The Comparative Advantage of Age

Joseph Kopecky*

5th April 2026

Abstract

I develop a quantitative multi-country, multi-sector trade model in which population age structure shapes comparative advantage through age-dependent skills. Workers of different ages supply heterogeneous bundles of cognitive and physical abilities that appreciate or depreciate over the life cycle, and sectors use these skills with differing intensities. I embed this mechanism into a Ricardian trade model calibrated to 30 countries and 20 manufacturing sectors. Reduced-form evidence from a panel of 204 countries (1995–2024) documents a "grey advantage": a country's share of older workers is associated with a shift in its export mix toward appreciating-skill-intensive sectors. Evaluated in partial equilibrium, the calibrated model recovers 84% of this empirical relationship. General equilibrium adjustment then compresses the trade-composition response by a factor of ten, as endogenous wage and price changes absorb the sectoral reallocation but generate welfare-relevant real-income effects. Forward projections decompose the total demographic effect into a workforce-size channel (shrinking populations produce less) and a skill-composition channel (aging shifts the mix of skills, altering sectoral comparative advantage). The composition channel generates welfare effects ranging from -0.9% to $+0.05\%$, comparable to standard trade-policy shocks, and exhibits a striking temporal pattern: historical demographic divergence supported positive composition gains, but as countries' age structures converge toward a common older profile, these gains are reversing. Bilateral trade flows reallocate accordingly, with the model projecting that China–US trade will fall by 10% and India–US trade will rise by 34% in the coming decades.

JEL classification codes: F11, F14, F17, J11

Keywords: Comparative advantage, Population aging, Trade, Demographics, Ricardian model

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

1. Introduction

The demographic transition is a slow-moving and largely predictable force that looks to reshape the global economy. The world's population aged 65 and older will more than double between 2024 and 2050, rising from roughly 800 million to over 1.6 billion (United Nations, Department of Economic and Social Affairs, Population Division, 2024). Working-age populations are already shrinking in much of the developed world. Countries like Japan and Italy have been well-known bellwethers for aging, but much of the rest of the developed world is rapidly joining the club. China, for example, saw its working-age population peak in 2015 and is projected to decline by nearly 200 million by mid-century. At the same time, countries like India, Indonesia, Mexico, and much of Sub-Saharan Africa continue to experience rapid growth in the numbers of young workers entering the labour force. The consequences of these shifts for domestic labour markets, public finances, and economic growth have received enormous attention (Bloom, Canning, and Fink, 2010; Maestas, Mullen, and Powell, 2023). Their consequences for international trade have not.

This paper argues that they should. Its main insight is relatively straightforward: workers of different ages supply different skills. A large body of evidence in cognitive neuroscience and gerontology documents that certain cognitive abilities appreciate with age and work experience. These include verbal fluency, vocabulary, and comprehension, which tend to peak in the mid-50s or later (Salthouse, 2009; Schaie, 2005). At the macro level, Feyrer (2007) demonstrates that the age composition of the workforce, particularly the share of workers aged 40–49, has a substantial effect on aggregate productivity. Other abilities, such as processing speed, divided attention, and perceptual acuity, begin declining from the late 20s onward. Physical capacities like strength, stamina, and manual dexterity, follow a similar depreciating trajectory (**?**Skirbekk, 2004). Because sectors differ systematically in their reliance on these skill types, a country's age structure determines its vector of effective skill endowments, which in turn determines its pattern of sectoral comparative advantage.

Consider two stylized countries, one young and one old, that are otherwise identical. The young country is abundantly endowed with age-depreciating cognitive skills and physical abilities: the skills that peak early in the working life. It will therefore have a comparative advantage in sectors that rely intensively on these skills, sectors such as routine assembly, machine operation, and manual production. The old country, conversely, is abundantly endowed with age-appreciating cognitive skills (communication, comprehension, and managerial judgment). The older economy will have a comparative advantage in sectors that use these intensively. This is what I call the "grey advantage."

The idea that demographics might affect comparative advantage is not new to the trade literature. Cai and Stoyanov (2016) provide the first systematic evidence, showing in a cross-section of bilateral trade flows that countries with older populations export relatively more in sectors intensive in age-appreciating skills. Gu and Stoyanov (2019) extend this by showing that aging erodes comparative advantage in sectors intensive in skill adaptability. Kopecky (2023) shows suggestive evidence in aggregate trade flows through a gravity framework, documenting that demographic similarity between trading partners predicts bilateral trade volumes. All three contributions are reduced-form estimates. This limits their ability to answer the quantitative questions that matter for policy, such as: how large are the welfare effects of demographic comparative advantage, how will bilateral trade between specific country pairs evolve, and which sectors are most exposed?

Answering these questions requires a calibrated structural model. I embed the demographic skill mechanism into the multi-sector Ricardian framework of Caliendo and Parro (2015), itself an extension of the foundational Eaton and Kortum (2002) model to multiple sectors with input–output linkages. The innovation relative to Caliendo and Parro (2015) is that I replace homogeneous labour with three heterogeneous skill types whose aggregate supplies are determined endogenously by the country's demographic structure. This allows me to embed the Cai and Stoyanov (2016) reduced-form relationship between age, ability and sectors into a fully specified general equilibrium model that accounts for input–output linkages, multilateral resistance, and the endogenous adjustment of wages and prices to demographic shocks.

The model is calibrated to 30 countries (the WIOD 2016 country sample, including a rest-of-world aggregate) and 20 manufacturing sectors matched to ISIC Revision 4. Skill intensities are constructed by combining O*NET ability importance scores with Bureau of Labor Statistics occupational composition data, following Cai and Stoyanov's (2016) methodology. Other trade parameters[1] come from the World Input–Output Database for the base year 2014. Age–productivity profiles are calibrated from the cognitive aging literature. Counterfactual exercises use the exact hat algebra of Dekle, Eaton, and Kortum (2008), which requires only observed trade shares and calibrated elasticities, not estimated levels of bilateral trade costs or absolute productivities.

The paper delivers four sets of results. First, I provide new reduced-form evidence on demographic comparative advantage using a comprehensive panel of 204 countries, 19 manufacturing sectors, and 30 years of trade data from BACI. The baseline specification regresses exporter–sector export shares on the interaction of sector skill intensity and country demographic structure, with exporter–sector and year fixed effects. The core finding is robust: a one-percentage-point increase in the share of workers aged 40–64

[1] Bilateral trade shares, input–output coefficients, value-added shares, and consumption shares.

increases a sector's export share by 0.55 percentage points for each unit of appreciating-skill intensity, and decreases it by 0.33 percentage points for each unit of depreciating-skill intensity. A Mundlak decomposition confirms that the effect operates both within and between countries, and long-difference estimates over the full 1995–2024 window remain significant. Figure 1 provides a visual preview. The 112,632 country-sector-year observations are sorted by the interaction of sector appreciating-skill intensity and country share aged 40–64 (after absorbing fixed effects), divided into 20 equal-sized bins, and the mean export share is plotted against the mean interaction within each bin. The positive slope confirms the grey advantage: country-sector-years where skill intensity and demographic aging are jointly high have larger export shares.

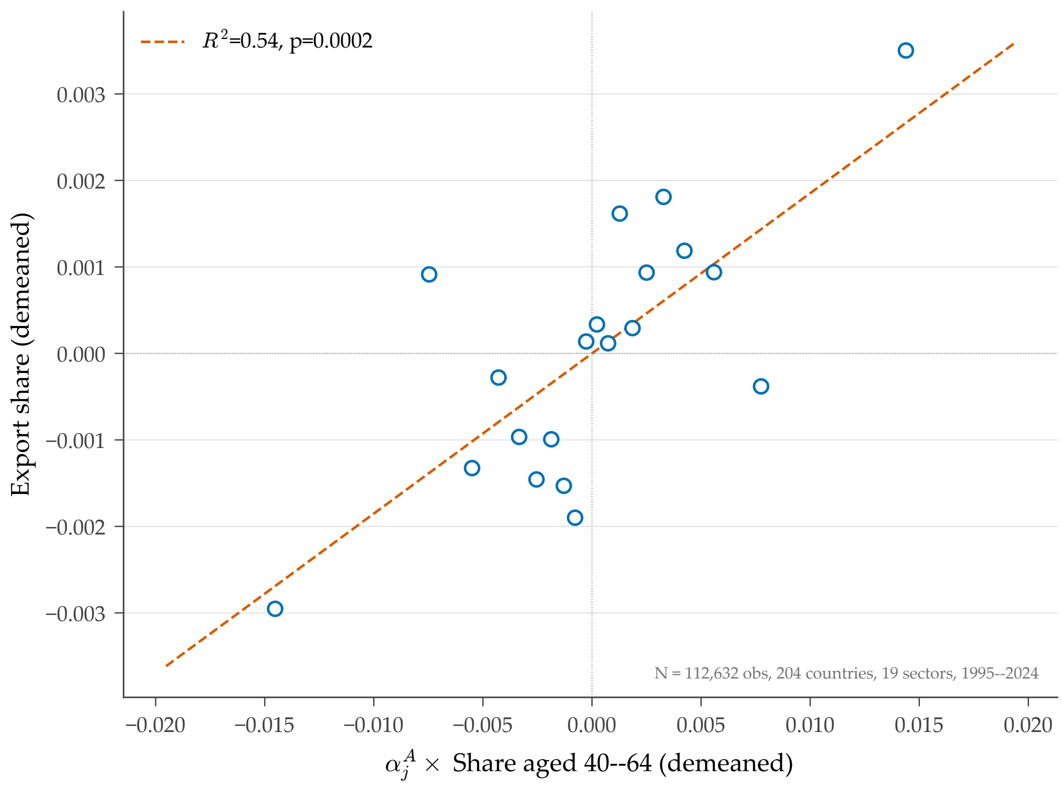

Figure 1: *Binscatter of sector export shares against the interaction of sector appreciating-skill intensity and country share aged 40–64, conditional on exporter–sector and year fixed effects. Each point is the mean of approximately 5,500 observations within an equal-sized bin. The dashed line is the OLS fit. Data: BACI (1995–2024), 204 countries, 19 manufacturing sectors, 112,632 observations.*

Second, a partial-equilibrium decomposition validates the full model's calibration. Evaluated on the same footing as the reduced-form regressions (holding wages and prices fixed), the model recovers 84% of the empirical Rybczynski coefficient. General equilibrium adjustment then compresses the trade-composition response by a factor of ten, as endogenous wage and price changes absorb the sectoral reallocation and generate the welfare effects that the reduced-form evidence alone cannot quantify. The skill-composition channel generates welfare effects ranging from -0.9% to $+0.05\%$ across countries, comparable in magnitude to Caliendo and Parro's (2015) estimate of NAFTA's effect on US welfare (0.11%). The total demographic effects are much larger (Japan -34%, India $+52\%$ by 2050), but these are dominated by the workforce-size channel (countries with shrinking working-age populations simply produce less). They complement estimates from the macroeconomic aging literature (Cooley, Henriksen, and Nusbaum, 2024; Kitao, 2015), which operate through different channels (capital accumulation, labour supply responses) but do not allow for trade adjustment. The total magnitudes isolate the demographic channel in a framework that holds technology and capital fixed and should be interpreted as partial-effect benchmarks accordingly. The composition effects exhibit a striking temporal pattern: counterfactual composition welfare was higher under earlier, more diverse demographic structures than under the 2014 baseline, but as countries' age structures converge toward a common older profile, the composition effects turn negative, suggesting that the gains from trade based on demographic comparative advantage are eroding.

Third, the model predicts a fundamental reorientation of bilateral trade flows away from rapidly aging economies that see larger shares of retirees and toward demographically young economies whose workforce is moving up this age-appreciation trajectory. By 2050 the China–US bilateral trade is predicted to fall by 10% while the India–US trade is expected to rise by 34% and Mexico–US trade by a similar 26%. At the sectoral level, the effects are heterogeneous in an economically meaningful way: China's exports of computers and electronics to the US are essentially unchanged (this sector's high appreciating-skill intensity insulates it from aging), while textiles, basic metals, and petroleum products decline by 17–67%. This sectoral pattern is qualitatively consistent with the ongoing "China-plus-one" narrative in global supply chains, suggesting that demographic forces may reinforce the economic incentives for supply-chain diversification alongside the geopolitical factors that have received more attention.

Fourth, I examine the role of retirement-age policy. Extending the retirement age by five years in aging OECD countries recovers substantial aggregate welfare: Japan gains 11.3 percentage points, Germany 8.7, and the United States 8.6. However, decomposing these gains reveals that they operate almost entirely through the workforce-size channel. The skill-composition effect is slightly negative. In other words, adding 65–69 year olds pushes aging economies' skill mix further from the global average, modestly worsening the comparative advantage distortion. Retirement extension may be an effective macroeconomic policy; it does not address the sectoral competitiveness challenge at the heart of this paper.

The paper bridges several literatures. It connects the quantitative trade literature (Eaton and Kortum, 2002; Caliendo and Parro, 2015; Costinot, Donaldson, and Komunjer, 2012; Costinot and Rodríguez-Clare, 2014) with the growing literature on the macroeconomic consequences of population aging (Bloom and Williamson, 1998; 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 the reduced-form evidence in Cai and Stoyanov (2016) and the gravity estimates in Kopecky (2023). It complements the work of Sposi (2022), who studies demographic effects on trade *imbalances* through aggregate saving channels but abstracts from the sectoral composition effects that are the focus here, and Brakman, Kohl, and van Marrewijk (2025), who use a gravity framework to project how demographic shifts will redistribute aggregate trade flows across regions but abstract from the within-country compositional channel through which the age structure of the workforce determines *which sectors* gain and lose. The skill-endowment mechanism relates to the broader literature on factor endowments and trade composition (Burstein and Vogel, 2017); Cai and Stoyanov (2023) show that progressive income taxation also shapes comparative advantage through the occupational channel, providing a policy-driven complement to the demographic forces studied here. The demand-side channels I abstract from are studied by Caron, Fally, and Markusen (2014). The China–US bilateral predictions connect to the large literature on trade with China (Autor, Dorn, and Hanson, 2013), offering a demographic channel that complements the geopolitical and policy factors emphasised in that literature. My work also relates to the literature on demographics and automation: Acemoglu and Restrepo (2022) show that aging countries adopt more robots, providing an alternative channel through which demographics reshape production patterns; the present paper isolates the skill-composition channel while holding technology fixed. Finally, the paper speaks to the dynamics of comparative advantage documented by Levchenko and Zhang (2016), offering a demographic explanation for part of the cross-country variation in how sectoral specialization evolves over time.

The remainder of the paper is organized as follows. Section 2 presents the model. Section 3 provides reduced-form evidence. Section 4 describes the calibration. Section 5 presents the quantitative results. Section 6 concludes.

2. Model

The economy consists of N countries indexed by $c \in \{1, \dots, N\}$ and J tradeable manufacturing sectors indexed by $j \in \{1, \dots, J\}$, plus a 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.

2.1. Demographics and Skill Endowments

Each country c at time t has a population distributed across A age groups, with $n_{c,a,t}$ denoting the mass of individuals in age group a.[2] There are K skill types indexed by k, partitioned into three categories:

- **Age-appreciating cognitive skills** ($k \in \mathcal{A}$): oral and written communication, comprehension
- **Age-depreciating cognitive skills** ($k \in \mathcal{D}$): memory, divided attention, perceptual speed, speed of closure
- **Physical abilities** ($k \in \mathcal{P}$): strength, stamina, coordination, dexterity

Following the neuroscience and gerontology literatures surveyed in Cai and Stoyanov (2016), these skill types evolve differently over the life cycle. For tractability, I aggregate the K individual skills into three composite skill types using principal component analysis, exactly as in Cai and Stoyanov (2016): age-appreciating cognitive ($k = A$), age-depreciating cognitive ($k = D$), and physical ($k = P$). In what follows I use $k \in \{A, D, P\}$.

Each worker in age group a is endowed with a bundle of effective skill units. The efficiency of an age-a worker in skill type k is given by:

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

where $\bar{e}^k > 0$ is a normalizing constant and $g^k : \{1, \ldots, A\} \to \mathbb{R}_+$ is the age–productivity profile for skill k. The profiles satisfy:

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

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

where $\underline{a}$ is the age group at which physical abilities peak. The profile $g^D(a)$ is hump-shaped, peaking in the late 20s to early 30s and declining thereafter. These shapes are calibrated from the cognitive aging literature (Section 4).

Workers supply their skill bundles inelastically. The aggregate effective supply of skill k in country c at time t is:

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

[2] In the calibration, I use 5-year age groups from 20–24 through 60–64, giving $A = 9$ working-age groups. Dependents (ages 0–19 and 65+) do not supply labour but affect demand.

Demographic change shifts the age distribution $\{n_{c,a,t}\}$ over time, which in turn alters the vector of effective skill 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 will have different effective skill endowments and therefore different comparative advantages.

2.2. Production

2.2.1 Tradeable Sectors

Within each tradeable sector $j \in \{1, \ldots, J\}$, there is a continuum of varieties $\omega \in [0, 1]$. Production of variety ω in sector j by country c requires a composite of the three skill types, combined with intermediate inputs.

The unit cost of producing any variety in sector j in country c is:

$$mc^j_{c,t} = \underbrace{\left[\frac{1}{\Phi^j_{c,t}} \prod_{k \in \{A,D,P\}} \left(w^k_{c,t} \right)^{\alpha^k_j} \right]^{\beta_j}}_{\text{value added}} \cdot \underbrace{\prod_{\ell=0}^{J} \left(P^\ell_{c,t} \right)^{\gamma^j_\ell}}_{\text{intermediates}}, \tag{5}$$

where $w^k_{c,t}$ is the price of skill k in country c, $\alpha^k_j \geq 0$ is the cost share of skill k in sector j's value added (with $\sum_k \alpha^k_j = 1$), $P^\ell_{c,t}$ is the price index of sector ℓ goods in country c, $\gamma^j_\ell \geq 0$ is the share of sector ℓ intermediates in sector j production, and $\Phi^j_{c,t}$ is a value-added-augmenting productivity term. Value-added and intermediates combine in Cobb–Douglas fashion with value-added share $\beta_j = 1 - \sum_\ell \gamma^j_\ell$.

The skill cost shares $\{\alpha^k_j\}$ are the key objects linking demographics to comparative advantage. They are constructed from O*NET skill importance scores and Bureau of Labor Statistics occupational composition data for each 4-digit NAICS manufacturing industry, following the methodology of Cai and Stoyanov (2016). Sectors with high α^D_j (e.g., yarn mills, wood products—dominated by machine operators for whom coordination and perceptual speed are critical) have comparative advantage in demographically young countries. Sectors with high α^A_j (e.g., printing, beverages—employing many technical writers and sales representatives) have comparative advantage in demographically old countries.

2.2.2 Non-Tradeable Sector

The non-tradeable sector ($j = 0$) uses the same Cobb–Douglas technology over skills and intermediates, with its own skill cost shares $\{\alpha^k_0\}$ and input–output coefficients $\{\gamma^0_\ell\}$. These are calibrated from WIOD in the same way as the tradeable sectors (Section 4.2). The non-tradeable sector enters household consumption (Section 2.4) and serves as an intermediate input for tradeable sectors, but is not traded internationally: $\pi^0_{cp} = 1$ if $c = p$ and zero otherwise.

2.2.3 Ricardian Productivity

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

$$\Pr\left[z^j_c(\omega) \leq z\right] = \exp\left(-T^j_{c,t} \cdot z^{-\theta_j}\right), \tag{6}$$

where $T^j_{c,t} > 0$ governs the location (absolute advantage) of country c in sector j and $\theta_j > 1$ governs the dispersion (scope for comparative advantage). I allow $T^j_{c,t}$ to depend on demographics through:

$$\ln T^j_{c,t} = \mu_c + \mu_j + \sum_{k \in \{A,D,P\}} \rho_k \cdot \alpha^k_j \cdot \ln L^k_{c,t} + \nu^j_{c,t}, \tag{7}$$

where μ_c and μ_j are country and sector fixed effects, ρ_k captures the Ricardian channel, which is the extent to which a country's effective skill endowment enhances productivity in skill-k-intensive sectors, and $\nu^j_{c,t}$ is a residual.

Proposition 1 (Demographic Comparative Advantage)**.** *Consider two countries c_1 and c_2 with identical populations, trade costs, and residual productivities, but where c_1 has a younger age distribution. If $\rho_k > 0$ for all k and the age–productivity profiles satisfy (2)–(3), then:*

1. *Country c_1 has lower relative unit costs in sectors with high α^D_j or α^P_j (age-depreciating skills);*
2. *Country c_2 has lower relative unit costs in sectors with high α^A_j (age-appreciating skills).*

The proof follows from noting that a younger age distribution implies $L^D_{c_1,t} > L^D_{c_2,t}$ and $L^P_{c_1,t} > L^P_{c_2,t}$ (more workers in age groups where depreciating skills and physical abilities are high), while $L^A_{c_1,t} < L^A_{c_2,t}$ (fewer workers in age groups where appreciating skills peak). Through both the Heckscher–Ohlin channel (lower skill prices due to higher endowments) and the Ricardian channel (higher $T^j_{c,t}$ via equation (7)), country c_1 gains comparative advantage in depreciating-skill-intensive sectors.

2.3. Trade

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

$$\pi^j_{cp,t} = \frac{T^j_{c,t} \cdot \left(mc^j_{c,t} \cdot d^j_{cp} \right)^{-\theta_j}}{\Omega^j_{p,t}}, \tag{8}$$

where $d^j_{cp} \geq 1$ is the iceberg trade cost for shipping sector j goods from c to p (with $d^j_{cc} = 1$), and

$$\Omega^j_{p,t} = \sum_{c'=1}^{N} T^j_{c',t} \cdot \left(mc^j_{c',t} \cdot d^j_{c'p} \right)^{-\theta_j} \tag{9}$$

is a multilateral resistance term. The sectoral price index in country p is:

$$P^j_{p,t} = \zeta_j \cdot \left(\Omega^j_{p,t} \right)^{-1/\theta_j}, \tag{10}$$

where $\zeta_j = \Gamma\left(\frac{\theta_j + 1 - \sigma_j}{\theta_j} \right)^{1/(1-\sigma_j)}$ is a constant involving the within-sector elasticity of substitution $\sigma_j > 1$ and the Gamma function.

Bilateral exports from c to p in sector j are:

$$X^j_{cp,t} = \pi^j_{cp,t} \cdot E^j_{p,t}, \tag{11}$$

where $E^j_{p,t}$ is total expenditure of country p on sector j goods.

2.4. Demand

Final demand in country c at time t is generated by a representative household with Cobb–Douglas preferences over goods from each sector $j \in \{0, 1, \ldots, J\}$, with expenditure shares $\{\eta^j_c\}$ satisfying $\sum_{j=0}^{J} \eta^j_c = 1$:

$$U_{c,t} = \prod_{j=0}^{J} \left(C^j_{c,t} \right)^{\eta^j_c}, \tag{12}$$

where $C^j_{c,t}$ is the composite consumption of sector j goods. Final demand for sector j is $F^j_{c,t} = \eta^j_c \cdot Y_{c,t}$, where $Y_{c,t}$ is total income. Total expenditure on sector j goods in country c

includes both final demand and intermediate demand:

$$E^j_{c,t} = F^j_{c,t} + \sum_{\ell=0}^{J} \gamma^\ell_j \cdot Y^\ell_{c,t}, \tag{13}$$

where γ^ℓ_j is the share of sector j intermediates in sector ℓ's production. It is this total expenditure $E^j_{p,t}$ that enters the gravity equation (11). Under this specification, demographics affect trade only through the supply side.[3]

2.5. Equilibrium

Given demographics $\{n_{c,a,t}\}$, trade costs $\{d^j_{cp}\}$, and exogenous productivities $\{\nu^j_{c,t}\}$, a competitive equilibrium at time t consists of skill prices $\{w^k_{c,t}\}$, sectoral price indices $\{P^j_{c,t}\}$, trade shares $\{\pi^j_{cp,t}\}$, and expenditures $\{E^j_{c,t}\}$ such that:

1. **Goods market clearing:** For each country c and sector j,

$$\sum_{p=1}^{N} X^j_{cp,t} = Y^j_{c,t}, \tag{14}$$

   where $Y^j_{c,t}$ is total revenue of sector j in country c.

2. **Skill market clearing:** For each country c and skill k,

$$L^k_{c,t} = \sum_{j=0}^{J} \frac{\alpha^k_j \cdot \beta_j \cdot Y^j_{c,t}}{w^k_{c,t}}. \tag{15}$$

3. **Trade balance:**[4]

$$\sum_{j=1}^{J} \sum_{p=1}^{N} X^j_{cp,t} = \sum_{j=1}^{J} \sum_{p=1}^{N} X^j_{pc,t} \quad \forall\, c. \tag{16}$$

4. **Income identity:** Total income in country c equals total factor payments plus tariff revenue (if any):

$$Y_{c,t} = \sum_k w^k_{c,t} \cdot L^k_{c,t}. \tag{17}$$

[3] An extension would allow expenditure shares to depend on the age distribution, capturing the empirical regularity that consumption baskets shift toward non-tradeable services as populations age. Calibrating such an extension would require age-specific expenditure data at the sectoral level across a broad set of countries, which is not currently available. I discuss this as a direction for future work in Section 6.

[4] In the baseline, I impose balanced trade. The extension to allow imbalances using the methodology of Sposi (2022) is straightforward and would allow the model to also address demographic effects on aggregate capital flows.

2.6. Counterfactual Analysis: Exact Hat Algebra

For counterfactual exercises, I employ the exact hat algebra approach of Dekle et al. (2008). Define $\hat{x} = x'/x$ as the ratio of a variable in the counterfactual to its value in the observed equilibrium. Given a change in demographics from $\{n_{c,a,t}\}$ to $\{n'_{c,a,t'}\}$, the implied changes in skill endowments are:

$$\hat{L}_c^k = \frac{\sum_a n'_{c,a,t'} \cdot e^k(a)}{\sum_a n_{c,a,t} \cdot e^k(a)}. \tag{18}$$

The counterfactual equilibrium is characterized by a system in the hat variables $\{\hat{w}_c^k, \hat{P}_c^j, \hat{\pi}_{cp}^j\}$ that can be solved using only observed trade shares $\{\pi_{cp,t}^j\}$, calibrated parameters $\{\alpha_j^k, \beta_j, \gamma_\ell^j, \theta_j, \eta_c^j\}$, and the demographic-induced changes $\{\hat{L}_c^k\}$, without requiring estimation of the levels $T_{c,t}^j$ or d_{cp}^j.

The system of equations in hat notation is:

$$\widehat{mc}_c^j = \prod_k \left(\hat{w}_c^k\right)^{\alpha_j^k \cdot \beta_j} \cdot \prod_\ell \left(\hat{P}_c^\ell\right)^{\gamma_\ell^j}, \tag{19}$$

$$\hat{\pi}_{cp}^j = \frac{\hat{T}_c^j \cdot \left(\widehat{mc}_c^j\right)^{-\theta_j}}{\sum_{c'} \pi_{c'p}^j \cdot \hat{T}_{c'}^j \cdot \left(\widehat{mc}_{c'}^j\right)^{-\theta_j}}, \tag{20}$$

$$\hat{P}_p^j = \left[\sum_{c'} \pi_{c'p}^j \cdot \hat{T}_{c'}^j \cdot \left(\widehat{mc}_{c'}^j\right)^{-\theta_j}\right]^{-1/\theta_j}, \tag{21}$$

where the change in technology due to demographics is:

$$\ln \hat{T}_c^j = \sum_k \rho_k \cdot \alpha_j^k \cdot \ln \hat{L}_c^k. \tag{22}$$

The value-added productivity parameter $\Phi_{c,t}^j$ from equation (5) is held fixed in all counterfactual exercises ($\hat{\Phi}_c^j = 1$); demographic shocks operate on sectoral productivity solely through $\hat{T}_c^j$ via equation (22). The goods and factor market clearing conditions close the system. This formulation makes the counterfactual exercises computationally tractable: for each projected future demographic structure, I compute $\hat{L}_c^k$ from equation (18) and solve the nonlinear system (19)–(22) for counterfactual wages, prices, and trade shares.

A note on dynamics. The model computes a sequence of static equilibria rather than a fully dynamic transition path. Caliendo, Dvorkin, and Parro (2019) show that in the Caliendo and Parro (2015) framework, dynamic labour reallocation frictions attenuate shortrun responses but leave long-run welfare effects largely unchanged. The static approach is appropriate here because demographic change is slow relative to trade adjustment. Caliendo et al.'s (2019) estimates imply adjustment half-lives of roughly ten years, while the counterfactual projections look 35 to 85 years ahead. Because the demographic channel of interest operates through the exogenous composition of the labour force, not through workers' endogenous sectoral choices, dynamics are likely less of a concern. The key dynamic channel the model misses is capital accumulation: aging affects saving rates, which affects the capital stock, which independently shifts comparative advantage (Sposi, 2022; Cooley et al., 2024). My aggregate welfare declines should be taken with this important caveat. The trade composition results (my main focus) are more robust to this omission: to the extent that aging economies accumulate capital, and capital-intensive sectors overlap with appreciating-skill-intensive sectors, capital deepening would reinforce rather than undermine the grey advantage identified here.

3. Reduced-Form Evidence

Before turning to the structural model estimation, I present reduced-form evidence that demographic age structure shapes the pattern of sectoral comparative advantage, as predicted by Proposition 1. This evidence serves two purposes. First, it updates and extends the original findings of Cai and Stoyanov (2016), who used a cross-section of bilateral trade flows, to a panel setting with 30 years of within-country demographic variation. Second, the estimated coefficients provide an external benchmark against which to validate the calibrated model's predictions.

3.1. Data

The reduced-form analysis combines several data sources. Trade data come from BACI (CEPII), which harmonizes UN Comtrade bilateral trade flows at the HS 6-digit level. I aggregate to 19 manufacturing sectors defined at the 2-digit ISIC Revision 4 level and construct exporter–sector export shares as the ratio of country c's exports in sector j to its total manufacturing exports.[5] The panel covers 1995–2024, with 204 exporting countries observed in each year, yielding a total of 112,632 exporter–sector–year observations.

Demographic data come from the United Nations World Population Prospects 2024 Revision (United Nations, Department of Economic and Social Affairs, Population Division, 2024), which provides age-specific population counts for 237 countries from 1950 to 2100. I

[5] I exclude ISIC 12 (Tobacco), which is heavily regulated and has limited cross-country variation in trade patterns.

construct the share of the working-age population in each 10-year age bin (20–29, 30–39, 40–49, 50–64), the share aged 40–64, and the old-age dependency ratio (65+ / 15–64). All demographic variables are measured at the start of the period to reduce simultaneity concerns.

Sector skill intensities α_j^k are the same O*NET-based measures used in the structural model (Section 4). The robustness specification in column 7 additionally uses sector-level capital shares from the WIOD Socio-Economic Accounts and country–year capital–labour ratios from the Penn World Table 10.01 (175 countries) to control for the standard Heckscher–Ohlin capital channel. The key variable of interest is the interaction $\alpha_j^k \times \text{DemoVar}_{ct}$: the prediction from the model is that sectors with high appreciating-skill intensity (α_j^A) should see rising export shares in countries whose working-age populations are shifting toward older age groups.

3.2. Specification

The baseline specification is:

$$s_{ct}^j = \sum_k \phi_k \cdot \alpha_j^k \cdot \text{DemoVar}_{ct} + \mu_{cj} + \delta_t + \varepsilon_{ct}^j, \tag{23}$$

where s_{ct}^j is the export share of sector j in country c's total manufacturing exports in year t, α_j^k is the intensity of skill k in sector j, DemoVar_{ct} is the demographic variable (share aged 40–64, or age-bin shares), μ_{cj} is an exporter–sector fixed effect, and δ_t is a year fixed effect. Standard errors are clustered at the exporter level to account for arbitrary serial correlation within countries.

The exporter–sector fixed effects absorb all time-invariant determinants of country c's specialization in sector j, including institutional quality, geography, endowments of physical capital, and historical patterns of industrial policy. Identification therefore comes from within-country, within-sector variation in demographic structure over the 30-year panel. The year fixed effects absorb global trends in sectoral trade patterns, including technology-driven shifts in comparative advantage that are common across countries.

3.3. Results

Table 1 reports the main results across nine specifications. In the baseline (column 1), the coefficient on $\alpha_j^A \times \text{share}_{40\text{–}64}$ is 0.553, indicating that an increase in the share of workers aged 40–64 is associated with a statistically significant increase in export shares for sectors intensive in appreciating skills. The coefficient on $\alpha_j^D \times \text{share}_{40\text{–}64}$ is -0.327, consistent with the symmetric prediction: countries with older workforces have lower export shares in sectors intensive in depreciating skills.

Columns 8–9 separate the working-age population into young (20–39) and experienced (40–64) groups, interacting both with α_j^A simultaneously. Both coefficients are positive and highly significant in the full sample (column 8): 0.305 for the young share and 0.457 for the experienced share. The larger coefficient on the experienced group is consistent with the appreciating-skill mechanism. The positive coefficient on the young group reflects that young workers do supply appreciating skills (they simply supply less than older workers), and that both groups contribute to overall manufacturing skill endowments relative to the dependent population. The key comparison is the magnitude: the experienced-share coefficient (0.457) exceeds the young-share coefficient (0.305), indicating that older workers are associated with a stronger shift toward appreciating-skill-intensive sectors.

Columns 3–4 implement a Mundlak (1978) decomposition, separating the within-country (time-demeaned) from the between-country (country-mean) variation. The within-country coefficient is 0.567, while the between-country coefficient is substantially larger at 1.145. Both are precisely estimated and positive, indicating that the association holds both within countries over time and across countries in the cross section.

Column 5 reports a long-difference specification, regressing the 1995–2024 change in export shares on the 1995–2024 change in demographics. The coefficient is 0.761, confirming that the relationship is robust to collapsing the panel to a single long difference that maximizes the signal-to-noise ratio of demographic change.

Columns 6–7 report two demanding robustness checks. First, replacing year fixed effects with exporter–year fixed effects (column 6) absorbs all time-varying country-level confounders, including GDP growth, exchange rate movements, trade policy changes, and capital accumulation. Identification comes purely from the *cross-sector* differential: is the cross-sector gradient in export shares with respect to α_j^A steeper in country-years with older working-age populations? The coefficient is essentially unchanged at 0.554, indicating that the grey advantage is not driven by country-year confounders. Second, column 7 tests whether demographics operate through the capital channel rather than the skill channel by adding the interaction of sector capital share (from WIOD) with country capital–labour ratio (from Penn World Table 10.01) alongside the skill interactions, for 175 countries. The capital interaction is insignificant, while the skill channel remains significant at 0.433. The demographic-skill interaction remains the dominant correlate of export composition, with no evidence that the capital channel is driving the results.

Table 1: *Demographic 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)				
$\text{share}_{20\text{–}39} \times \alpha_j^A$								0.305***	0.481***
								(0.075)	(0.135)
$\text{share}_{40\text{–}64} \times \alpha_j^A$								0.457***	0.555***
								(0.097)	(0.180)
Fixed effects	Exp–sec, Year	Exp–sec, Year	Exp–sec, Year	Exp–sec, 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	3,876	112,632	80,187	112,632	21,508

Notes: Dependent variable is sector export share (levels). Standard errors clustered by exporter in parentheses. $^{***}\, p < 0.01$, $^{**}\, p < 0.05$, $^{*}\, p < 0.10$. Columns (1)–(7): the demographic variable is $\text{share}_{40\text{–}64}$ (or mean age in column 2), interacted with sector skill intensities α_j^A and α_j^D. (3)–(4): Mundlak decomposition into within-country and between-country variation. (5): long difference, 1995–2024 change. (6): exporter–year FE absorb all country–year variation. (7): adds sector capital share × country K/L ratio (WIOD SEA; PWT 10.01). In (6)–(7), α_j^P is omitted ($\alpha_j^A + \alpha_j^D + \alpha_j^P = 1$). Columns (8)–(9): both age-group shares interacted with α_j^A simultaneously; (9) restricts to 38 OECD countries. Sample: BACI, 1995–2024.

3.4. Identification and Interpretation

Several features of the research design address potential threats to identification. The exporter–sector fixed effects are demanding: they absorb not only time-invariant country characteristics but also time-invariant country–sector interactions, such as a country's historical specialization in particular industries. The identifying variation is therefore the within-country demographic change over the 30-year panel. Because demographic transitions are slow-moving and driven by fertility decisions made decades earlier, they are largely predetermined with respect to contemporaneous trade shocks.

One remaining concern is reverse causality: perhaps sectors experiencing export growth attract workers of particular ages, generating a spurious correlation. Do, Levchenko, and Raddatz (2016) document this channel, showing that trade openness affects fertility decisions. This concern is mitigated by two features of the specification. First, the demographic variables are measured at the country level. Sector-level labour demand shocks are unlikely to meaningfully alter a country's aggregate age distribution. Second, the long-difference specification in column 5 uses demographic changes over the full 30-year window, which are dominated by cohort-size effects from past fertility, not by labour market conditions. A related concern is international migration, which responds to economic conditions and directly affects age structure. This channel is most relevant for countries with large immigration flows (e.g., Gulf states, Singapore) but contributes modestly to national-level age structure change for most countries in the sample.

The magnitudes are economically meaningful. The baseline coefficient implies that a country undergoing the demographic transition from the 25th to the 75th percentile of the share aged 40–64 would see its export shares in the most appreciating-skill-intensive sector increase by approximately 1.2 percentage points relative to the least intensive sector. This is a meaningful reallocation given that the median sector export share is approximately 5%.

These reduced-form results are consistent with the core prediction of the model and provide the quantitative benchmark against which Section 4.5 evaluates the structural model.

4. Calibration

The model is calibrated to the base year 2014, the last year of the WIOD 2016 release. I describe each component in turn.

4.1. Skill Intensities

The skill intensity matrix $\{\alpha_j^k\}$ maps each of the 20 manufacturing sectors to its cost shares in the three composite skills. Following Cai and Stoyanov (2016), I construct these from two data sources. First, O*NET version 29.1 provides importance ratings for 18 cognitive and physical abilities across 763 occupations defined at the SOC 2018 8-digit level.[6] Second, the Bureau of Labor Statistics Occupational Employment and Wage Statistics (OEWS) for May 2014 provides employment counts by detailed occupation for each 4-digit NAICS manufacturing industry, allowing me to aggregate occupation-level skill scores to the industry level using employment weights.

The 18 O*NET abilities are reduced to three composite skills via principal component analysis. Before applying PCA, I normalize each ability by its loading on Inductive Reasoning, following Cai and Stoyanov (2016), to ensure that the first three principal components correspond interpretably to the three skill categories. The first component loads positively on oral and written comprehension, expression, and speech clarity—the appreciating skills. The second loads on memory, perceptual speed, speed of closure, and selective attention—the depreciating cognitive skills. The third loads on trunk strength, stamina, static strength, and manual dexterity—the physical abilities.

[6] I use the "Importance" ratings (Scale ID = IM) rather than "Level" ratings, following Cai and Stoyanov's (2016) methodology.

For each 4-digit NAICS industry, I compute the employment-weighted average score on each of the three composites across all occupations employed in that industry. These are then aggregated to the 20 ISIC Revision 4 manufacturing sectors using a NAICS-to-ISIC concordance. The resulting sector-level scores are normalized to sum to unity within each sector, yielding $\alpha_j^A + \alpha_j^D + \alpha_j^P = 1$ for each j. These α_j^k should be interpreted as task-intensity indices rather than structural cost shares or production elasticities; the same caveat applies to Cai and Stoyanov's (2016) original construction. A further maintained assumption is that the skill-intensity structure of each sector is common across countries. In practice, the same ISIC sector may employ different occupational mixes in different countries, reflecting differences in factor prices and technology. To the extent that developing-country manufacturing relies more heavily on physical tasks than its US counterpart, the model may understate cross-country differences in skill intensity. Consistent with this, the reduced-form coefficients are larger when the sample is restricted to OECD countries (Table 1, column 9), where the US-derived skill intensities are a closer match. Two features mitigate these measurement concerns: the reduced-form results in Section 3 validate the mechanism independently of the structural calibration, and the partial-equilibrium decomposition in Section 4.5 confirms that the calibrated model recovers 84% of the empirical Rybczynski coefficient.

The resulting α matrix reveals substantial cross-sector variation. Computers and electronics (ISIC 26) has the highest appreciating-skill intensity ($\alpha^A = 0.586$), reflecting its heavy reliance on engineers, technical writers, and managers whose communication and comprehension skills improve with experience. At the other extreme, wood products (ISIC 16) and textiles (ISIC 13) have the highest depreciating-skill and physical intensities, reflecting their reliance on machine operators and manual labourers whose speed and dexterity peak early in the working life.

4.2. 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 World Input–Output Database (WIOD) 2016 release for the year 2014. WIOD provides inter-country input–output tables for 43 countries and 56 sectors; I aggregate to the 30-country, 20-sector classification used in the model, collapsing the remaining 13 countries into a rest-of-world composite.

The trade elasticities $\{\theta_j\}$ are taken from Caliendo and Parro (2015), who estimate sector-specific elasticities using tariff variation. Following their approach, I impose a lower bound of $\theta_j \geq 2$ for sectors where the point estimate is imprecisely estimated or implausibly low, as trade elasticities below 2 generate unrealistically large trade responses to small cost changes.

4.3. Age–Productivity Profiles

The age–productivity profiles $\{g^k(a)\}$ govern how each skill type evolves over the working life. I calibrate these from three sources in the cognitive aging and occupational health literatures.

For appreciating cognitive skills, Schaie (2005) documents in the Seattle Longitudinal Study that verbal ability, verbal memory, and inductive reasoning show minimal decline through age 60, with vocabulary continuing to increase through the mid-60s. I fit a log-linear profile that increases at roughly 0.5% per year from age 20 to 55, then flattens.

For depreciating cognitive skills, Salthouse (2009) shows that processing speed, working memory, and spatial visualization begin declining from the late 20s at a rate of roughly 0.5–1% per year, with acceleration after age 50. I fit a quadratic profile peaking at age 27 and declining monotonically thereafter.

For physical abilities, **?** and the occupational health literature document that grip strength, trunk strength, and aerobic capacity peak in the mid-to-late 20s and decline at roughly 1% per year thereafter. I fit a profile peaking at age 25 and declining linearly.

A maintained assumption is that the age–productivity profiles are common across countries. Lagakos, Moll, Porzio, Qian, and Schoellman (2018) show that experience–wage profiles are substantially steeper in rich countries, suggesting that the appreciating-skill gradient may be stronger in advanced economies. To the extent that this is the case, the uniform profiles used here may understate the grey advantage for OECD countries and overstate it for developing economies.

4.4. Demographics

Demographic data for the 30 model countries come from the United Nations World Population Prospects 2024 Revision (United Nations, Department of Economic and Social Affairs, Population Division, 2024), which provides population counts by 5-year age group for 237 countries from 1950 to 2100 (medium variant projections). For each country and year, I compute effective skill supplies $L^k_{c,t}$ by applying the calibrated age–productivity profiles to the observed (1950–2024) or projected (2025–2100) age distribution. Historical counterfactuals use demographic data from 1970, 1980, 1990, 2000, and 2010; forward projections use 2020, 2030, 2040, 2050, and 2060.

4.5. Reconciling the Empirical and Structural Estimates

A key question is whether the apparent gap between the empirical and structural Rybczynski coefficients reflects model misspecification or the expected consequence of general equilibrium adjustment. The empirical coefficient from Table 1 is 0.553; the full-GE structural coefficient is 0.048, roughly one-tenth as large. I show that this gap is almost entirely attributable to general equilibrium dampening by solving the model under progressively stronger equilibrium restrictions.

Table 2 reports the results of this exercise. I solve the historical decomposition (1970–2014) at three levels. At *Level 1* (pure Ricardian PE), only sectoral productivities $\hat{T}_c^j$ change through equation (22); wages and prices are held fixed. This isolates the Ricardian channel and generates a coefficient of just 0.006, indicating that the productivity channel alone has negligible quantitative bite. At *Level 2* (Heckscher–Ohlin PE), wages additionally adjust to the demographic shock at fixed world prices via the factor-price-insensitivity result ($\hat{w}_c^k = 1/\hat{L}_c^k$), while intermediate-good prices and multilateral resistance remain fixed. This captures the standard HO mechanism: an increase in the supply of skill k lowers its price, reducing unit costs in k-intensive sectors. The Level 2 coefficient is 0.464, capturing 84% of the empirical estimate. At *Level 3* (full GE), the model solves for the complete equilibrium adjustment of wages, intermediate prices, price indices, expenditures, and multilateral resistance. The coefficient compresses to 0.048.

The progression from Level 2 to Level 3 reveals that GE dampening compresses the trade-composition response by a factor of approximately ten. This dampening has a clear economic interpretation. When a country accumulates appreciating skills, the HO channel lowers their price, conferring comparative advantage in appreciating-skill-intensive sectors. But as the country specialises in those sectors, the increased demand for appreciating skills pushes wages back up, while the expansion of supply lowers the world price of those goods through multilateral resistance. These endogenous responses absorb the trade-composition shift but generate the real-income changes that constitute the welfare effects reported in Section 5.1.

The Level 2 coefficient of 0.464 provides the appropriate benchmark against which to evaluate the empirical estimate of 0.553. The exporter-sector and year fixed effects in the empirical specification absorb country-year wage and price adjustments, making the empirical estimate a partial-equilibrium object that is comparable to the Level 2 exercise. The remaining gap (84% rather than 100%) is consistent with differences in sample composition (30 model countries vs. 204 in the empirical panel), specification (changes vs. levels), and the fact that the fixed effects do not absorb all GE channels. The model's calibrated skill intensities and age-productivity profiles thus generate demographic comparative advantage effects of the right order of magnitude when evaluated on a comparable footing.

The decomposition also clarifies the role of the Ricardian channel parameter ρ. The near-zero Level 1 coefficient indicates that the Ricardian channel (demographics $\to T^j_{c,t}$) contributes minimally to the PE trade-composition response; the quantitative heavy lifting is done by the HO factor-price channel. This insulates the core results from the fact that ρ is not independently estimated: even at $\rho = 0$, the HO channel alone would generate most of the trade-composition response. Table 3 confirms this from a different angle: as ρ varies across a wide grid, the composition-only welfare effects change gradually, and the qualitative results are preserved throughout.

Table 2: *Partial-Equilibrium Decomposition of the Rybczynski Gap*

	α^A_j coef.	t	α^D_j coef.	t
Model (historical decomposition, 29 countries $\times$ 20 sectors $\times$ 5 decades)				
Level 1: Ricardian PE ($\hat{T}$ only)	0.006	0.62	-0.003	-0.60
Level 2: HO + Ricardian PE ($\hat{w} = 1/\hat{L}$)	0.464	2.06	-0.296	-2.08
Level 3: Full GE	0.048	3.50	-0.030	-3.07
Empirical (Table 1)				
Column 1 (204 countries)	0.553	3.68	-0.327	-3.23
Level 2 / Empirical	0.839		0.905	
Level 3 / Empirical	0.087		0.091	

Notes: Rybczynski coefficients from regressing model-implied export share changes on $\alpha^k_j \times \Delta\text{share}_{40\text{–}64}$ with sector and decade fixed effects, clustering by country. Levels 1–3 are defined in the text.

5. Quantitative Results

I now use the calibrated model to quantify the effects of demographic change on trade patterns and welfare. I begin by decomposing the welfare effects of projected demographic change into a workforce-size channel and a skill-composition channel, then examine how demographic divergence and convergence shape these composition effects over time. I then turn to bilateral trade reallocation, a sectoral decomposition, and a retirement-age policy counterfactual.

5.1. Welfare Decomposition: Level vs. Composition

The central exercise uses United Nations population projections (medium variant) to compute counterfactual equilibria at each decade from 2020 to 2060. For each projection

Table 3: *Rybczynski Coefficient and Welfare Decomposition: Sensitivity to ρ*

ρ	Ryb. coef.	t	Composition-only welfare (%) Japan	China	Korea
0.10	0.068	2.10	−0.06	−0.69	−0.39
0.25	0.060	2.83	−0.07	−0.78	−0.44
0.50	0.048	3.50	−0.07	−0.92	−0.53
0.75	0.037	1.30	−0.07	−1.06	−0.62
1.00	0.026	0.55	−0.08	−1.20	−0.70
2.00	0.006	0.06	−0.09	−1.75	−1.05
Empirical coefficient					
Fullsample	0.553	3.68			

Notes: Each row re-solves the model at the indicated ρ (symmetric across skill types). Rybczynski coefficient estimated from the historical decomposition (29 countries × 20 sectors × 5 decades). Composition-only welfare isolates the skill-mix shift by removing the workforce-size change (see text). Baseline ($\rho = 0.5$).

year, I compute the change in effective skill endowments $\hat{L}_c^k$ implied by the projected demographic structure and solve the hat-algebra system for the new equilibrium.

The headline welfare effects combine two distinct channels. The first is an aggregate *workforce-size* channel: countries with shrinking working-age populations have fewer workers and therefore produce less. The second is the *skill-composition* channel that is the focus of this paper: aging shifts the mix of effective skills across workers, altering sectoral comparative advantage. To separate the two, I decompose each country's demographic shock $\hat{L}_c^k$ into a level component (the geometric mean across skill types, representing the uniform workforce size change) and a composition component (the residual, representing the pure skill-mix shift), and solve the model under each component separately.

The composition-only welfare effects, ranging from -0.9% to $+0.05\%$ across countries, are the model's estimate of the pure comparative advantage channel of demographic change. These magnitudes are comparable to the welfare effects of major trade agreements: Caliendo and Parro's (2015) estimate of NAFTA on US welfare is 0.11%. Japan's composition-only effect of -0.07% is of the same order, and China's -0.9% is roughly eight times as large. The skill-composition channel is economically meaningful even though it is a small fraction of the total demographic effect.

The aggregate workforce-size channel accounts for nearly all of the headline welfare effects. Japan's 34% welfare loss decomposes into 34.1% from workforce shrinkage and just 0.07% from skill-composition shifts.[7] Korea's 40% loss is 39.8% level and 0.5% composition.

[7] The level–composition decomposition is multiplicative: $\hat{W}_c = \hat{W}_c^{\text{level}} \times \hat{W}_c^{\text{comp}}$. The reported percentages

For India, the 52% gain is 53.4% level and -0.9% composition; the composition channel actually works against India because its workforce is aging slightly even as it grows. Figure 2 plots the total welfare trajectories for selected countries across both the historical and forward horizons.

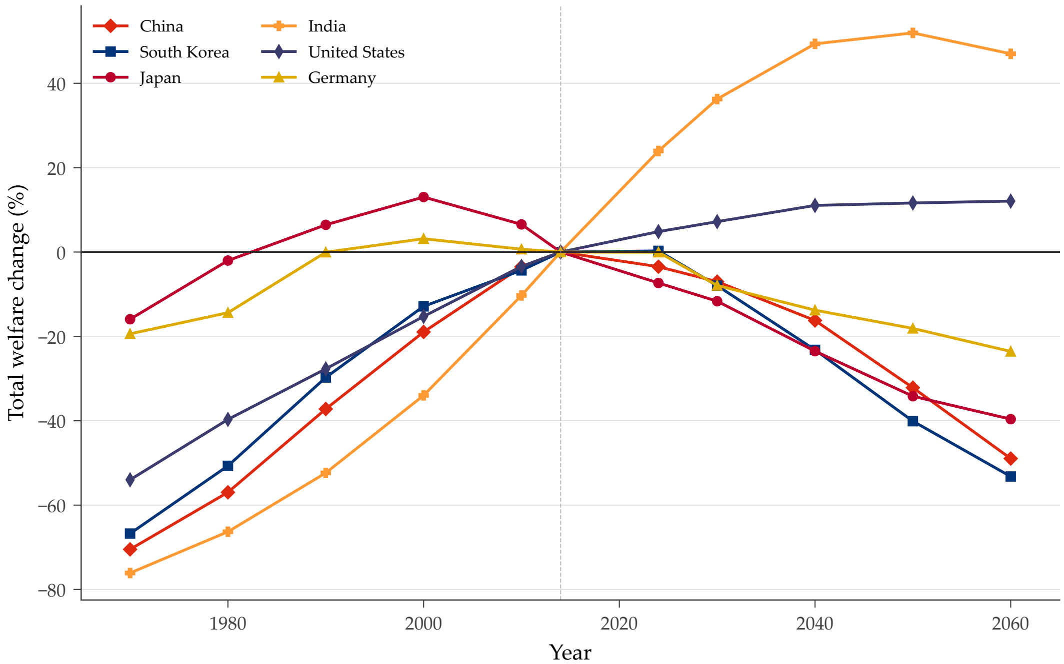

Figure 2: *Total demographic welfare, 1970–2060. Each line traces the total welfare effect of replacing 2014 demographics with the demographic structure of the indicated year, holding the 2014 trade structure fixed. Points left of the dashed line are backward-looking counterfactuals; points to the right are forward projections (UN WPP 2024, medium variant).*

The total welfare effects fall into three broad groups.

Rapidly aging economies. Japan, Korea, Taiwan, Italy, and Germany face large welfare losses from continued aging. By 2050, Japan's welfare falls by 34%, Korea's by 40%, Taiwan's by 44%, Italy's by 34%, and Germany's by 18%. These losses reflect the erosion of these countries' effective skill endowments as their working-age populations shrink and shift toward older age groups. The losses are largest for Korea and Taiwan, which are experiencing the world's most rapid fertility declines and are projected to have median ages exceeding 55 by mid-century.

Transitioning economies. China occupies an intermediate position: its welfare declines by 32% by 2050, driven by the legacy of the one-child policy, which will cause its working-age

therefore do not sum exactly.

population to shrink by nearly 200 million between 2014 and 2060. The United States, by contrast, gains 12% by 2050, reflecting its relatively favorable demographics, a combination of higher fertility and immigration that keeps its working-age population growing, albeit slowly. France loses a modest 3%, consistent with its demographic profile being close to replacement fertility.

Young economies. India (+52%), Mexico (+35%), and Indonesia (+30%) are the major beneficiaries. These countries are in the early-to-middle stages of their demographic transitions, with large cohorts of young workers entering the labour force, the classic "demographic dividend" (Bloom, Canning, and Sevilla, 2003). Their effective supplies of depreciating cognitive and physical skills are growing rapidly, conferring comparative advantage in precisely the labour-intensive manufacturing sectors that have historically driven export-led development.

Brazil and Turkey represent an instructive intermediate case. Both experience welfare gains through approximately 2040 (Brazil peaks at +1% and Turkey somewhat higher) before their gains begin to reverse as their own demographic transitions mature and their working-age populations stabilize or begin to shrink. Their demographic dividends are time-limited, a finding with important implications for development strategy.

The rest-of-world aggregate, which is dominated by Sub-Saharan African and South Asian countries with the youngest demographics globally, gains 101% by 2050. While this figure should be interpreted cautiously given the heterogeneity within this aggregate, it underscores that the global centre of gravity for labour-intensive manufacturing is shifting toward the demographically youngest regions.

The total welfare effects, which combine both channels, are at the upper end of the macroeconomic aging literature, consistent with the model's role as a benchmark that holds technology and capital fixed. Cooley et al. (2024) find cumulated demographic effects of 17–20% on European GDP per capita through mid-century, close to this model's estimate for Germany (−18%). Kitao (2015) estimates Japan's GDP per capita declines by 5–7% in the long run from aging, substantially smaller than this model's −34% for Japan, reflecting that her OLG framework allows for endogenous capital accumulation and labour supply responses that offset the raw workforce decline. The gap between these estimates and the present model's total effects quantifies the role of adjustment margins (capital deepening, automation, behavioural responses) that attenuate the pure demographic shock. On the positive side, Table 6 implies that India's welfare under 1970 demographics would have been 76% below the 2014 baseline, and Korea's 67% below, corresponding to annualised demographic effects of approximately 3.3 and 2.5 percentage points per year, respectively, over the 1970–2014 period. These are at the upper end of the demographic dividend literature: Bloom and Williamson (1998) estimate that the dividend contributed 1.4–1.9 percentage points per year to East Asian growth during 1965–1990, cumulating to 40–60% over the period. The model's larger estimates are expected, as it captures the full demographic effect on output with no endogenous capital or labour supply offsets.

5.2. Divergence and Convergence

The welfare decomposition reveals a striking temporal pattern in the skill-composition channel. Looking backward from the 2014 baseline, the counterfactual composition effects are *positive* for all countries, meaning that earlier demographic structures, which were more diverse across countries, supported higher composition welfare than the 2014 baseline. This reflects the fact that during the era of demographic divergence, cross-country skill diversity was greater, creating more scope for gains from trade based on age-composition. Looking forward, the composition effects become *negative* for nearly all countries: demographic *convergence*, as previously young countries (China, Korea, India) age toward the global mean, further reduces cross-country skill diversity. Figure 3 traces this pattern for selected countries. The gains from trade based on demographic comparative advantage are eroding as countries' age structures converge. The era of demographic divergence that generated the composition gains documented by Cai and Stoyanov (2016) is giving way to a period in which the comparative advantage channel works against, rather than for, global welfare.

5.3. Bilateral Trade Reallocation

The welfare effects operate through a reallocation of bilateral trade flows. Table 4 reports the projected percentage change in bilateral exports for key country pairs by 2050.

The most striking pattern is the projected decline in trade from rapidly aging East Asian economies to major import markets, and the offsetting rise from young economies. China's exports to the United States fall by 10%, and to Germany by 27%. Japan's exports to the United States decline by 16%, and Korea's by 20%. These declines reflect the erosion of these countries' comparative advantage in labour-intensive manufacturing as their effective skill endowments shift.

Simultaneously, India's exports to the United States rise by 34%, and Mexico's by 26%. Indonesia's exports to China rise by 11%, partly offsetting China's own production decline in labour-intensive sectors. This pattern parallels the "China-plus-one" reorientation that firms and policymakers have pursued largely for geopolitical reasons, suggesting that structural demographic forces may independently reinforce the same shift.

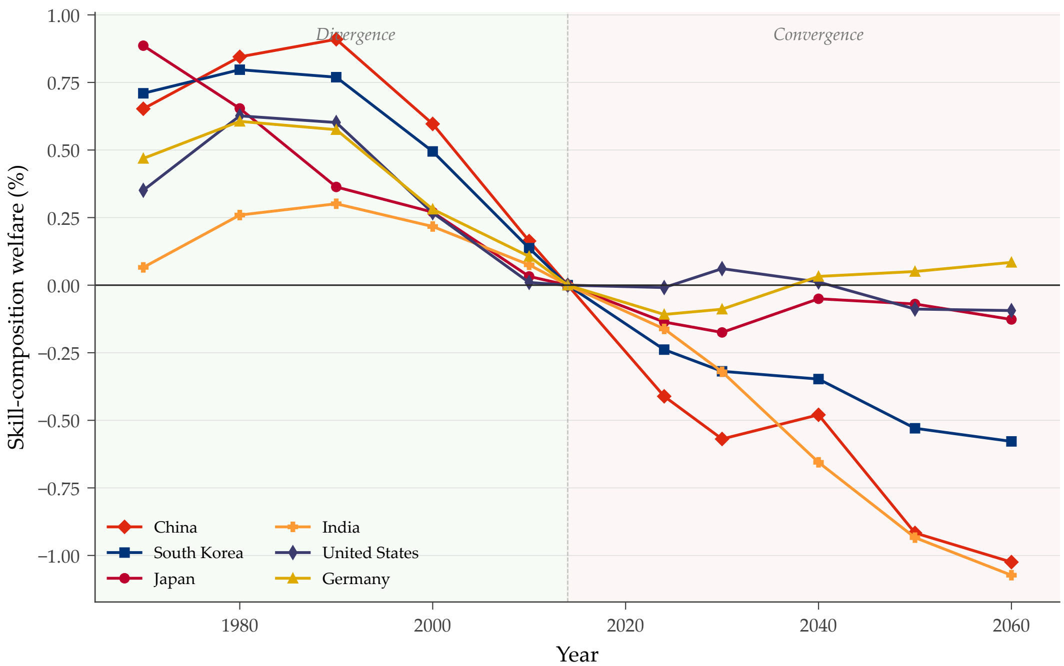

Figure 3: *Skill-composition welfare relative to 2014 baseline. Each line traces the composition-only welfare effect of replacing 2014 demographics with the demographic structure of the indicated year, holding the 2014 trade structure fixed. Points left of the dashed line are backward-looking counterfactuals (historical demographics); points to the right are forward-looking (projected demographics).*

Applying the same level–composition decomposition to bilateral trade flows reveals that the comparative advantage channel contributes meaningfully to specific corridors even though it is a small fraction of the total. China's exports to the United States fall by 10% overall, of which 10.9% reflects the shrinking Chinese workforce (level) and +0.9% reflects China's improving comparative advantage in appreciating-skill sectors (composition). China's aging actually *helps* its position in the high-skill exports that dominate the US corridor. Indonesia's exports to China show the largest composition effect (+1.1%), reflecting Indonesia's growing young-worker advantage in sectors that aging China increasingly demands. These composition effects range from 0.2% to 1.1% across corridors, a small fraction of the total bilateral changes but at the upper end of existing estimates for major trade-policy shocks.

Figure 4 provides a look at the China–US bilateral relationship. Panel (a) shows the difference in working-age population shares (WAP). China's WAP share exceeded that of the US from roughly 1990 to the mid-2040s, a window created by the one-child policy's compression of China's age structure. Panel (b) overlays the model's predictions with actual bilateral trade data. The model isolates the pure demographic channel: all non-demographic parameters (trade costs, technology, input–output structure) are held fixed at

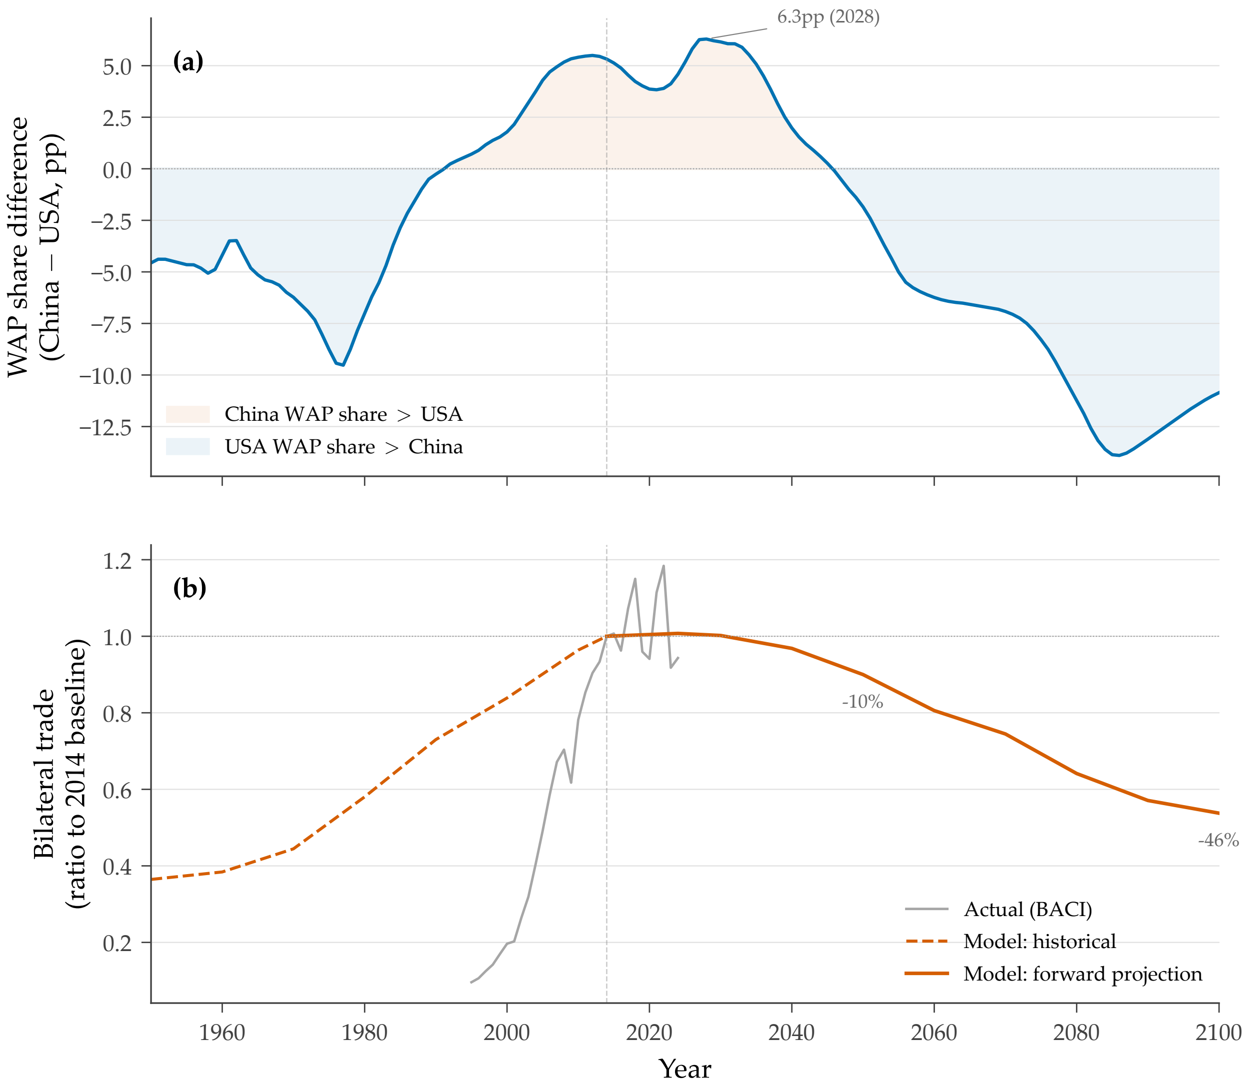

Figure 4: *China–US: Working-Age Population Shares and Bilateral Trade. Panel (a) plots the difference between China's and the US's working-age population shares (ages 20–64 as a fraction of total population), from UN WPP 2024. Shading indicates which country has the larger share. Panel (b) shows model-predicted bilateral trade (orange, dashed for historical counterfactuals, solid for forward projections) alongside actual trade data from BACI (gray), both normalized to 2014. The gap between the actual and model series reflects non-demographic factors (WTO accession, technology transfer, trade policy); the model isolates the pure demographic channel.*

their 2014 values via the hat algebra. Actual CHN–US trade grew roughly tenfold between 1995 and 2014, driven overwhelmingly by China's WTO accession and integration into global supply chains. The model's demographic channel accounts for a more modest arc, ranging from 64% below the 2014 baseline with 1950 demographics to the peak around 2014–2024.[8] The model generates substantial forward predictions: a 10% decline by 2050 and a 46% decline by 2100 as China's demographic advantage reverses.

One surprising finding is that Brazil's exports to China fall by 20% despite Brazil being a younger economy than China. This reflects the input–output linkages in the model: Brazil is a major supplier of intermediate inputs (iron ore, processed metals) to Chinese manufacturing, and as Chinese manufacturing output declines due to aging, so does its demand for Brazilian intermediates. This highlights the importance of general equilibrium effects that would be missed by a partial equilibrium analysis.

Table 4: *Projected Change in Bilateral Exports (% Change from 2014 Baseline)*

Corridor	2024	2030	2040	2050	2060
CHN → USA	+0.7	+0.2	−3.2	−10.0	−19.4
IND → USA	+15.2	+22.9	+31.2	+34.4	+33.1
MEX → USA	+11.1	+16.9	+23.2	+25.8	+25.4
DEU → USA	+1.7	−2.4	−4.9	−5.6	−7.5
JPN → USA	−2.6	−4.1	−10.4	−16.2	−18.5
KOR → USA	+2.3	−1.4	−9.5	−19.7	−28.7
IDN → CHN	+9.9	+14.5	+16.2	+11.1	+0.8
BRA → CHN	+2.0	+0.1	−7.3	−20.0	−34.5

Notes: Each entry reports the percentage change in total bilateral exports from the row exporter to the indicated importer, relative to the 2014 baseline. Counterfactual demographics are from UN World Population Prospects 2024 medium variant. All other model parameters held at 2014 values.

5.4. Sectoral Decomposition

The aggregate bilateral trade effects mask substantial sectoral heterogeneity. Table 5 decomposes the change in China's exports to the United States by sector.

The key finding is that China's aging affects different sectors very differently, and the pattern maps directly onto the skill-intensity matrix $\{\alpha_j^k\}$. Computers and electronics (ISIC 26) is essentially unaffected: China's exports in this sector decline by only 0.2% by 2050.

[8] One would not expect demographics to be the sole driver of this trade relationship.

This sector has the highest appreciating-skill intensity ($\alpha^A = 0.586$) of any manufacturing sector, meaning that China's shift toward older workers actually reinforces its comparative advantage. The sector's heavy reliance on engineers, managers, and technical specialists—whose communication, comprehension, and judgment skills improve with experience—insulates it from the demographic headwinds facing Chinese manufacturing more broadly.

At the other extreme, petroleum products (ISIC 19) declines by 67%, textiles (ISIC 13) by 19%, and basic metals (ISIC 24) by 17%. These sectors have low appreciating-skill intensity and high reliance on physical abilities and depreciating cognitive skills, making them vulnerable to workforce aging. The pattern is consistent across all sectors: the correlation between α_j^A and the change in exports is strongly positive (Figure 5).

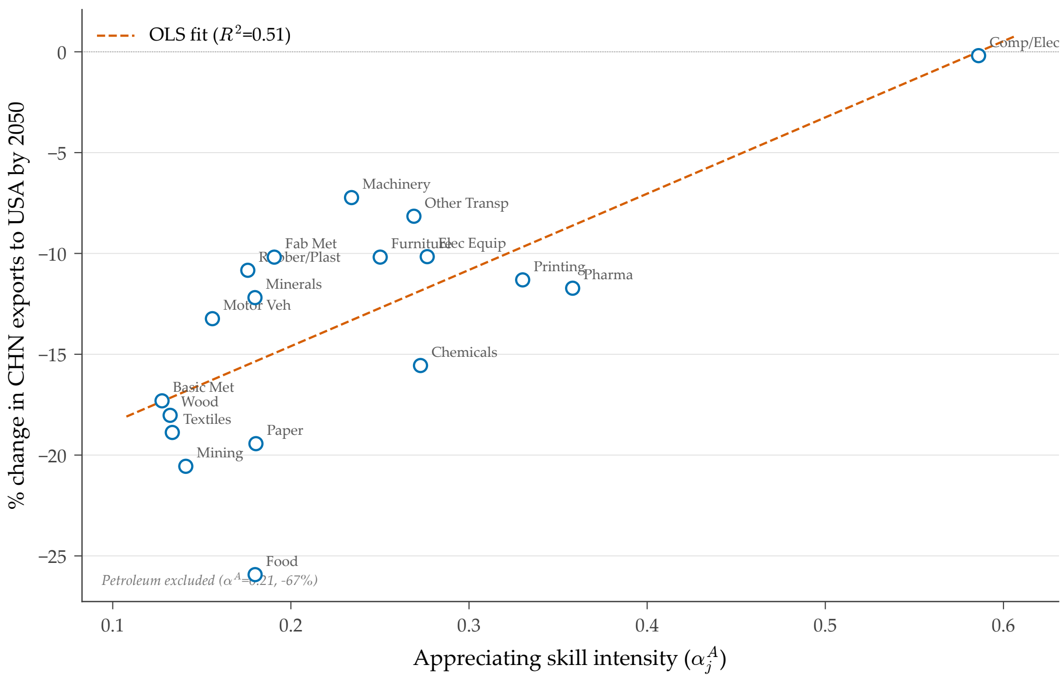

Figure 5: *Sector Trade Changes vs. Appreciating Skill Intensity: China to USA, 2050. Each point represents a tradeable goods sector, plotting the projected percentage change in China's exports to the US (2050 vs. 2014 baseline) against the sector's appreciating-skill intensity α_j^A. The 20 model sectors comprise 19 manufacturing industries (ISIC Rev 4 Divisions 10–33, excluding Tobacco) plus Mining (ISIC Section B). Computer and electronics ($\alpha^A = 0.59$) is essentially unaffected. The dashed line shows the OLS fit ($R^2 = 0.51$). Petroleum is excluded from the figure and regression as an outlier ($\alpha^A = 0.21$, -67%); its extreme decline reflects exceptionally high physical-skill intensity combined with a small base trade share.*

This sectoral heterogeneity has practical implications for narratives around recent structural shifts in trade relationships. The model suggests that the demographic pressure to diversify supply chains away from China is concentrated in routine, labour-intensive manufacturing—textiles, metals, wood products—while high-skill sectors such as electronics and pharmaceuticals face little demographic pressure. To the extent that demographic forces contribute to supply-chain diversification, the model predicts that their effect will be selective rather than across-the-board.

Table 5: *Sectoral Decomposition: China's Exports to the United States by 2050*

Sector	% Change	Base share (%)	Mechanism
Computer & Electronics	−0.2	25.2	High α^A
Machinery	−7.2	9.0	High α^A
Electrical Equipment	−10.2	9.4	High α^A
Rubber & Plastics	−10.8	2.9	Moderate α^A
Furniture & Other Mfg	−10.2	4.9	Moderate α^A
Fabricated Metals	−10.2	3.5	Moderate α^A
Pharmaceuticals*	−11.7	0.4	High α^D
Non-metallic Minerals	−12.2	1.2	High α^D/α^P
Motor Vehicles	−13.2	3.1	High α^D/α^P
Chemicals	−15.6	7.0	High α^D/α^P
Basic Metals	−17.3	3.4	High α^P
Wood	−18.0	0.7	High α^P
Textiles & Apparel	−18.9	14.0	High α^D/α^P
Paper	−19.4	1.0	High α^D
Food & Beverages	−25.9	1.7	High α^P
Petroleum	−67.1	0.4	High α^P

Notes: % change in CHN→USA exports by sector, 2050 vs. 2014 ($\rho = 0.5$). Mechanism: α^A = appreciating cognitive, α^D = depreciating cognitive, α^P = physical. Four sectors with <1% combined base share omitted. *Pharmaceuticals' high α^D reflects manufacturing occupations (production-line workers), not R&D.

5.5. Policy Counterfactual: Retirement Age Extension

I also consider a policy-relevant counterfactual: extending the effective retirement age by five years in aging OECD countries. This is implemented by expanding the working-age population to include the 65–69 age group, with skill endowments determined by the calibrated age–productivity profiles evaluated at those ages. This is a stylised experiment: in practice, effective retirement ages vary across countries (Japan's is already 66.7 for men, above the OECD average of 64.4), and the policy instruments used to achieve later retirement differ substantially—ranging from statutory age increases to employer mandates and pension incentives. The exercise should therefore be interpreted as illustrating the potential magnitude of the trade channel of retirement extension, not as a precise policy recommendation for any individual country.

The total welfare gains from retirement extension are substantial. Japan recovers 11.3 percentage points of its total welfare loss from aging, equivalent to offsetting roughly one decade of demographic decline. Germany gains 8.7 percentage points, the United States 8.6 points, and Korea 7.0 points.

Applying the same level–composition decomposition as above reveals a nuanced picture. Nearly all of the welfare gain comes from the workforce-size channel: retaining the 65–69 age group in the labour force increases aggregate output. The skill-composition channel, however, is slightly *negative* (-0.17% for Korea to -0.37% for the United States). This is because adding workers at the extreme of the appreciating-skill profile pushes aging economies' skill mix further from the global average, exacerbating their over-specialisation in appreciating-skill sectors. The resulting terms-of-trade deterioration, as the unilateral increase in appreciating-skill supply depresses the return to these countries' already-abundant factor, more than offsets any gains from greater skill diversity. Retirement extension is thus an effective policy for offsetting the aggregate workforce decline, but it does not resolve the comparative advantage distortion that is the focus of this paper. Addressing the composition channel would instead require policies that rebalance the age structure of the workforce, such as immigration of younger workers or investments in retraining that shift older workers toward depreciating-skill tasks.

This finding has direct policy relevance. Many OECD countries are actively debating or implementing increases in statutory retirement ages in response to fiscal pressures from population aging. The model suggests that such policies may carry large aggregate welfare benefits, but that the sectoral competitiveness implications are more subtle than the headline numbers suggest.

6. Conclusion

This paper develops the first calibrated, multi-country, multi-sector trade model in which demographic age structure is a source of comparative advantage. The mechanism is intuitive: workers of different ages supply different skills, sectors use these skills with different intensities, and a country's age distribution therefore determines its pattern of sectoral specialization. By embedding this demographic channel into the workhorse Caliendo and Parro (2015) Ricardian trade framework, the model generates quantitative predictions about how the global demographic transition will reshape international trade patterns through mid-century.

Reduced-form evidence from 204 countries confirms the core mechanism, and a partial-equilibrium decomposition shows that the calibrated model recovers 84% of the empirical

Rybczynski coefficient when evaluated on the same footing as the regressions. General equilibrium adjustment compresses the trade-composition response by a factor of ten, converting it into the welfare-relevant price changes that the reduced-form evidence cannot quantify. The skill-composition channel generates welfare effects ranging from -0.9% to $+0.05\%$, comparable to standard trade-policy benchmarks (Arkolakis, Costinot, and Rodríguez-Clare, 2012; Caliendo and Parro, 2015), and exhibits a striking temporal pattern: historical demographic divergence supported positive composition gains, but as countries' age structures converge these gains are reversing. The total demographic welfare effects are much larger (up to -44% for Taiwan, $+52\%$ for India), but these are dominated by the workforce-size channel and serve as partial-effect benchmarks that complement estimates from the macroeconomic aging literature. At the sectoral level, the model identifies which sectors and bilateral corridors are most demographically exposed: high-skill sectors like electronics are largely insulated from aging, while labour-intensive sectors like textiles and metals bear the brunt.

Several limitations point toward productive extensions. First, the model abstracts from current account dynamics: countries cannot run trade surpluses or deficits, ruling out the saving channel through which aging affects international capital flows (Auclert et al., 2025; Sposi, 2022). Incorporating current account imbalances would allow the model to jointly address the saving and comparative-advantage channels of demographics. Second, the baseline model treats demand as independent of demographics. Allowing expenditure shares to depend on the age distribution would add a demand-side channel through which aging shifts consumption toward non-tradeables and reduces aggregate trade volumes. Calibrating such an extension credibly would require age-specific expenditure data at the sectoral level across a broad set of countries—a data construction exercise that is beyond the scope of this paper but represents a promising complement to the supply-side mechanism studied here. Third, the model does not include capital accumulation or automation. Acemoglu and Restrepo (2017) show that aging countries have not experienced slower aggregate growth, potentially because they adopt labour-saving technologies more aggressively (Acemoglu and Restrepo, 2022). To the extent that automation substitutes for the depreciating skills of young workers, it could partially offset the comparative advantage shifts predicted by the model. Fourth, expanding the country sample beyond the current 30 countries would enable a richer analysis of the developing world. The present model aggregates all non-WIOD countries into a single rest-of-world composite, which receives a population-weighted average demographic shock. This is a standard simplification that does not affect the base-year calibration: the hat algebra operates on observed bilateral trade shares, and adding countries to the model would not alter the base-year shares for existing countries. However, because the RoW composite is demographically heterogeneous, applying a population-weighted average demographic shock introduces aggregation bias into the counterfactual responses, as the non-linear price changes from an averaged shock differ from those of disaggregated regions. This bias propagates to all trading partners through multilateral resistance terms, though its magnitude is likely small relative to the direct demographic effects on individually modeled countries. However, it does limit the questions the model can answer. In particular, the model cannot separately identify welfare or trade effects for countries inside the rest-of-world composite. Sub-Saharan Africa, the world's youngest and fastest-growing population, is the most important omission. Africa's working-age population share is projected to rise sharply through mid-century, and the reduced-form evidence in Section 3 (which does cover African countries) suggests that this demographic dividend should shift African export composition toward depreciating-skill-intensive sectors. Whether Africa can realize this potential and become the next frontier of labour-intensive manufacturing as East Asia ages out of it remains an open question and is a question that would require not only expanding the country sample but also confronting the fact that many African economies have not yet translated favourable demographics into manufacturing export capacity. This is a natural direction for future work.

These extensions notwithstanding, the central message of the paper is clear. The demographic transition is not merely a domestic challenge of pension reform and healthcare provision. It is a force that will meaningfully reshape who trades what with whom. A country's age structure, determined decades earlier by fertility decisions, is a significant predictor of its future trade patterns, and the quantitative effects are large enough to warrant serious attention from both trade economists and policymakers.

References

Acemoglu, D. and P. Restrepo (2017). Secular stagnation? The effect of aging on economic growth in the age of automation. *American Economic Review: Papers and Proceedings 107*(5), 174–179.

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

Arkolakis, C., A. Costinot, and A. Rodríguez-Clare (2012). New trade models, same old gains? *American Economic Review 102*(1), 94–130.

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.

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.

Bloom, D. E., D. Canning, and J. Sevilla (2003). *The Demographic Dividend: A New Perspective on the Economic Consequences of Population Change*. Santa Monica, CA: RAND Corporation.

Bloom, D. E. and J. G. Williamson (1998). Demographic transitions and economic miracles in emerging Asia. *World Bank Economic Review 12*(3), 419–455.

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.

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

Cai, J. and A. Stoyanov (2023). Progressive income tax and comparative advantage in trade. *Economic Modelling 129*, 106553.

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.

Caron, J., T. Fally, and J. R. Markusen (2014). International trade puzzles: A solution linking production and preferences. *Quarterly Journal of Economics 129*(3), 1501–1552.

Cooley, T. F., E. Henriksen, and C. Nusbaum (2024). Demographic obstacles to European growth. *European Economic Review 169*, 104852.

Costinot, A., D. Donaldson, and I. Komunjer (2012). What goods do countries trade? A quantitative exploration of Ricardo's ideas. *Review of Economic Studies 79*(2), 581–608.

Costinot, A. and A. Rodríguez-Clare (2014). Trade theory with numbers: Quantifying the consequences of globalization. In G. Gopinath, E. Helpman, and K. Rogoff (Eds.), *Handbook of International Economics*, Volume 4, pp. 197–261. Elsevier.

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 (1995). Gender differences in upper extremity musculoskeletal complaints in the working population. *International Archives of Occupational and Environmental Health 68*(3), 189–196.

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

Do, Q.-T., A. A. Levchenko, and C. Raddatz (2016). Comparative advantage, international trade, and fertility. *Journal of Development Economics 119*, 48–66.

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

Feyrer, J. (2007). Demographics and productivity. *Review of Economics and Statistics 89*(1), 100–109.

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

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

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

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.

Levchenko, A. A. and J. Zhang (2016). The evolution of comparative advantage: Measurement and welfare implications. *Journal of Monetary Economics 78*, 96–111.

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.

Papke, L. E. and J. M. Wooldridge (1996). Econometric methods for fractional response variables with an application to 401(k) plan participation rates. *Journal of Applied Econometrics 11*(6), 619–632.

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

Santos Silva, J. M. C. and S. Tenreyro (2006). The log of gravity. *The Review of Economics and Statistics 88*(4), 641–658.

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

Skirbekk, V. (2004). Age and individual productivity: A literature survey. *Vienna Yearbook of Population Research 2*, 133–153.

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

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

A. Data Sources

This appendix describes the six primary data sources used in the paper.

BACI (CEPII). Bilateral trade data at the HS 6-digit product level, 1995–2024. BACI harmonizes UN Comtrade mirror flows, reconciling discrepancies between reported exports and imports. I aggregate from HS 6-digit products to 19 ISIC Revision 4 manufacturing sectors using the standard HS–ISIC concordance. The sample includes 204 exporting countries in each year.

United Nations World Population Prospects 2024. Population estimates and projections by 5-year age group for 237 countries, 1950–2100. I use the medium-variant projections for forward counterfactuals. Historical estimates (1950–2024) are based on observed data; projections (2025–2100) incorporate assumptions about future fertility, mortality, and migration.

World Input–Output Database (WIOD) 2016. Inter-country input–output tables for 43 countries and 56 sectors, 2000–2014. I extract bilateral trade shares, input–output coefficients, value-added shares, and final consumption shares for the base year 2014, aggregated to 30 countries and 20 manufacturing sectors plus one non-tradeable sector.

O*NET 29.1. The Occupational Information Network provides detailed ability and skill descriptors for 923 occupations defined at the SOC 2018 8-digit level. I use the 18 cognitive and physical ability importance ratings (Scale ID = IM) for 763 occupations with non-missing data, and apply principal component analysis to construct three composite skill types.

Bureau of Labor Statistics OEWS May 2014. The Occupational Employment and Wage Statistics survey provides employment counts by detailed occupation (SOC 2010, 6-digit) for each 4-digit NAICS industry. I use these to construct employment-weighted industry-level skill scores from the O*NET occupation-level data. Because O*NET 29.1 uses the SOC 2018 classification while the OEWS 2014 uses SOC 2010, the merge is performed via the BLS SOC 2010–2018 crosswalk; in cases of one-to-many mappings, employment counts are apportioned equally across the successor codes.

Penn World Table 10.01. Capital stock and employment data for 183 countries, 1950–2019. I use these to construct country–year capital–labour ratios for the capital-intensity robustness check in Table 1, column 7. Sector-level capital shares are computed from the WIOD Socio-Economic Accounts.

B. Historical Welfare Decomposition

Table 6: *Historical Welfare Decomposition: $\hat{W}_c$ Relative to 2014 Baseline*

	1970	1980	1990	2000	2010
Large demographic transitions					
India	0.239	0.337	0.476	0.660	0.897
Mexico	0.224	0.329	0.497	0.701	0.910
Indonesia	0.264	0.365	0.534	0.735	0.930
China	0.295	0.430	0.628	0.811	0.965
Korea	0.333	0.493	0.703	0.871	0.957
Mature economies					
Japan	0.841	0.980	1.065	1.130	1.066
Germany	0.806	0.856	1.000	1.032	1.007
USA	0.460	0.603	0.723	0.847	0.966
France	0.663	0.757	0.870	0.933	1.005
UK	0.721	0.756	0.838	0.894	0.982

Notes: Baseline calibration ($\rho = 0.5$). Each entry reports the counterfactual welfare ratio $\hat{W}_c \equiv \hat{Y}_c / \prod_j (\hat{P}^j_c)^{\eta^j_c}$ (the change in real income) when the 2014 trade structure is held fixed but country c's demographic structure is replaced by its value from the indicated year. Values below 1 indicate that the historical demographic structure would have yielded lower welfare than the actual 2014 demographics; values above 1 indicate higher welfare. The 2014 baseline is normalized to 1.