. Across these alternatives the composition-welfare magnitude ranges from 0.55 to 1.54 times the baseline (a factor of roughly 2.8 from the gentlest to the steepest), but the sign of every country's effect and the cross-country ordering are preserved. The differentiation therefore affects only the magnitude of the composition-welfare decomposition, leaving the sign of every country's effect and the cross-country ordering intact. It does not disturb the corridor reorientation, which Section 4.1 shows is carried almost entirely by the workforce-scale margin and so is insensitive to how the declining-input composition is split. This is why the baseline collapses the two declining inputs.

B. The Reduced-Form Life-Cycle Comparative Advantage

The mechanism this paper models leaves a clear footprint in the raw trade data: across countries, older workforces export relatively more in appreciating-skill-intensive sectors. This appendix documents that footprint. The estimates are descriptive conditional correlations, not causal effects: demographic structure moves slowly and travels with income, institutions, capital, and policy, none of which the fixed effects fully strip out. They motivate the exercise; nothing in the paper rests on the magnitudes reported here.

With that understood, the evidence is as follows. The closest antecedent, Cai and Stoyanov (2016), established the relationship on bilateral trade flows; I document it on the export-composition object the model speaks to, across a long panel. Trade data come from BACI (CEPII) (Gaulier and Zignago, 2010), aggregated to 19 manufacturing sectors (ISIC Revision 4) for 204 countries over 1995–2024; the dependent variable is a sector's share of a country's manufacturing exports. Demographic data are from UN WPP 2024 (United Nations, 2024), and the sector skill intensities are the O*NET-based α_j^k of Section 3.1, concorded to this coarser BACI/ISIC aggregation (the structural model uses twenty WIOD manufacturing sectors). The specification regresses the export share on the interaction of sector skill intensity with country demographic structure, absorbing exporter–sector and year fixed effects, so the variation comes from within-country demographic change over

the panel:

$$s_{ct}^{j} = \sum_{k} \phi_k \, \alpha_j^k \, \text{Demo}_{ct} + \mu_{cj} + \delta_t + \varepsilon_{ct}^{j}. \tag{15}$$

Table 6 reports the estimates. In the baseline (column 1), a higher share of workers aged 40–64 is associated with higher export shares in appreciating-intensive sectors ($\phi_A = 0.553$) and lower shares in depreciating-intensive sectors ($\phi_D = -0.327$); because the three intensities sum to one, physical intensity is the omitted reference and these coefficients are read relative to it. The most credible cut is the long difference (column 5): regressing the 1995–2024 change in shares on the change in demographics gives 0.761, and because demographic change is glacial it is best seen over long horizons. The association is stable across demanding fixed effects. Replacing year with exporter–year fixed effects, so the variation is purely cross-sector within a country-year, leaves it essentially unchanged (column 6, 0.554), and adding a Heckscher–Ohlin capital interaction does not displace it (column 7). A Mundlak (1978) decomposition (columns 3–4) splits the interaction into within-country and between-country components, 0.567 and 1.145, both precisely estimated, so the regularity holds both over time within countries and across countries in the cross section; this specification uses sector and year fixed effects, since exporter–sector effects would absorb the between component. The age-group split (columns 8–9), the annual-panel estimates, and a fractional-logit treatment of the share dependent variable all leave it intact. None of it sharpens into a causal claim; it establishes that the pattern is real and robust, which is all the model takes from it. Figure 7 shows the relationship.

C. Services Calibration

The non-tradeable sector aggregates WIOD industries 24–56. Its domestic trade share is forced to unity ($\pi_{cc}^0 = 1$), and consumption shares are renormalised over the 21 sectors, giving a non-tradeable consumption weight η_c^0 that averages 0.79 (range 0.65 to 0.88). The build preserves the base-year accounting identity: forced-domestic services output equals services expenditure to numerical precision, with the base-year fixed point converging to roughly 10^{-13} and the identity test passing at 10^{-8}. The manufacturing skill intensities are unchanged from the 20-sector calibration; the α block for the tradeable sectors is byte-identical across the two builds, so the 21-sector results add the services sector without disturbing the goods block. The change in the goods-channel magnitudes between the two models is therefore pure general equilibrium plus the added sector, not a renormalisation artifact. The age-varying services *demand* channel, through which aging would shift consumption toward non-tradeables, is left for future work.

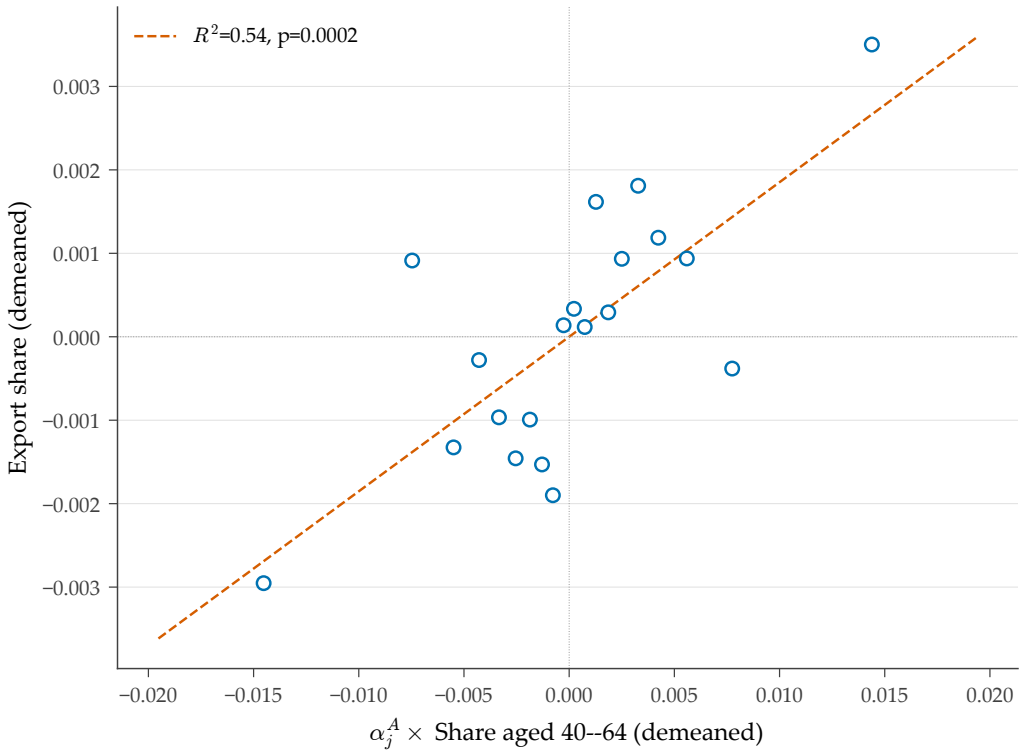


Figure 7: *Binscatter of sector export shares against the interaction of sector appreciating-skill intensity and country share aged 40–64, conditional on exporter–sector and year fixed effects. Twenty equal-sized bins; the dashed line is the OLS fit. Data: BACI 1995–2024, 204 countries, 19 sectors, 112,632 observations.*

Table 6: *Life-Cycle Comparative Advantage: Main Rybczynski Results*

	Baseline		Mundlak		Long	Robustness		Age-group split	
	(1) share$_{40+}$	(2) Mean age	(3) Within	(4) Between	(5) diff.	(6) Exp–yr FE	(7) +Capital	(8) All	(9) OECD
$\alpha_j^A \times$ Demo	0.553***	0.020***	0.567***	1.145***	0.761**	0.554***	0.433**		
	(0.150)	(0.005)	(0.148)	(0.150)	(0.329)	(0.150)	(0.176)		
$\alpha_j^D \times$ Demo	−0.327***	−0.012***	−0.344***	−0.736***	−0.526**				
	(0.101)	(0.003)	(0.100)	(0.100)	(0.212)				
share$_{20–39} \times \alpha_j^A$								0.305***	0.481***
								(0.075)	(0.135)
share$_{40–64} \times \alpha_j^A$								0.457***	0.555***
								(0.097)	(0.180)
Fixed effects	Exp–sec, Year	Exp–sec, Year	Sector, Year	Sector, Year	Sector	Exp–sec, Exp–yr	Exp–sec, Exp–yr	Exp–sec, Year	Exp–sec, Year
N	112,632	112,632	112,632	112,632	3,462	112,632	80,187	112,632	21,508

Notes: Dependent variable is sector export share. Standard errors clustered by exporter in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The demographic variable is share$_{40–64}$ (mean age in column 2). (3)–(4): Mundlak (1978) within/between decomposition, estimated jointly with sector and year fixed effects (exporter–sector effects would absorb the between term). (5): long difference, 1995–2024. (6): exporter–year fixed effects absorb all country–year variation. (7): adds sector capital share × country K/L ratio. (8)–(9): both age-group shares interacted with α_j^A; (9) restricts to OECD. Sample: BACI, 1995–2024, 204 countries, 19 sectors.

D. Data Sources

BACI (CEPII). Bilateral trade at the HS 6-digit level, 1995–2024, aggregated to 19 ISIC Revision 4 manufacturing sectors; 204 exporters per year (Gaulier and Zignago, 2010).

UN World Population Prospects 2024. Population by five-year age group for 237 countries, 1950–2100; medium-variant projections for the forward counterfactuals (United Nations, 2024).

World Input–Output Database 2016. Inter-country input–output tables for 43 countries and 56 industries; I extract 2014 trade shares, input–output coefficients, value-added shares, and consumption shares, aggregated to 30 countries and 20 manufacturing plus one non-tradeable sector (Timmer et al., 2015).

O*NET 29.1 and BLS OEWS 2014. Ability importance ratings by occupation and employment counts by occupation within 4-digit NAICS manufacturing industries, combined into industry-level skill intensities (National Center for O*NET Development, 2024; U.S. Bureau of Labor Statistics, 2014).

Penn World Table 10.01. Capital stocks and employment (Feenstra, Inklaar, and Timmer, 2015), used for the capital-intensity control in Table 6.

External Wealth of Nations. Net international investment positions (year-end 2014 vintage), used to validate the model's demographic net foreign asset positions in Section 5.2 (Lane and Milesi-Ferretti, 2007, 2018).

E. Historical Welfare Decomposition

Table 7 reports the counterfactual welfare ratio $\hat{W}_c$ under each historical demographic vintage, relative to the 2014 baseline. The values trace the demographic dividend in reverse: the now-aging fast transitioners (India, China, Korea) would have had far lower welfare under their younger historical age structures, while the mature economies sit closer to one, with the United States, still demographically young in 1970, an intermediate case.

Table 7: *Historical Welfare Decomposition: $\hat{W}_c$ Relative to 2014 Baseline*

	1970	1980	1990	2000	2010
Large demographic transitions					
India	0.235	0.332	0.471	0.655	0.895
Mexico	0.223	0.328	0.495	0.698	0.909
Indonesia	0.259	0.360	0.529	0.730	0.928
China	0.291	0.423	0.630	0.802	0.962
Korea	0.333	0.489	0.724	0.864	0.940
Mature economies					
Japan	0.861	0.983	1.052	1.123	1.062
Germany	0.791	0.841	0.988	1.012	0.986
USA	0.458	0.600	0.723	0.830	0.957
France	0.652	0.745	0.861	0.924	1.000
UK	0.711	0.745	0.830	0.886	0.979

Notes: Baseline calibration ($\rho = 0.5$). Each entry is the counterfactual welfare ratio $\hat{W}_c$ when the 2014 trade structure is held fixed but country c's demographics are replaced by the indicated year's; the 2014 baseline is normalised to one. Values below one indicate that the historical age structure would have yielded lower welfare than 2014.