Blueprint Labs

Discussion Paper #2026.08

Baby Busts and Growth Booms: Demographic Change and the Macroeconomy

Daron Acemoglu
David Autor
Keelan Beirne
Andrew Scott

June 2026

The views expressed in this paper are those of the authors and do not necessarily reflect the views of MIT Blueprint Labs, the Massachusetts Institute of Technology, or any affiliated organizations. Blueprint Labs working papers are circulated to stimulate discussion and invite feedback. They have not been peer-reviewed or subject to formal review processes that accompany official Blueprint Labs publications.

MIT Department of Economics
77 Massachusetts Avenue, Bldg. E53-390
Cambridge, MA 02139

National Bureau of Economic Research
1050 Massachusetts Avenue, 3rd Floor
Cambridge, MA 02138

Baby Busts and Growth Booms: Demographic Change and the Macroeconomy*

Daron Acemoglu — MIT
David Autor — MIT
Keelan Beirne — MIT
Andrew Scott — LBS and EIT

June 24, 2026

Abstract

The secular decline in birth rates across the globe over the past seven decades has slowed population growth, raised average ages, and reshaped labor markets and the macroeconomy. Contrary to the widespread expectation that these trends hamper economic growth, we find lower birth rates are associated with higher growth in GDP per working-age adult across countries and higher wage growth across US commuting zones, with no negative impact on aggregate GDP or earnings. These patterns are not explained by educational upgrading, rising female labor force participation, the declining importance of agriculture, or neoclassical-Solow mechanisms. We argue that they reflect the endogenous, labor-saving response of technology to the scarcity of younger workers. Consistent with this interpretation, countries and regions with lower birth rates exhibit more labor-saving patents and growing high-tech activity. There is also higher TFP growth across countries and industries. Exploiting cross-country variation in WWII military and civilian deaths, we find that declines in younger population, rather than population size per se, drive our results.

Keywords: aging, automation, birth rates, demographic change, economic growth, directed technological change, labor scarcity, population growth.
JEL Classification: J11, J31, O33, O40, E24.

*We thank Lauren Qu and Joanne Yongyin Liang for excellent research assistance. Acemoglu and Autor gratefully acknowledge financial support from the Hewlett Foundation, the Smith Richardson Foundation, and the James and Cathleen D. Stone Foundation. Autor also thanks the Google Technology and Society Visiting Fellows Program, the NOMIS Foundation, and the Schmidt Sciences AI2050 Fellowship for research support. Scott acknowledges ESRC grant T002204. We are also grateful to Iris Kesternich for generously sharing with us data on World War II deaths.

1 Introduction

Birth rates have declined sharply across the world over the past seven decades, driving a historic demographic transition. The global average (crude) birth rate was 3.78 per 100 people in 1950; by 2025, it had fallen by more than half to 1.71.[1] This decline is evident across all continents, as shown in Figure 1, although its pace and magnitude are not uniform. The United Nations projects that population on all continents except Africa and North America will level off and then fall in the second half of this century, generating the first sustained decline in global population since the bubonic plague of the 14th century. The countries with the lowest birth rates today—China at 0.63, Japan at 0.6, and the Republic of Korea at 0.46—are projected by 2050 to experience population declines of 20–30% relative to 2023; to see their share of the population aged 65+ reach 30–40%; and to have older workers (aged 45+) comprise 60% of the labor force.[2]

Figure 1: Birth Rates across Continents and Selected Countries, 1950–2020

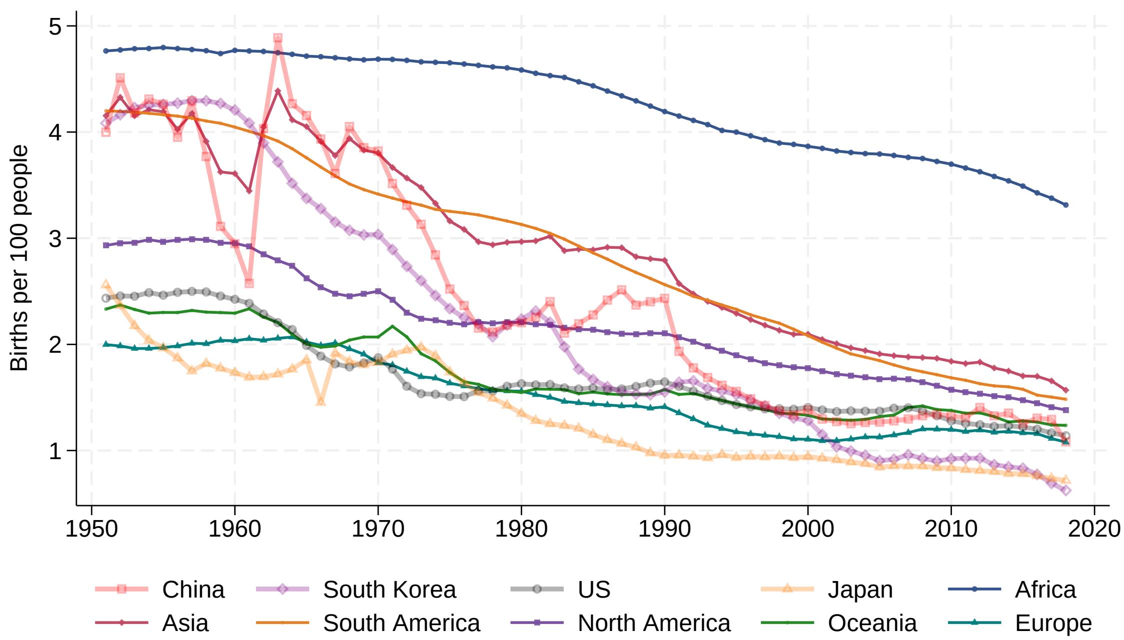

Notes: This figure plots birth rates over time across continents. The birth rate is defined as births per 100 people. The figure also reports birth rates for a select set of individual countries. Continents follow the UN classification; the named country series are also included in their continent's aggregate. Data are taken from the UN Population Prospects (United Nations, Department of Economic and Social Affairs, Population Division, 2024).

[1] We focus on crude birth rates (which we refer to as "birth rates") though similar trends hold for total fertility rates (expected number of lifetime births per woman in the population). We use birth rates rather than fertility rates because of longer-run data availability, the lack of total fertility rate data for US commuting zones, and the direct link between current birth rates and current population change, as we show in Section 4.

[2] All numbers in this paragraph are from the UN Population Prospects (United Nations, Department of Economic and Social Affairs, Population Division, 2024).

These demographic trends are a source of alarm among economists and policymakers (International Monetary Fund, 2025; Organisation for Economic Co-operation and Development, 2025; European Bank for Reconstruction and Development, 2025) because older workers are widely assumed to be less productive, less innovative, less entrepreneurial, and less technologically adept than younger adults (see discussion Section 2). Smaller and aging populations are also expected to generate weaker aggregate demand, a shortfall of able-bodied workers, and fewer new ideas (Summers, 2016; Gordon, 2017; Jones, 2022; Hayashi, 2025). Together with rising pension and healthcare costs, these concerns fuel pessimism about economic growth prospects in both the industrialized and the developing world.[3]

While these concerns rest largely on forward projections, the global economy has already experienced sharp demographic changes of historic magnitude over several generations. Beginning in the early 1940s, the baby boom temporarily raised birth rates across much of the industrialized world, followed by the baby bust of subsequent decades. Many developing economies also underwent their own rapid demographic transitions, with dramatic declines in fertility rates often compressed into just a few decades.

This paper analyzes seven decades of demographic change to assess the impact of aging and declining populations on economic performance across countries and within the United States. Our findings challenge the prevailing pessimism: lower birth rates, and the aging and shrinking populations they have produced, have raised rather than lowered GDP per "worker" (our shorthand for GDP per working-age adult) during these decades. The gain in GDP per worker has been large enough to fully offset the negative effect of population decline, leaving aggregate GDP broadly unaffected.

Figure 2 summarizes our main result by plotting the relationship between 1950 birth rates and subsequent economic growth from 1970 to 2020 across countries (left panels) and US commuting zones (right panels). For countries, we use GDP per worker and aggregate GDP as outcome variables. For commuting zones, we use the average (composition-adjusted) weekly wage rate as a proxy for income per worker and total earnings as the analogue of aggregate GDP. In each case, we control for initial log GDP per worker (or wage) and log numbers of younger and older workers (see equation (10) in Section 6 for the exact specification). Our main analysis focuses on birth rates on the right-hand side because, as shown in Section 4, changes in population and age composition are driven by birth rates, rather than changes in birth rates.

The first two panels of Figure 2 reveal a striking, quantitatively large, and robust *negative* relationship between birth rates and growth in GDP per worker (Panel A) or growth in composition-adjusted wages in commuting zones (Panel B). That is, lower birth rates—and hence aging and slower population growth—are associated with higher GDP per worker. For example, across countries, a one percentage point lower birth rate is associated with 26.8% higher GDP per worker.

[3] Media coverage echoes the same concerns. See: *The Economist*: https://www.economist.com/interactive/briefing/2025/09/11/humanity-will-shrink-far-sooner-than-you-think, *The Atlantic*: https://www.theatlantic.com/ideas/archive/2025/06/birth-rate-population-decline/683333/, and the *Financial Times* https://www.ft.com/content/19cea1e0-4b8f-4623-bf6b-fe8af2acd3e5.

Figure 2: 1950 Birth Rates and Income Growth from 1970 to 2020

Notes: This figure plots Frisch-Waugh residuals from the regressions $\Delta_{50}y_{c,2020} = \alpha + \beta b_{c,1950} + \boldsymbol{x}'_{c,1970}\boldsymbol{\gamma} + \varepsilon_c$ where $\Delta_{50}y_{c,2020}$ is income growth between 1970 and 2020, $b_{c,1950}$ is the birth rate in 1950 and $\boldsymbol{x}_{c,1970}$ is a vector of controls including initial income in 1970 alongside the log number of younger workers (20–45) and the log number of older workers (45–70). In Panel A y is log GDP per working-age (20–70) adult across countries, while in Panel B y is the composition-adjusted wage rate across commuting zones. In Panel C y is log aggregate GDP across countries and in Panel D y is the aggregate earnings across commuting zones. To minimize the impact of the COVID-19 pandemic, for countries we proxy for 1970–2020 income growth using 1970–2019 income growth, re-scaled to a 50-year period. Cross-country regressions are unweighted, while commuting-zone regressions are weighted by working-age population in 1970.

The results are qualitatively similar across commuting zones, though smaller in magnitude. The bottom two panels show the implications for overall GDP and aggregate earnings: despite declining or slower-growing populations, *aggregate* incomes do not fall.[4] The patterns summarized in Figure 2 are robust to adjusting for the initial level of education, urbanization, differential regional trends, and initial sectoral composition across commuting zones.

Why do slowdowns in population growth and aging workforces predict faster economic growth? We hypothesize that this reflects the endogenous response of technology to *labor scarcity*. We present a version of the framework in Acemoglu (2010), showing that labor scarcity triggers the adoption of *labor-saving technologies* that can lead to increases in GDP per worker and even aggregate GDP. Our evidence supports this interpretation. In cross-country data, declining birth rates lead to higher total factor productivity (TFP), larger capital stocks, a shift toward exports in high-tech industries, and more labor-saving patenting. In US commuting zone data, lower birth rates spur an employment shift toward high-tech industries and an increase in labor-saving patenting. In US industry data, we additionally find positive TFP and capital stock responses (though the latter are typically not statistically distinguishable from zero). We argue that it is this technological response that produces the positive relationship between baby busts and subsequent growth booms.

We consider four alternative channels through which lower birth rates might raise GDP per worker. First, the standard neoclassical growth model predicts that falling population increases average income in the short run, but lowers aggregate GDP and investment; the evidence contradicts this prediction (see footnote 4). Second, lower birth rates may facilitate female labor force participation, but in the data we find no such response. Third, lower birth rates may increase human capital investments by causing substitution of quality for quantity in child-rearing or increasing per-pupil resources in schools; the evidence does not support a major role for this channel. Fourth, lower birth rates may foster a transition from agriculture to manufacturing; this channel too receives limited support. Overall, the data do not point to any of these alternative mechanisms as an important contributor to the results we document in this paper.

Because lower birth rates simultaneously reduce population growth and raise the share of older people, our primary results are silent on the relative importance of population decline versus population aging. In the final section of the paper, we explore this question by studying changes in population resulting from WWII casualties. This approach allows us to distinguish civilian deaths, which primarily reduce population without changing the age structure, from military deaths, which both reduce population *and* raise the average age of the workforce. We show that total war deaths are associated with lower future GDP per worker, while military deaths predict *higher* GDP per worker, suggesting that it is the scarcity of younger workers that induces the positive productivity

[4] The finding of zero impact on aggregate incomes distinguishes our results from neoclassical-Solow growth models, which predict that declining population raises GDP per worker in the short run (as the capital-labor ratio increases) but always lowers aggregate GDP. We find that aggregate output does *not* decrease with lower birth rates. Moreover, the neoclassical-Solow models predict that capital stocks should decline in response to lower birth rates; we show in Section 7 that they increase.

response. This pattern parallels Bergeaud et al. (2025), who document that French regions with higher WWI military mortality experienced increased patenting in labor-saving technologies.

The remainder of the paper proceeds as follows. Section 2 situates our contribution in the literatures on the macroeconomic consequences of demographic change and on endogenous, labor-saving technical change. Section 3 presents a model that illustrates how labor scarcity induces adoption of labor-saving technologies and can increase aggregate income. Section 4 derives our key empirical specifications by formally showing that population dynamics are determined by birth rates rather than changes in birth rates. Section 5 describes data sources. Sections 6 and 7 present our main results and evidence that they operate through labor-saving technology adoption. Section 8 explores alternative mechanisms, and Section 9 uses WWII casualties to distinguish the effects of population size from population aging. Section 10 concludes. The (online) Appendix contains more information on data construction, robustness checks and additional results.

2 Related Literature

Neoclassical growth models predict that declining populations increase GDP per capita in the short run while reducing aggregate GDP. In the longer run, declining populations are expected to deter investment, depressing aggregate GDP even more (Solow, 1956).[5] A rising share of older people could further reduce GDP and GDP per capita by increasing dependency ratios and lowering aggregate savings (Fernández-Villaverde et al., 2025). A separate literature points in the same direction, emphasizing that larger populations contribute to higher GDP by fostering faster technological progress through increasing returns in production or idea creation (Kremer, 1993; Jones, 2022; Eden and Kuruc, 2023). A related argument holds that population pressure itself induces innovation (Boserup, 2014; Lee, 1988).

More recent work identifies additional channels through which aging may reduce GDP beyond simple reductions in working-age population. First, older workers are argued to have lower productivity than their younger counterparts (Murphy and Welch, 1990; MacDonald and Weisbach, 2004; Gordon, 2017). Older populations may also be less entrepreneurial and innovative, dampening an economy's growth potential (Liang et al., 2018; Peters and Walsh, 2021; Hopenhayn et al., 2022; Karahan et al., 2024). Finally, aging populations could create an imbalance of savings relative to investment opportunities, triggering "secular stagnation" (Hansen, 1939; Baldwin and Teulings, 2014). Our findings that aging populations lead to higher GDP per capita and broadly unchanged GDP indicate that if these channels are present, they are not the dominant ones.

Several empirical studies report negative effects of aging on economic growth, including Yoon et al. (2014); Aksoy et al. (2019); Maestas et al. (2023); Hayashi (2025). These studies typically

[5] Several papers present evidence broadly consistent with this prediction. Acemoglu and Johnson (2007) show that the spread of antibiotics and other drugs and chemicals to the developing world reduced mortality and increased population, but not GDP. Young (2005) finds that the AIDS epidemic may have increased GDP per capita in parts of Africa. None of these papers study variation in birth rates or document the offsetting productivity and investment responses we find here.

examine shorter time horizons than we do and do not focus on variation in birth rates. This distinction matters because, particularly across commuting zones, migration makes age composition endogenous to current and future growth outcomes. We document in Appendix C that the results in Maestas et al. (2023) are sensitive to the choice of control variables and the inclusion of Alaska, Hawaii, and DC. By contrast, when included in our commuting-zone specifications, the Maestas et al. (2023) variables do not affect our estimates and do not reproduce their state-level findings in our setting. We find that the results in Hayashi (2025) are likewise sensitive to the precise sample and specification used. The main explanatory variables used in Hayashi (2025) are not statistically significant in our specification; moreover, their inclusion leaves our estimates of the birth-rate effect unchanged.

Beyond the technological channel that is our focus, lower birth rates could raise GDP per worker through other indirect mechanisms. Lower birth rates may create an initial "demographic dividend" by increasing female labor force participation and reducing the dependency rate (Bloom et al., 2003). They may alternately encourage greater investment in child "quality" rather than quantity, raising human capital levels (Ashraf et al., 2013; Karra et al., 2017). We find no evidence that these channels contribute to our results.

Instead, the operative channel appears to be labor scarcity: changes in the ratio of older to younger workers stimulate the development and adoption of labor-saving technologies. The notion that factor scarcity and prices direct the course of innovation traces back to Hicks (1932). Habakkuk (1962) proposed that labor scarcity explained why the United States adopted labor-saving technologies more rapidly than Britain during the 19th century. Applying the same logic to an earlier period, Allen (2009) argued that Britain industrialized first, because it faced labor shortages and high wages, creating strong incentives to adopt labor-saving technologies.

Acemoglu (2010) formalizes the mechanism whereby labor scarcity influences the direction of technological change, showing that the impact on innovation at the margin depends on whether technological change is labor-saving or labor-complementary. When technology is labor-saving, scarcity encourages technological progress, as hypothesized by Habakkuk and Allen; when it is labor-complementary, scarcity discourages innovation. Building on this framework, Acemoglu and Restrepo (2022a) develop a model of automation and show that aging leads to faster adoption of automation technologies. Using both cross-country and cross-commuting-zone data, they provide detailed evidence that aging is associated with more rapid adoption of robots and other automation technologies, with especially pronounced effects in industries that employ a high share of younger and middle-aged workers. This conceptual framework rationalizes a range of findings, both prior and subsequent. Clemens et al. (2018) show that reduced supplies of young and less-educated workers (coming from immigration restrictions) induce the adoption of more mechanized agricultural technologies. Numerous other papers find that shortages or high prices of key inputs lead to faster technology development or adoption, including Lewis (2011), Dechezleprêtre et al. (2021), and Bergeaud et al. (2025) for labor; Hanlon (2015) for cotton during the American Civil War; Newell et al. (2010), Popp (2002), and Acemoglu et al. (2023) for energy; and Moscona and Sastry

(2023) for agriculture.

Prior work on birth rates and economic growth has been limited in geographic and temporal scope. To our knowledge, no prior work empirically documents the striking long-run increase in income per worker that accompanies falling birth rates. The two most closely related papers are Li and Zhang (2007) and Acemoglu and Restrepo (2017). Li and Zhang (2007) show that China's one-child policy, which applied only to the ethnic majority *Han* Chinese, spurred faster growth between 1978 and 1998 in provinces with larger Han majorities. Their paper does not explore mechanisms and is limited to a short time window and the intrinsically small set (28) of Chinese provinces. Consonant with our findings, Acemoglu and Restrepo (2017) show that aging between 1990 and 2015 is not associated with slower growth across countries. However, given their earlier sample period, they do not observe the sharp increase in GDP per worker that we document. Neither paper develops the broader cross-country and cross-commuting-zone evidence on patenting, sectoral reallocation, and productivity through which we link labor scarcity to labor-saving technology.

3 Conceptual Framework

We begin with a general accounting framework for the channels through which demographic change affects output, and then develop a model in which technology responds endogenously to labor scarcity.

3.1 Macroeconomic Responses to Demographic Change

Suppose aggregate output is produced by the constant returns to scale function

$$Y = F(K, L, H, \theta),$$

where K is capital, L denotes younger workers, H is older workers (whom we treat as equivalent to skilled workers, since they perform different tasks than younger workers), and θ is an index of technology (to be specified further in the next subsection).[6]

We are interested in the implications of a demographic change, such as a decline in birth rates, that reduces the supply of younger workers. For simplicity, we hold the supply of older/skilled workers fixed over the horizons we consider,[7] and focus on the effect of a decline in L. We assume throughout that all factor markets are competitive. The direct (proportional) effect of a change in the supply of younger labor on GDP, Y, is given by

$$\begin{aligned} \frac{\partial \log Y}{\partial \log L} &= \frac{\partial F(K, L, H, \theta)}{\partial L} \frac{L}{Y} \\ &= \frac{w_L L}{Y} = s_L, \end{aligned} \tag{1}$$

[6]See Acemoglu and Restrepo (2022a) for evidence that younger workers are more likely to perform tasks that can be automated by industrial machinery.

[7]As we show below, this is a good approximation for about five decades following a birth rate decline.

where the second equality uses the fact that, under competitive markets, the marginal product of labor L equals its wage, w_L. This relationship holds for any production function and implies that a proportional decline in the supply of younger workers leads to a proportional decline in GDP equal to this factor's share in GDP.

This is not the total impact of a change in L, however, since capital and technology can also respond. Taking these indirect responses into account, we can write

$$\frac{d\log Y}{d\log L} = \frac{\partial \log Y}{\partial \log L} + \frac{\partial \log Y}{\partial \log K}\frac{d\log K}{d\log L} + \frac{\partial \log Y}{\partial \theta}\frac{d\theta}{d\log L}. \tag{2}$$

Consider now how these indirect effects operate in the neoclassical growth model. Since technology is exogenous in that model, the last term equals zero. The second term is positive because lower labor supply reduces the capital stock in the long run, with or without endogenous savings (see Acemoglu, 2010, for a proof). Hence:

$$\left.\frac{d\log Y}{d\log L}\right|_{\text{neoclassical}} = s_L + \frac{\partial \log Y}{\partial \log K}\left.\frac{d\log K}{d\log L}\right|_{\text{neoclassical}} > s_L.$$

Therefore, in neoclassical-Solow models, the medium- and longer-term impacts of labor scarcity on GDP are even more negative than labor's direct share.

With endogenous technology, in contrast, both capital and technology can increase even as labor falls, and consequently, declines in the size of younger cohorts have smaller negative, or even *positive*, effects on GDP.

3.2 Endogenous Technology Choice

We now present a model that illustrates how technology responds to labor scarcity induced by lower birth rates, and how this can ultimately increase rather than reduce GDP. The model is static for simplicity and is a special case of Acemoglu (2007, 2010).

There is a unique final good produced competitively, with the constant returns to scale aggregate production function

$$Y = q(\theta)^{\alpha}\left(F(L,K,H,\theta)\right)^{1-\alpha}. \tag{3}$$

Here $q(\theta)$ is the quantity of machines, embodying technology $\theta \in (0,1)$, supplied by a technology monopolist and combined with value added from the other factors of production, $F(L,K,H,\theta)$, which again exhibits constant returns to scale in L, K and H. The parameter $\alpha \in (0,1)$ captures the importance of machines relative to the aggregate of tasks. The inclusion of H highlights that the relevant notion of labor scarcity in our model is the scarcity of unskilled labor or younger workers, who compete with capital for tasks. We initially assume that the economy has a fixed supply of capital K that can be used for task production. We also assume F is increasing, differentiable, strictly concave in θ and satisfies Inada conditions $\lim_{\theta\to 0} \partial F/\partial\theta = \infty$ and $\lim_{\theta\to 1} \partial F/\partial\theta = 0$.[8]

[8] Note that there are two kinds of capital in this model: machines supplied by the technology monopolist and capital

Machines $q(\theta)$ are produced by a technology monopolist at constant marginal cost γ in units of the final good. The monopolist will therefore produce at the profit-maximizing price $p(\theta)$. This monopolist also chooses θ, where producing (or "designing") technology θ costs $\Gamma \cdot \theta$ in units of the final good (which is without loss of generality, since the dependence of F on θ is unrestricted).

An equilibrium consists of a quantity of machines $q(\theta)$, a technology choice $\theta \in (0,1)$, and a machine price $p(\theta) > 0$, such that: (i) given θ and $p(\theta)$, the machine quantity $q(\theta)$ maximizes profits for the competitive final good sector; and (ii) the technology monopolist's choices of θ and $p(\theta)$ maximize its profits.

The demand for machines can be obtained from the profit-maximization problem of the final good sector: $\max_{q(\theta)\geq 0}\ q(\theta)^{\alpha}\left(F\left(L,K,H,\theta\right)\right)^{1-\alpha} - p(\theta)q(\theta)$, which yields an isoelastic demand for $q(\theta)$ given by $q(\theta) = \left(\alpha^{-1}p(\theta)\right)^{-\frac{1}{1-\alpha}} F\left(L,K,H,\theta\right)$. This implies a constant markup for the monopolist. Let us also normalize the marginal cost of machine production to $\gamma = \alpha^2$, so that the monopoly markup is $1/\alpha$, and the monopoly price is $p(\theta) = \alpha$. Substituting this into the demand for machines, we obtain

$$q(\theta) = F\left(L,K,H,\theta\right). \tag{4}$$

This implies a simple form for GDP: $Y = F\left(L,K,H,\theta\right)$.

Next, consider the profit-maximization problem of the technology monopolist, which can be written as $\max_{\theta\in[0,1]} \Pi(\theta) = \left(p(\theta)-\gamma\right)q(\theta) - \Gamma\theta$. Using (4) and $p(\theta)=\alpha$, this can be simplified to $\max_{\theta\in[0,1]} \Pi(\theta) = \alpha(1-\alpha)\,F\left(L,K,H,\theta\right) - \Gamma\theta$. Using the linearity of $\Gamma(\theta)$, θ^* satisfies the first-order condition:

The equilibrium choice of technology is then characterized by the monopolist's first order condition

$$(1-\alpha)\ \alpha\,\frac{\partial F(L,\theta^*)}{\partial\theta} = \Gamma, \tag{5}$$

where we suppress dependence on K and H for notational simplicity.

Given the monopoly markup, this equilibrium is inefficient. The following lemma shows that this inefficiency takes the form of insufficient technological investment:

Lemma 1 *An increase in θ starting from the equilibrium θ^* increases GDP and net output.*

Intuitively, as in other models of endogenous technology choice, the technology monopolist cannot appropriate the full returns from its technology choice, and as a result chooses a lower than efficient level of technology.

We now turn to the characterization of when labor scarcity induces technological investments. The following lemma shows this is the case, whenever labor L and technology θ are substitutes—or using the terminology in Acemoglu (2010), when technology is "(strongly) labor-saving." To state this result in the simplest form, let us define $\varepsilon_{L\theta} \equiv -(\theta\partial^2 Y/\partial L\partial\theta)/\partial Y/\partial L$ as the elasticity of the

used for task production. The latter type of capital is held fixed for convenience, and we discuss how it can be endogenized later in this section. One could alternatively treat the machines embedding technology θ as intermediate goods, in which case GDP would be $(1-\alpha)\,Y$, and all of the results derived here would still apply.

marginal product of labor with respect to technology. This measures how the marginal product of L changes with θ. In particular, $\varepsilon_{L\theta} > 0$ if and only if $\partial^2 F/\partial\theta\partial L < 0$, or equivalently if and only if technology is labor-saving at the margin.

Lemma 2 *Labor scarcity changes technology in the direction of increasing GDP (or equivalently $d\theta^*/dL < 0$) whenever technology is labor-saving at the margin or equivalently $\varepsilon_{L\theta} > 0$.*

In Appendix B, we show that this condition is satisfied in task-based models of automation (Autor et al., 2003; Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018, 2022a,b). Intuitively, automation lets capital take over tasks previously performed by younger workers, lowering their marginal product; the resulting productivity gains accrue instead to capital and to older, skilled labor, which are complementary to automation.

Although labor scarcity—meaning in particular younger worker scarcity—may increase technology in the direction of increasing GDP, the total impact on GDP (value added) also includes the direct negative impact of reduced labor supply from fewer younger workers. Mathematically:

$$\begin{aligned}
\frac{d\log Y}{d\log L} &= \frac{\partial \log Y}{\partial \log L} + \frac{\partial \log Y}{\partial \theta}\frac{d\theta}{d\log L} \\
&= \left[\frac{\partial Y}{\partial L} + \frac{\partial Y}{\partial \theta}\frac{d\theta}{dL}\right]\frac{L}{Y} \\
&= \left[1 - \frac{\frac{\partial Y}{\partial \theta}}{\frac{\partial^2 Y}{\partial \theta^2}\theta}\frac{\frac{\partial^2 Y}{\partial L\partial\theta}\theta}{\frac{\partial Y}{\partial L}}\right]\frac{w_L L}{Y} \\
&= \left[1 - \frac{\varepsilon_{L\theta}}{\varepsilon_{\theta\theta}}\right]s_L.
\end{aligned}$$

Here, the third line divides through by $\partial Y/\partial L = w_L$, substitutes for $d\theta/dL$ using the implicit function theorem (see equation (B-2) in the Appendix), and divides and multiplies through by θ. The elasticity $\varepsilon_{\theta\theta} \equiv -(\theta\partial^2 Y/\partial\theta^2)/\partial Y/\partial\theta > 0$ measures how rapidly the marginal product of θ declines with θ, while the elasticity $\varepsilon_{L\theta}$ was defined above and is informative on how rapidly the marginal product of θ declines with L. The total impact of labor scarcity from falling birth rates on GDP is positive if and only if $\varepsilon_{L\theta} > \varepsilon_{\theta\theta}$:

Proposition 3 *Labor scarcity increases GDP if and only if $\varepsilon_{L\theta} > \varepsilon_{\theta\theta}$.*

Labor scarcity can thus raise GDP even when its direct effect is to reduce production: the indirect technology response dominates when the marginal product of (younger) labor is more sensitive to automation than is the marginal product of θ itself.

3.3 The Response of Capital

The preceding analysis held the supply of capital used for task production, K, fixed. Following Acemoglu and Restrepo (2022b), it is straightforward to endogenize K by assuming that one unit of final good produces χ units of capital. In this case, and in contrast to the neoclassical growth model, the supply of capital also increases in response to labor scarcity.

The mechanism is that labor scarcity, by spurring automation, raises the demand for and return to capital, encouraging more capital to be produced. Both the technology and capital supply responses thus create indirect effects that oppose the direct impact of labor scarcity. Since the technology effect alone can overturn the direct impact, the same is *a fortiori* true when capital also adjusts.

We will see that, in the data, both technology and capital respond in the direction predicted by this framework, though the technology response is more pronounced and important.

4 Deriving the Empirical Specification

The preceding section showed that labor scarcity can raise output through induced automation. This section formalizes how demographic dynamics—specifically population size and age composition—impact labor scarcity. It also elucidates why birth rates are the appropriate source of identifying variation in our empirical work.

To clarify how birth rates shape population dynamics and, in turn, GDP, we start with a setting in which only population matters, so the economic relationship of interest is

$$\Delta_h y_{c,t+h} = \text{time effect}^h_{t+h} + \alpha^h \, \Delta_h \log P_{c,t+h} + \rho^h \, y_{c,t} + \pi^h \log P_{c,t} + \eta^h_{c,t+h}, \tag{6}$$

where $y_{c,t}$ denotes log GDP per worker in country c at time period t, h denotes the horizon on which we focus, $P_{c,t}$ is working-age population ("workers") in country c at time t, and $\eta^h_{c,t+h}$ is an error term capturing all omitted influences, again for country c, time t and horizon h. The operator Δ_h is defined as the h period difference $\Delta_h x_{c,t+h} = x_{c,t+h} - x_{c,t}$. The presence of $y_{c,t}$ on the right-hand side allows for mean-reverting dynamics and ensures that we do not (inadvertently) impose permanent growth differences between countries (in particular, there will be permanent differences in growth rates across countries if $\rho^h = 0$). Including $\log P_{c,t}$ on the right-hand side, in turn, allows more general error-correction dynamics to be incorporated, so that past population levels can also affect growth.

A key challenge in estimating this equation is the possible endogeneity of population changes. The population structure is determined by three components: births, deaths, and migration. Both death rates and net migration could respond to contemporaneous income levels, as higher-income areas tend to have better healthcare and also attract more migrants. Entry into the workforce, on the other hand, is determined by past birth rates, which depend on past income and demographic variables and are also shaped by social and cultural factors. This observation implies that, conditional on past incomes, past birth rates can be taken as (weakly) exogenous.[9]

[9] In Appendix B we show more formally that this assumption is satisfied even when parents make fertility choices based on expectations of future incomes, provided that productivity follows a Markov process and parents use past and current variables to form their expectations—rather than having more information about the future than is

The dynamics of working-age population $P_{c,t}$ over an interval of length h can be written as:

$$P_{c,t+h} - P_{c,t} = \text{surviving births}^h_{c,t-n} - \text{net exits}^h_{c,t+h},$$

where surviving $\text{births}^h_{c,t-n}$ denotes births from n years ago that survive to period $t+h$, given that individuals enter the working-age population at age n. Net $\text{exits}^h_{c,t+h}$ incorporate both deaths, transitions out of the working-age group, and net migration between dates t and $t+h$. Dividing both sides of this equation by $P_{c,t}$ and using the approximations: $(x_{t+h} - x_t)/x_t \simeq \Delta_h \log x_{t+h}$, surviving $\text{births}^h_{c,t-n}/P_{c,t} \simeq \text{birth rate}_{c,t-n}$, we obtain

$$\Delta_h \log P_{c,t+h} = \text{birth rate}_{c,t-n} - \mathsf{e}^h_{c,t+h}. \tag{7}$$

where $\mathsf{e}^h_{c,t+h} = \text{exits}^h_{c,t+h}/P_{c,t}$ is the net exit rate, inclusive of deaths and migration. This term could be correlated with determinants of income and thus with the error term $\eta^h_{c,t+h}$ (for example, because higher incomes could attract migrants or improve healthcare, reducing mortality). For this reason, we cannot identify α^h in equation (6) using variation in population growth, $\Delta_h \log P_{c,t+h}$, though variation in birth rates is orthogonal to $\eta^h_{c,t+h}$. Combining equations (6) and (7), we obtain:

$$\Delta_h y_{c,t+h} = \text{time effect}^h_{t+h} + \alpha^h \text{ birth rate}_{c,t-n} + \rho^h y_{c,t} + \pi^h \log P_{c,t} + \varepsilon^h_{c,t+h}. \tag{8}$$

where $\varepsilon^h_{c,t+h} = \eta^h_{c,t+h} - \alpha^h \mathsf{e}^h_{c,t+h}$ is a composite error term including endogenous death rates and migration $\mathsf{e}^h_{c,t+h}$. Additionally, time effect^h_{t+h} absorbs the common components of population and income growth across countries. This equation shows that the link between growth in GDP per worker and population growth directly translates into a relationship between growth in GDP per worker and *the flow* of past births. This observation determines the form of our empirical specification, which we estimate for multiple values of t and h. Throughout we impose $n = 20$, assuming workers enter the labor force at age 20.

Our main specification separates the working-age population into younger and older workers, following Acemoglu and Restrepo (2017, 2022a) who documented the importance of age composition for economic growth. To construct two broadly equal-sized groups, we define the younger-age group to be between 20 and 45 and the older-age group to be between 45 and 70, and we allow these two groups to have separate effects on GDP dynamics. In Appendix B, we show that in this case equation (8) generalizes to

$$\Delta_h y_{c,t+h} = \text{time effect}^h_{t+h} + (\beta^h_L + \beta^h_H) \text{ birth rate}_{c,t-n} + \rho^h y_{c,t} + \pi^h \log L_{c,t} + \mu^h \log H_{c,t} + \eta^h_{c,t+h}, \tag{9}$$

where $L_{c,t}$ is younger population in country c at time t and $H_{c,t}$ denotes older population, and β^h_L and β^h_H are the separate effects of younger- and older-population growth on GDP per worker (derived in Appendix B.1). Because a single birth-rate cohort feeds both age groups over the horizon, the

available in the data. Specifically, the key assumption is that parents *20 years ago* would have no more information about future growth than an econometrician who observes income and population *today*.

two effects enter only as their sum, $\beta^h = \beta^h_L + \beta^h_H$. Their initial levels on the right-hand side allow for error-correction dynamics.

At first glance, our use of birth rates as exogenous variation suggests an instrumental variable (IV) strategy. Although birth rates are argued to be exogenous, a full IV approach is not feasible because economic outcomes depend on both younger and older workforces, and we have only a single instrument. While this means that we cannot separately identify β^h_L and β^h_H in equation (9) using variation in birth rates alone, we can estimate their sum β^h as the reduced-form response of GDP per worker growth to past birth rates. In Section 9 we address this issue by separately identifying the impact of population size and population aging using exogenous variation in deaths from WWII.

5 Data Sources and Measurement

Our analysis draws on country-level and US commuting-zone data.

5.1 Country-Level Measures

Our country sample includes all nations with population exceeding one million in 2019, excluding tax havens (Ireland, Singapore, Luxembourg, and Hong Kong) and crisis-affected states (Yemen, Syria, and Venezuela). The sample period runs from 1950 to 2020. To limit the influence of the COVID-19 pandemic, we impute t_0-to-2020 log GDP growth by linear extrapolation of t_0-to-2019 log GDP growth, adding one year of average growth.

Our primary income measure is GDP per working-age adult (ages 20–70) or simply "per worker." We use aggregate real GDP, in constant prices, converted to 2017 US dollars from version 10 of the Penn World Tables (variable rgdpna) (Feenstra et al., 2015), available from 1950 to 2019. Demographic data come from the 2024 UN Population Prospects (United Nations, Department of Economic and Social Affairs, Population Division, 2024), which reports population by age and births from 1950 onward. We measure fertility using the (crude) birth rate, defined as births per 100 people.

Capital stock estimates are taken from the Penn World Tables, and we construct TFP via the growth accounting formula: $\Delta \text{TFP}_{c,t+1} = \Delta \log \text{GDP}_{c,t+1} - (1 - \text{labor share}_{c,t})\Delta \log \text{Capital}_{c,t+1} - \text{labor share}_{c,t}\Delta \log \text{Labor}_{c,t+1}$, where the labor share and employment data are from the Penn World Tables. The h-period change in TFP is the sum of one-year changes over a horizon of h years.

5.2 Commuting Zone Data

Within the United States, our unit of analysis is the commuting zone, a consistently defined local labor market area where most residents both live and work. We use the 722 commuting zones covering the contiguous United States constructed by Tolbert and Sizer (1996), based on journey-to-work data from the 1990 Census. Our primary sources for US demographic data are the US Census of Population and the American Community Survey (ACS). We obtain demographic structure, employment and aggregate earnings from the Census Survey Tabulation Files (STF) accessed via Schroeder et al. (2025). These files are constructed from the universe of Census long-form responses, aggregated to the county-by-decade level. Births come from the US Vital Statistics.[10] Because the STF data are discontinued after 2000, for the years 2010 and 2020 we use five-year ACS averages centered on 2010 and 2020 (averaging 2008–2012 in the former case and averaging 2018–2022 in the latter case). As detailed below, we calculate composition-adjusted log weekly wages using Census Integrated Public Use Microdata Series (IPUMS) samples for each decade between 1960 and 2020 (Ruggles et al., 2025).

Our key demographic controls are the population aged 20 to 45 ("younger workers") and the population aged 45 to 70 ("older workers"). The STF report population only for the 65–75 age group, so we impute the 65–70 subgroup as follows: from 1970 to 2010, we apply the shares available from the NBER intercensal population estimates; before 1970 and after 2010, we impute the 65–70 share of the 65–75 bin using the corresponding proportion of the cohort that was aged 55–60 ten years earlier, implicitly assuming uniform mortality within the 55–65 age band. When studying the response of employment by age group, we consider employment among younger workers (aged 20 to 45) and among older workers (aged 45 or older). In the STF we only see total employment for people aged over 65. Since the bulk of workers over age 65 are between 65 and 70, we measure employment of older workers as all employees aged 45 and older.

Our baseline measure of wages across commuting zones is the composition-adjusted weekly wage, constructed as follows. We define w_{igt} as the average wage of group g in commuting zone i at time t, and let s_{igt} be the population share of that group.[11] The actual wage series for a commuting zone is the population-weighted average of group-specific wages, $w_{it}^{\text{raw}} = \sum_g s_{igt} w_{igt}$. Our composition-adjusted measure uses group weights in the base year of 1990, so $w_{it}^{\text{adj}} = \sum_g s_{ig,1990} w_{igt}$. Demographic cells corresponding to $s_{ig,1990}$ are defined by age group (25–34, 35–44, 45–54, 55–64, and 65+), gender (male, female), race (white or non-white), and education (less than high school, high school, some college, college, and post-college). We also use average log annual wages among workers observed in the IPUMS samples (or composition-adjusted versions thereof) either as a control or as an outcome variable, analogous to overall GDP in the cross-country sample.

5.3 Export and Employment Shares

In our analysis of mechanisms, we look at sectoral shifts in response to birth rates across countries and commuting zones. Our first measure is the high-tech share of exports across countries and high-tech share of employment across commuting zones. Our data on exports by industry are

[10] We access Vital Statistics and Census tabulation files via the NHGIS (Schroeder et al., 2025). The 1970 birth rates are derived from individual-level microdata, which is not available for the five counties within New York City. For these five counties we compute birth rates using the National Center for Health Statistics (NCHS) estimates, which are also available from NHGIS.

[11] For 1960, 1970, and 2010, the number of weeks worked is available in six discrete intervals. We set the weeks worked equal to the midpoint in each interval, with the exception of the 1–13 weeks bin, which we set to 10 weeks worked.

taken from the UN Comtrade database, which we access via the World Integrated Trade Solutions (WITS) platform. These data are available from 1990 forward for 26 countries, with additional countries covered beginning in 1995. We classify high-tech exports following the Eurostat high-tech industry classification, which includes pharmaceuticals, computers and electronics, chemicals, electrical equipment, machinery, and transport equipment.[12] For commuting zones, we use the cleaned County Business Patterns (CBP) dataset, prepared by Eckert et al. (2021). These data provide employment counts by NAICS industry for US counties, which we aggregate to commuting zones. Because county-by-industry cells in the CBP are heavily censored after 2016, we assume industry structure is constant between 2016 and 2020. We classify industries as high-tech using an analogous definition, but now based on NAICS codes.[13]

Second, for our commuting-zone sample, we classify US industries as labor-intensive or capital-intensive using the Bureau of Economic Analysis (BEA) KLEMS accounts, which report the compensation shares of capital and labor at the national level for industries since 1987. We define an industry as labor-intensive if its labor share of value added is above the value-added-weighted median across all industries. Third, we define an analogous measure of R&D-intensive industries, based on the share of value added paid to R&D capital, as recorded by the BEA.

We also compute manufacturing and agricultural employment across both countries and commuting zones for our study of alternative mechanisms. Manufacturing employment across countries comes from the International Labor Organization (ILO), which provides data for only 46 countries in 1990. For commuting zones, we measure manufacturing employment using the CBP. Because this data source does not cover agriculture, we obtain the agricultural employment share from Census Public Use Microdata Samples.

5.4 Industry Data

Because TFP and capital data are not available at the industry-by-commuting-zone level, we use national industry-level data instead. Our industry TFP data come from the BEA KLEMS accounts, which contain TFP measures alongside factor shares and quantity indices for 60 industries. These factors include five distinct capital classes: intellectual property (entertainment-related, distinct from R&D), R&D, IT, software, and other capital. We aggregate these capital classes into a single capital stock for each industry j using a Divisia index, given by $\Delta \ln k_{j,t} = \sum_g s_{j,g,t-1} \Delta \ln k_{j,g,t}$, where the $s_{j,g,t}$'s are factor shares of the different types of capital (summing to one) and the $k_{j,g,t}$'s denote quantity indices of capital type g for industry j. The h-period change in industry capital is obtained as the sum of one-year changes.

[12] This classification is also used by Goos et al. (2015). The precise classification is available at https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Glossary:High-tech_classification_of_manufacturing_industries

[13] High-tech industries are: chemicals (NAICS 325), machinery (NAICS 333), computers and electronics (NAICS 334), electrical equipment (NAICS 335), and transportation equipment (NAICS 336).

5.5 Patent Data

We use patent data to construct a more direct measure of labor-saving innovation. We obtain data from the universe of patents granted by the United States Patent and Trademark Office (USPTO) and assign patents to countries based on the country of the patent assignee, which is typically the firm that owns the patent. For commuting zones, we assign patents to the city of the listed assignee and map cities to counties using the crosswalk provided by the PatentsView database.[14] We then aggregate counties to the commuting-zone level. We consider granted patents, which are assigned to their year of application. To correct for lags in patent grants and citations, we consider patent applications through 2015. We weight patents by citations received in the five years after application, as is standard in the literature (Hall et al., 2001). To reduce noise, we compute patenting for ten-year periods centered on decennial census years, considering only patents granted in 1975 or later, making 1980 our first observation, and 2010 our last observation (which includes patents with application date between 2005 and 2015).

We classify patents using three measures of labor-saving technology. First, we assign patents to NAICS industries via their International Patent Classification (IPC) codes using the crosswalk of Lybbert and Zolas (2014), and define a patent as labor-intensive if its industry has a labor share of value added above the value-added-weighted median for all US industries. Second, we classify patents as ICT based directly on IPC codes, using subcategories G05–G16 (computing, controlling, and information storage devices).[15] Third, we classify patents as automation-related following Autor et al. (2024), who compute a similarity score between each patent and occupation task descriptions in the BLS's *Dictionary of Occupational Titles*, classifying a patent as automating for an occupation if its similarity score is in the top 15% of all patent-occupation pairs in a given decade. We define a patent as an automation technology if the number of occupations it automates is in the top decile among all patents in our sample. This implies that 10% of patents in our sample are classified as automation patents.

5.6 World War II Casualties

We consider the impact of World War II casualties on subsequent economic outcomes, using GDP per capita from the Maddison Project database (Bolt and van Zanden, 2025). Our data on World War II casualties come from Kesternich et al. (2014), which we extend using counts provided by the Encyclopedia Britannica and Wikipedia.[16] Our data on early 20th century population by age group come from the International Historical Statistics database (Mitchell, 2013). These data are

[14] Available at `https://s3.amazonaws.com/data.patentsview.org/download/g_location_disambiguated.tsv.zip`

[15] Specifically, we define all subcategories from G05–G16 as ICT. These are: G05: controlling; regulating, G06: computing or calculating, G07: checking-devices, G08: signaling, G09: educating; cryptography; display; advertising; seals, G10: musical instruments; acoustics, G11: information storage, G12: instrument details, G16: information and communication technology specifically adapted for specific application fields.

[16] The precise numbers of WWII deaths are unknown. Encyclopedia Britannica and Wikipedia compile statistics from multiple sources into a single harmonized dataset. The data are available at: `https://www.britannica.com/event/World-War-II/Costs-of-the-war` and `https://en.wikipedia.org/wiki/World_War_II_casualties`.

available at irregular intervals across countries, so we impute the 1935 share of younger workers using linear interpolation where necessary. The share of young workers is defined as the ratio of the population aged between 20 and 45 to total population over 20 (a broader denominator than elsewhere, dictated by the limited age detail in pre-1950 historical sources). For this analysis, we aggregate former Soviet Union countries by taking the population-weighted average of GDP per capita from the Maddison Project database.

5.7 Other Data Sources

We use several additional measures as controls and/or in our analysis of alternative mechanisms. For education, we use average years of schooling from Barro and Lee (2013) at the country level and the college-attainment share from the Census Tabulation Files for commuting zones. When education serves as a control, wherever possible we measure it among the older cohorts—among those aged 45–65—to mitigate concerns about endogeneity of education to birth rates and cohort sizes.[17] When education is used as an outcome variable, we want to capture *the flow* of education, and hence we focus on education among those aged 25–45. Data on female labor force participation come from the ILO for countries and from the Census Tabulation Files for commuting zones. Urbanization rates come from the World Bank for countries. For commuting zones, we proxy urbanization with log population density.

6 Birth Rates and Economic Growth: Evidence

This section presents our main results, first at the country level and then using more detailed commuting-zone data. Throughout, we estimate the following equations based on (9):

$$\Delta_h y_{c,t+h} = \alpha_t^h + \beta^h \text{birth rate}_{c,t-n} + \boldsymbol{x}'_{c,t}\boldsymbol{\gamma}^h + \varepsilon_{c,t+h}^h, \tag{10}$$

where $y_{c,t}$ is log GDP per worker in country c, or the composition-adjusted average wage in commuting zone c; birth rate$_{c,t-n}$ is the crude birth rate in year $t-n$; and $\boldsymbol{x}_{c,t}$ is a vector of controls that includes $y_{c,t}$ (log GDP or log average wage) and initial labor supply, measured as the log numbers of younger (aged 20–45) and older (aged 45–70) workers. Our main cross-country specifications are unweighted (so the results are not overwhelmingly driven by larger countries), while our main US commuting-zone and industry regressions weight observations by working-age population in year t. We also report population-weighted specifications for countries and unweighted specifications for commuting zones as robustness checks.

Our key identifying assumption is that, conditional on year t, average incomes and population structure $\boldsymbol{x}_{c,t}$, the 20-year lagged birth rate, $b_{c,t-20}$, is orthogonal to unobserved determinants of

[17] Education by age is available for commuting zones only after 1970, and only for the age groups 25–45 and 45–65. To avoid dropping data for 1960, we use education among all adults aged 25 or older as a control for our commuting-zone analysis.

income growth between t and $t + h$ (for example, in our long-differences specification, birth rates in 1950 are orthogonal to omitted influences on GDP or wage growth between 1970 and 2020). We view this as a compelling exclusion restriction, as explained in Section 4 and Appendix B.

We use three estimation strategies. Our first, long-differences, selects a single value of both t and h, then estimates equation (10) using differences between dates t and $t + h$. For countries we take $t = 1970$, $h = 50$, and $n = 20$. Hence, the outcome year is 2020 and the key right-hand side variable is the birth rate in 1950, which exhibits considerable variation across countries.

Our second strategy, stacked-differences, uses two shorter intervals: 25-year intervals across countries and 30-year intervals across commuting zones. The stacked-differences specification is useful because it captures variation in birth rates in 1950 and then again in 1975 (1940 and 1970 for commuting zones), the former corresponding to the midpoint of the baby boom era and the latter representing the early stages of the subsequent declines in birth rates for many countries. On the other hand, this specification imposes linearity and symmetry in dynamics, which may not be warranted (and are not imposed by our third approach).

Our final strategy follows Jordà (2005) and uses the local projection method and estimates equation (10) for all possible h using all time t periods between 1970 and 1990. This specification exploits the full variation in birth rates across countries and commuting zones, while also enabling us to inspect the time pattern of responses, summarized by $\left\{\beta^h\right\}_h$. Additionally, as in standard event studies, we explore whether there is a response in GDP or wages before the variation in birth rates reaches working-age.[18]

6.1 Country-Level Results

Table 1 presents country-level results for GDP per worker. In all specifications initial (1970) log GDP per worker, log younger population (aged 20–45), and log older population (aged 45–70) are included as controls. Panel A reports (unweighted) long-difference estimates of equation (10), where the outcome is the 50-year change in log GDP per worker between 1970 and 2020. Panel B reports stacked-difference estimates with two 25-year changes: 1970–1995 and 1995–2020.[19] In Panel A, standard errors are robust to heteroskedasticity. In Panel B, we cluster standard errors at the country level to account for within-country serial correlation. Panels C and D report identical models but using country population in the year t as weights.

Column 1 of Panel A reports the baseline estimate from equation (10). The coefficient of -0.23 (standard error $= 0.11$) indicates that a one percentage-point lower birth rate in 1950 is associated with 23 log points higher GDP per worker over the subsequent 50-year period.[20] Because initial GDP per worker on the right-hand side has a negative coefficient, this estimate can be interpreted

[18] See Dube et al. (2025) for the correspondence between local projection estimates and event studies.

[19] Our cross-country dataset does not allow us to work with two 30-year periods.

[20] This translates into 0.46 percentage-points faster annual growth between 1970 and 2020, which is large but not implausible. For example, the US had a 1.6 percentage-point lower birth rate in 1950 than Mexico. Our estimate implies that this difference explains 74% of the one percentage-point difference in annual growth rates between the US and Mexico between 1970 and 2020.

Table 1: Effect of Birth Rates on Growth in GDP per Working-Age Adult (Countries)

	(1)) Base	(2) Educ	(3) Urban	(4) Region	(5) All	(6) FE
	Panel A: Long Differences					
birth rate (t-20)	-0.23	-0.17	-0.21	-0.17	-0.04	
	(0.11)	(0.11)	(0.11)	(0.12)	(0.12)	
N	107	107	107	107	107	
R-sq	0.494	0.518	0.506	0.575	0.622	
	Panel B: Stacked Differences					
birth rate (t-20)	-0.22	-0.18	-0.21	-0.15	-0.11	-0.09
	(0.05)	(0.05)	(0.05)	(0.06)	(0.06)	(0.08)
N	229	229	229	229	229	214
R-sq	0.406	0.415	0.408	0.443	0.460	0.838
	Panel C: Long Differences - weighted					
birth rate (t-20)	-0.36	-0.17	-0.36	-0.30	-0.04	
	(0.15)	(0.11)	(0.15)	(0.19)	(0.12)	
N	107	107	107	107	107	
R-sq	0.751	0.518	0.752	0.768	0.622	
	Panel D: Stacked Differences - weighted					
birth rate (t-20)	-0.34	-0.35	-0.34	-0.32	-0.32	-0.20
	(0.05)	(0.05)	(0.05)	(0.05)	(0.06)	(0.16)
N	229	229	229	229	229	214
R-sq	0.642	0.642	0.642	0.656	0.658	0.841

Notes: This table reports estimates of β^h in equation (10), estimated in our cross-country dataset. The dependent variable is $\Delta_h y_{c,t+h}$, the change in log GDP per worker over horizon h. In each panel we restrict the sample to be consistent across specifications. Panels A and C set $h = 50$ and restrict the sample to be a single cross section, where $t = 1970$. Panels B and D set $h = 25$ and restrict the sample to two non-overlapping cross sections, where $t \in \{1970, 1995\}$. Panels A and B weight countries equally, while Panels C and D weight countries by working-age population in year t. The vector $\boldsymbol{x}_{c,t}$ includes initial log GDP per working-age adult and labor supply in t, measured by log population aged 20–45 and log population aged 45–70. Column 2 additionally controls for average years of schooling among the population aged 45–55 in year $t-20$. Column 3 controls for urbanization, measured by the share of people residing in urban areas in t. Column 4 controls for continent fixed effects, while Column 5 adds all controls simultaneously. Column 6 (Panel B only) includes country fixed effects, not identified in the long-difference specifications. Standard errors, clustered by country, are reported in parentheses.

as the impact of the 1950 birth rate on the 2020 *level* of log GDP per worker. Column 1 of Panel B, which uses stacked 25-year differences, yields a coefficient of -0.22 (standard error = 0.05). This estimate is similar in magnitude to the long-difference estimate in Panel A.[21]

[21] The coefficient of -0.22 measures the impact on growth of GDP per worker conditioning on initial log GDP per worker, which itself has a coefficient of approximately -0.3. This implies a more muted impact from birth rates in 1950 on GDP per worker growth between 1995 and 2020 than on GDP per worker growth between 1970 and 1995.

Columns 2 through 5 assess the robustness of these estimates to alternative controls. Column 2 adds education in year $t-20$ to test whether human capital trends confound our results. We use education among older (aged 45–55) workers in year $t-20$ to avoid the endogeneity of education to birth rates or cohort size. Controlling for education yields comparable point estimates across panels: -0.17 (standard error $= 0.11$) in Panel A and -0.18 (standard error $= 0.05$) in Panel B. Column 3 adds urbanization, and Column 4 adds a full set of region dummies, allowing, respectively, for differential trends by urbanization and across major world regions. The estimated coefficients remain stable across columns, except in Column 5, where we lose identification for the long-difference model when all controls are added simultaneously.

Country fixed effects are not identified in the long-differences specifications. Column 6 of Panel B adds country fixed effects to the stacked-difference models, which is equivalent to allowing linear country trends. Identification now comes from within-country variation in birth rates and growth over time. The point estimate remains negative but is attenuated and is no longer significant. Given the difficulty of identifying country-specific trends from only two observations per country, the lack of significance is unsurprising. We show below that the same specification in the commuting-zone analysis suffers much less attenuation.

Panels C and D document that the results are broadly similar in population-weighted specifications, though the magnitudes are somewhat larger. In the remainder of the paper, we focus on unweighted cross-country specifications.

The stacked-differences models exploit cross-country variation in birth rates in 1950 (for the 1970–1995 period) and 1975 (for 1995–2020). Appendix Table A1 examines whether baby-boom-era birth-rate variation has a different effect than baby-bust-era variation, by reporting the impact of birth rates in $t - 20 = 1950$, 1960, 1970, and 1980 on GDP per worker growth between t and $t + 30$, so that, for example, column 3 of Panel A of Table A1 regresses the 30-year growth of GDP per worker between 1980 and 2010 on the 1960 birth rate. The results show that the effects of birth-rate differences across countries are quite similar between the baby boom and the baby bust eras.

In Appendix Table A2, we extend our data to 1930 to show that adding birth rates from both 1930 and 1950 leads to similar results across specifications. Appendix Table A3 additionally confirms that our cross-country results are robust to using log birth rates rather than levels as the regressor. We focus our analysis on birth rates in levels, since this is the specification derived from population dynamics in Section 4.

Figure 3 presents local-projection impulse responses. We estimate equation (10) separately for each horizon h using all available (and potentially overlapping) time periods t, with standard errors corrected for overlapping observations. The coefficients $\{\beta^h\}_h$ trace out how GDP per worker evolves following a one percentage-point increase in birth rates at $t-20$. These estimates align

As our local-projection estimates show, the time pattern of effects differs from what this linear model imposes. Because of the linearity imposed by the stacked-differences model, we place more emphasis on the quantitative magnitudes implied by the long-differences and local-projection estimates.

Figure 3: Cross-Country Impulse Response of Income to Birth Rates in $t{-}20$

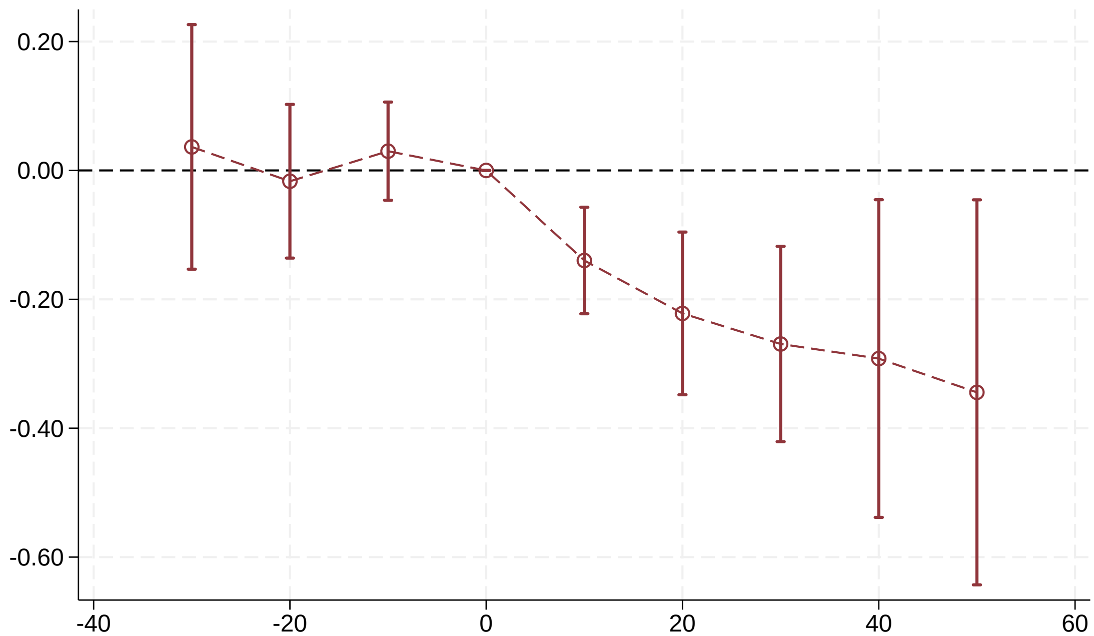

Notes: This figure reports estimates of β^h in equation (10) alongside 95% confidence intervals, estimated in our cross-country dataset. The dependent variable is $\Delta_h y_{c,t+h}$, the change in log GDP per worker over horizon h. The vector of controls $\boldsymbol{x}_{c,t}$ includes initial log GDP per worker and labor supply in year t, measured by log population aged 20–45 and log population aged 45–70. To keep the sample consistent across horizons, we exclude any countries with missing GDP data in 1960. All regressions weight countries equally. Standard errors are clustered by country.

closely with both Table 1 and theoretical expectations. Since cohorts born at $t{-}20$ (e.g., 1970) do not enter the working-age population until $t = 0$ (e.g., 1990), absent confounding effects, we should observe no response before $t = 0$. In lineit's with this expectation, the estimated coefficients are close to zero and statistically insignificant before $t = 0$. Beginning at $t = 0$, however, higher birth rates are associated with lower GDP per worker, and the negative effects grow progressively larger over time. The magnitudes are broadly consistent with Panel A of Table 1. For example, a one percentage-point lower birth rate at $t{-}20$ is associated with 22 log points higher GDP per worker at $t{+}20$ and 29 log points higher by $t{+}40$. The latter corresponds to roughly 0.73 percentage points faster annual growth over the 40-year interval.

In summary, cross-country estimates show that lower birth rates robustly predict subsequent growth in GDP per worker. The estimated effects are large, but are not implausible, as we explain next.

6.2 Population Responses and Economic Magnitudes

Birth rates affect GDP per worker via their impact on both working-age population and age composition. Figure 4 presents our local-projection impulse response estimates on overall working-age population, younger population, and older population (all in logs). In Panel A, we see that, reassuringly, log working-age population shows no movement before time t and begins to decline thereafter, falling by about 25 log points after 70 years for a one percentage-point lower birth rate. Panel B shows that the younger population, which dominates the working-age aggregate over this horizon, follows the same pattern: flat before t, then declining. Panel C turns to the older population, where the pre-trends are somewhat noisier, but still indistinguishable from zero. Consistent with our expectations, the older population does not begin to decline until about 50 years after the birth-rate change, once the smaller cohorts reach ages 45–70. Therefore, lower birth rates reduce the working-age population and the younger population, and have a smaller and more delayed impact on the older population, confirming the empirical specification of Section 4.

Figure 4: Effect of Birth Rates on Working-Age Population and Aggregate GDP (Countries)

Notes: This figure reports estimates of β^h in equation (10) alongside 95% confidence intervals, estimated in our cross-country dataset. The dependent variable is $\Delta_h y_{c,t+h}$. Panel A reports estimates where $y_{c,t}$ is log population aged 20 to 70 years old, which we call working-age population. Panel B reports estimates where $y_{c,t}$ is the log population aged between 20 and 45 years old. Panel C reports estimates where $y_{c,t}$ is the log of population aged 45 to 70. Panel D reports estimates where $y_{c,t}$ is the real aggregate GDP. The vector of controls $\boldsymbol{x}_{c,t}$ includes initial log GDP per worker and labor supply in year t, measured by log population aged 20–45 and log population aged 45–70. All regressions weight countries equally. Standard errors are clustered by country.

Panel D of Figure 4 examines the relationship between birth rates and *aggregate* GDP (see also Appendix Table A4). It shows that there is no statistically or economically significant relationship between birth rates and aggregate GDP. This implies that the positive impact of lower birth rates on GDP per worker reflects a reduction in working-age population (the denominator) that is not accompanied by a decline in aggregate GDP (the numerator).

This finding structures the interpretation of our estimates. From equation (2), lower birth rates reduce the working-age population, which directly lowers GDP. For aggregate GDP to remain approximately unchanged, the induced responses of technology and capital must fully offset this direct negative effect. As we show below, lower birth rates raise the capital stock, but this capital deepening alone is insufficient to offset the direct effect of reduced labor supply. The remainder of the offset must come from endogenous technology responses, which we examine in the next section.

Proposition 3 provides a theoretical foundation for such a strong technology response: it shows that labor scarcity can raise GDP when the induced technology effect is sufficiently powerful. A simple calculation illustrates the required magnitude. We abstract from age-composition effects, so the direct impact operates entirely through a smaller working-age population. Panel A of Figure 4 implies a decline of approximately 16 log points in the working-age population at horizon $t+30$ (50 years after the birth-rate change). From equation (1) in Section 3.1, this translates to a direct GDP loss of $s_L \times 0.16 \approx 9.6$ log points, given a labor share of about 0.6. For aggregate GDP to remain unchanged, the induced technology and capital responses must therefore raise GDP by roughly 9.6 log points in the 30 years after the cohort enters the labor force. This is a large effect, but is consistent with our subsequent estimates.

6.3 Commuting-Zone Results

Table 2 presents analogous results for US commuting zones. These results follow the cross-country structure, with two additional control columns (initial manufacturing and agricultural employment shares) made possible by the richer US data. The principal gain from the commuting-zone analysis is that the dependent variable, composition-adjusted real weekly wages, nets out demographic shifts in age, gender, education, and race, isolating composition-constant wage changes. The richer US data also let us construct two 30-year stacked differences (Panel B), versus two 25-year periods cross-nationally. We begin with the baseline specification from the cross-country analysis, replacing initial GDP per worker with mean wages among employed workers. Our main specifications in Panels A and B now weight date t observations by commuting-zone working-age population in year t. This weighting accounts for heteroskedasticity arising from Census and ACS samples offering greater measurement precision in larger commuting zones. Panels C and D show the robustness of these results to unweighted specifications.

Column 1 in Panel A presents our 60-year baseline long-differences results. It shows that lower birth rates in 1940 are associated with statistically significant and sizable increases in wage growth between 1960 and 2020. The point estimate of -0.15 (standard error = 0.04) indicates that a

Table 2: Effect of Birth Rates on Wage Growth in Commuting Zones

	(1)) Base	(2) Educ	(3) Urban	(4) Region	(5) Manuf	(6) Agri	(7) All	(8) FE
Panel A: Long Differences								
birth rate (t-20)	-0.15	-0.12	-0.15	-0.13	-0.17	-0.14	-0.07	
	(0.04)	(0.03)	(0.04)	(0.03)	(0.04)	(0.04)	(0.02)	
N	722	722	722	722	722	722	722	
R-sq	0.166	0.432	0.168	0.452	0.345	0.204	0.694	
Panel B: Stacked Differences								
birth rate (t-20)	-0.11	-0.12	-0.11	-0.10	-0.13	-0.12	-0.08	-0.09
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)
N	1,438	1,438	1,438	1,438	1,438	1,438	1,438	1,432
R-sq	0.956	0.965	0.956	0.962	0.959	0.957	0.973	0.986
Panel C: Long Differences - unweighted								
birth rate (t-20)	-0.05	-0.04	-0.06	-0.04	-0.06	-0.04	-0.03	
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	
N	722	722	722	722	722	722	722	
R-sq	0.202	0.280	0.214	0.310	0.275	0.237	0.435	
Panel D: Stacked Differences - unweighted								
birth rate (t-20)	-0.11	-0.12	-0.11	-0.10	-0.13	-0.12	-0.08	-0.09
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)
N	1,438	1,438	1,438	1,438	1,438	1,438	1,438	1,432
R-sq	0.956	0.965	0.956	0.962	0.959	0.957	0.973	0.986

Notes: This table reports estimates of β^h in equation (10), estimated in our cross-commuting-zone dataset. The dependent variable is $\Delta_h y_{c,t+h}$, the change in log composition-adjusted weekly wage over horizon h. In each panel we restrict the sample to be consistent across specifications. Panels A and C set $h = 60$ and restrict the sample to be a single cross section, where $t = 1960$. Panels B and D set $h = 30$ and restrict the sample to two non-overlapping cross sections, where $t \in \{1960, 1990\}$. Panels A and B weight commuting zones by working-age population in year t. Panels C and D present unweighted estimates. The vector $\boldsymbol{x}_{c,t}$ includes log average wages and labor supply in year t, measured by log population aged 20–45 and log population aged 45–70. Column 2 controls for the share of adults with a college education in $t - 20$. Column 3 controls for urbanization, measured by initial population density in year t. Column 4 controls for Census Region fixed effects. Column 5 controls for the initial manufacturing share of employment in t, while Column 6 controls for the initial agricultural share of employment. Column 7 includes all of the aforementioned controls simultaneously. Column 8 includes commuting-zone fixed effects. Standard errors, clustered by commuting zone, are reported in parentheses.

one percentage-point lower birth rate in 1940 is associated with 15 log points higher (composition-adjusted real) wage growth over the subsequent 60 years, 1960–2020. This coefficient is substantially smaller in magnitude than the cross-country estimate of -0.23 over 50 years in Panel A of Table 1, despite spanning a longer (60- versus 50-year) horizon. As we show below, this difference is explained by the much smaller response of population to birth rates across commuting zones, which in turn can be explained by lower persistence in birth rates. The stacked-differences specification in Panel B yields a slightly smaller but more precisely estimated coefficient of -0.11 (standard error $= 0.02$). As in the cross-country analysis, this estimate is broadly consistent with the long-differences estimates in Panel A.

Columns 2 through 7 assess robustness. Column 2 adds initial education (the share of workers with a college degree or higher in $t-20$); Column 3 adds log population density; and Column 4 adds Census Region dummies to allow for differential trends across major regions. Columns 5 and 6 add the initial manufacturing and agricultural employment shares (controls unavailable at the country level) to absorb confounding from the secular decline of manufacturing- and agriculture-dependent local economies. Coefficients remain stable and statistically significant across all six specifications. Column 7 includes all controls simultaneously; the estimate attenuates to -0.07 (standard error $= 0.02$), but remains negative and precise.

Column 8 in Panel B includes commuting-zone fixed effects, allowing for differential linear trends across commuting zones between the two 30-year periods (analogous to Column 6 of Panel B in Table 1). Even in this demanding specification, the estimate remains negative and highly significant (-0.09, standard error $= 0.02$). This is comparable in magnitude to Column 7, which includes all controls.[22]

The unweighted results in Panels C and D are also precise, but the point estimates are smaller, reflecting the fact that the birth-rate–wage relationship is stronger among larger commuting zones, as we document in greater detail in Appendix A. For example, when we restrict the unweighted local-projection impulse responses to large commuting zones (population above 300 thousand in 1960), the estimates closely match the weighted results in Figure 5 (Appendix Figure A3). The wedge between weighted and unweighted estimates is therefore driven by smaller commuting zones. In the remainder of our commuting-zone analysis, we focus on the weighted specifications.

Panel B of Table A1 explores differential effects over time, looking at the impact of birth-rate differences across commuting zones on subsequent wage growth, separately for 1940, 1950, 1960, 1970, and 1980, once again showing that baby-boom-era and baby-bust-era variation in birth rates leads to qualitatively similar results.

As with the cross-country analysis, the results are also robust to using log birth rates instead of birth rates in levels, as shown in Appendix Table A5, and they are robust to using alternative earnings measures, as documented in Appendix Table A6.

Figure 5 presents the commuting-zone local-projection impulse responses, as in Figure 3. The results align with theoretical predictions and mirror the cross-country patterns. Reassuringly, birth rates in $t-20$ have no statistically significant effect on wages before t, but beginning in $t+10$, higher birth rates are associated with significantly lower wage growth. The magnitudes closely match the regression estimates in Table 2.

[22]Appendix Table A2 also shows that controlling for birth rates in 1930 leads to qualitatively similar results in the baseline commuting-zone specification, though the coefficient is again attenuated when all controls are included. The low precision in this specification partly reflects the fact that we have to drop approximately half of our sample to include these lags.

Figure 5: Commuting-Zone Impulse Response of Wages to Birth Rates in $t-20$

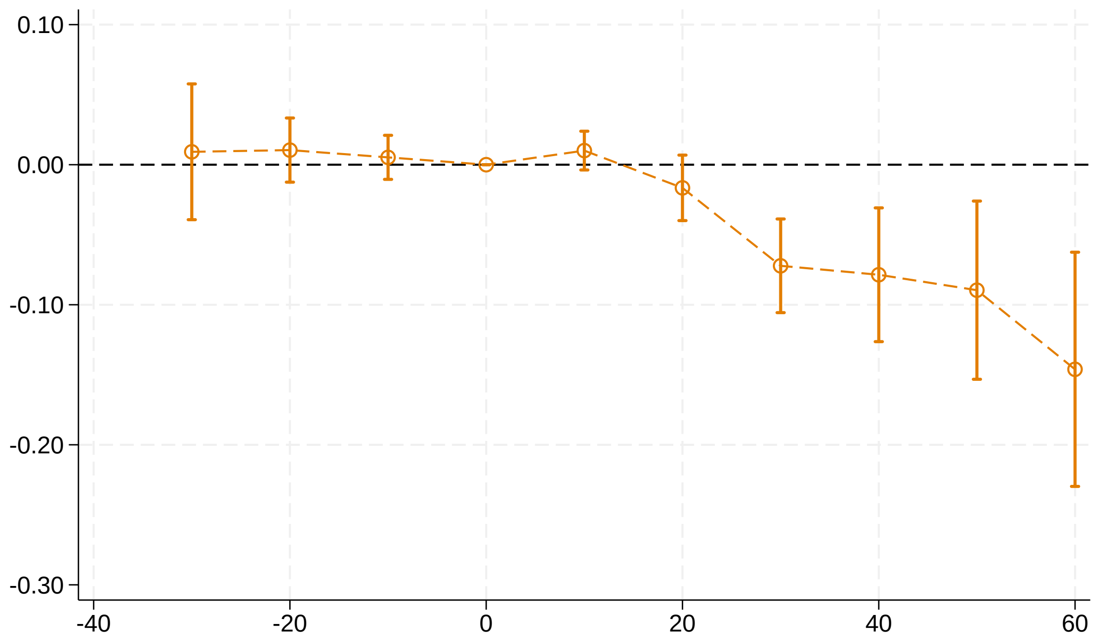

Notes: This figure reports estimates of β^h in equation (10) alongside 95% confidence intervals, estimated in our cross-commuting-zone dataset. The outcome $y_{c,t}$ is composition-adjusted log weekly wages, while the vector of controls $\boldsymbol{x}_{c,t}$ includes initial log average wages and labor supply in year t, measured by log population aged 20–45 and log population aged 45–70. To keep the sample consistent across horizons, we exclude any commuting zones with missing births data in 1970. All regressions weight commuting zones by working-age population in year t. Standard errors are clustered by commuting zone.

6.4 Interpreting Commuting-Zone Estimates

The estimated effect of birth rates on wage growth is considerably smaller for commuting zones than for countries. We now show that this gap reflects a much smaller population response to birth rates at the commuting-zone level.

Figure 6 is the analogue of Figure 4. Panels A through C report the local-projection impulse responses for the effect of 20-year-lagged birth rates on total, younger, and older employment.[23] While the estimates are positive, indicating that higher birth rates increase all three employment variables, they are generally small and imprecisely estimated. Total employment, for example, increases by about 7% at the 30-year horizon in response to a one percentage-point higher birth rate.

Panel D additionally shows the response of aggregate (composition-adjusted) earnings, which is once more small and statistically indistinguishable from zero, similarly to the aggregate GDP

[23] Because wages are observed only for working individuals, employment is the relevant labor-supply margin in the commuting-zone analysis; we rely on population at the country level because age-disaggregated employment is not consistently available internationally. Using population rather than employment outcomes in the commuting-zone analysis produces similar results to Figure 6.

Figure 6: Effect of Birth Rates on Employment and Aggregate Earnings (Commuting Zones)

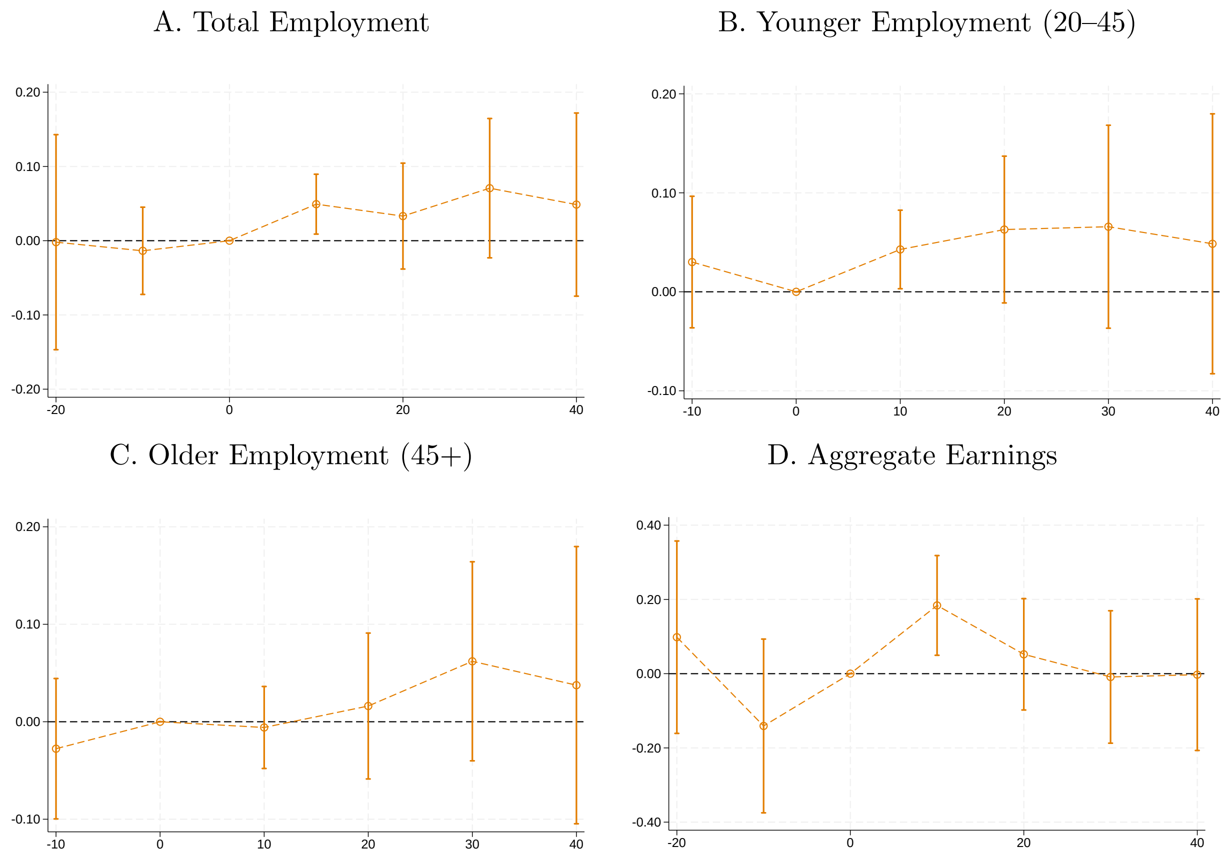

Notes: This figure reports estimates of β^h in equation (10) alongside 95% confidence intervals, estimated in our cross-commuting-zone dataset. The dependent variable is $\Delta_h y_{c,t+h}$. Panel A reports estimates where $y_{c,t}$ is the log of total commuting-zone employment. Panel B reports estimates where $y_{c,t}$ is log employment of those aged 20–45. Panel C reports estimates where $y_{c,t}$ is log employment of those aged 45 and over. Panel D reports estimates where $y_{c,t}$ is log aggregate earnings. These variables are all taken from the Census Tabulation Files. The vector of controls $\boldsymbol{x}_{c,t}$ includes initial log average wages and labor supply in year t, measured by log population aged 20–45 and log population aged 45–70. All regressions weight commuting zones by working-age population in year t. Standard errors are clustered by commuting zone.

results in the cross-country analysis.[24] This result again suggests that the estimated wage response to lower birth rates reflects a combination of unchanged aggregate earnings and a smaller workforce.

The estimated employment response of about 7% at horizon $t+30$ (50 years after the birth-rate change) is less than half of the corresponding country-level population response of approximately 16%. The primary reason for this difference is that birth rates are much less persistent across commuting zones than across countries. Figure 7 establishes that the persistence of birth rates after 40 years is less than half as large in the commuting-zone sample as in our country sample. Figure A4 shows that this difference can account for the contrast between the workforce responses

[24]The results are also similar using unadjusted aggregate earnings in the commuting zone as shown in Table A6 in the Appendix.

in Figures 4 and 6. It plots the response of a (counterfactual) native-born population, computed using age-specific national survival rates and actual local birth rates, and this response closely tracks the observed employment response. The smaller commuting-zone employment estimates, relative to their cross-country counterparts, are therefore driven by the dynamics of birth rates.

Figure 7: Birth Rate Dynamics in Countries and Commuting Zones

A. Countries B. Commuting Zones

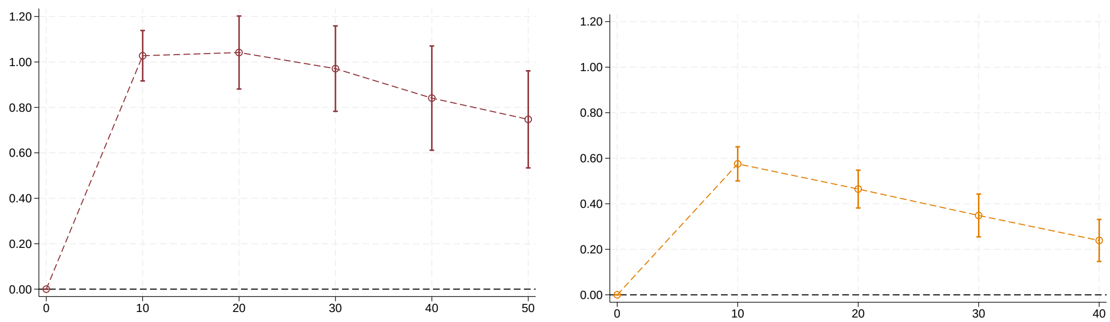

Notes: This figure reports estimates of β^h from the regression $\Delta_h b_{c,t+h} = \alpha^h + \beta^h b_{c,t} + \boldsymbol{x}'_{c,t}\boldsymbol{\gamma}^h + \varepsilon^h_{c,t+h}$ alongside 95% confidence intervals. Panel A reports estimates in our cross-country dataset, while Panel B reports estimates in our cross-commuting-zone dataset. The vector $\boldsymbol{x}_{c,t}$ includes initial average income and labor supply, measured by log population aged 20–45 and log population aged 45–70. Average income is measured as GDP per working-age adult across countries and average wage per worker across commuting zones. All regressions weight countries equally, and commuting zones by working-age population in year t. Standard errors are clustered by geographic unit c.

To complete the picture, we repeat the back-of-the-envelope calculation from the cross-country case. Employment in commuting zones decreases by only about 7% in response to a one percentage-point lower birth rate. Consequently, the offsetting technology and capital responses should compensate for an effect of $0.07 \times$ labor share $\simeq 4.2\%$, which contrasts with the implied 9.6% effect in our cross-country sample. This is in line with the smaller workforce response; since technology and capital respond to labor supply, and the labor-supply response is about half as large across commuting zones, we would expect the induced responses to be commensurately smaller.

Overall, although magnitudes differ across countries and commuting zones, both indicate that lower birth rates raise per-worker income while leaving aggregate GDP and earnings essentially unchanged.

7 Technology Responses

We propose that the development and adoption of labor-saving technology accounts for our main result that lower birth rates raise GDP per worker and composition-adjusted wages while leaving aggregate GDP and earnings broadly unaffected. This section investigates this mechanism by examining the effect of birth rates on five margins of technological adjustment across countries and commuting zones: high-tech exports and employment, industry composition, capital, total factor productivity (TFP), and patenting.

7.1 High-Tech Exports and Employment

Our first measure of labor-saving technology is the prominence of high-tech sectors, measured by the export share of high-tech products across countries and the high-tech employment share across commuting zones.[25]

Figure 8: Effect of Birth Rates on High-Tech Exports and Employment Shares

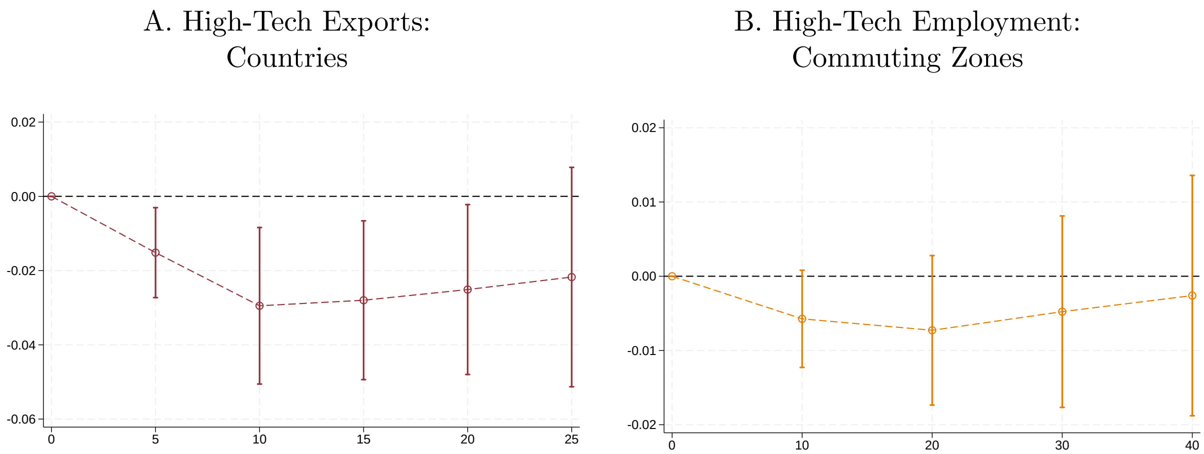

Notes: This figure reports estimates of β^h in equation (10) alongside 95% confidence intervals. Panel A reports estimates where the dependent variable $y_{c,t}$ is the share of exports in high-tech products across countries, while Panel B reports estimates where $y_{c,t}$ is the share of employment in high-tech industries across commuting zones. The vector $\boldsymbol{x}_{c,t}$ includes initial average income and labor supply, measured by log population aged 20–45 and log population aged 45–70. Average income is measured as GDP per working-age adult across countries and average wage per worker across commuting zones. All regressions weight countries equally and commuting zones by working-age population in year t. Standard errors are clustered by geographic unit c.

Figure 8 presents local-projection impulse responses for these outcome measures. Panel A depicts results for the high-tech export share across countries.[26] Lower birth rates at $t-20$ are associated with a significantly higher high-tech export share, with the effect emerging by $t+10$, and remaining stable over the subsequent 15 years. Panel B shows analogous results for the high-tech employment share across commuting zones. The pattern closely mirrors Panel A: lower birth rates lead to increases in the high-tech employment share beginning roughly one decade after cohort entry, with effects remaining stable over time. Appendix Table A7 reports long-difference estimates

[25] While high-tech sectors do not exclusively produce labor-saving technologies, many frontier technologies incorporate advanced machinery, digital equipment, and information and communication tools that tend to be labor-saving. See, for example, Berman et al. (1994) on labor-saving technical change in manufacturing; Autor et al. (2003) on computers; Graetz and Michaels (2018) and Acemoglu and Restrepo (2022b) on industrial robots; Acemoglu and Restrepo (2022b) on digital manufacturing equipment and software systems; Kogan et al. (2023) on patent-level evidence of worker displacement; and Autor et al. (2024) on automation technologies broadly. Section 5 details our classification of sectors into high-tech.

[26] Throughout this section and the next, we use the same covariates as in the main analysis: initial GDP per worker or wage and labor supply. Including the initial values of the dependent variables does not materially change the results. Detailed export data from Comtrade begin in 1995, so we estimate the local projection at five-year intervals up to a 25-year horizon and report no pre-periods.

Figure 9: Effect of Birth Rates on Employment Shares in Labor-Intensive and R&D-Intensive Industries

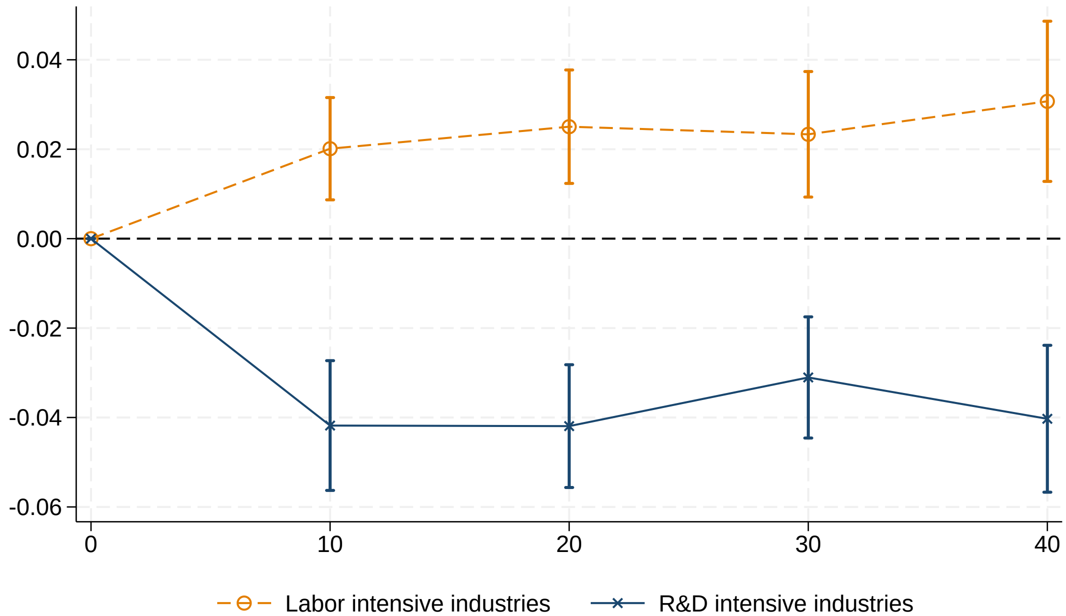

Notes: This figure reports estimates of β^h in equation (10) alongside 95% confidence intervals, estimated in our cross-commuting-zone dataset. The orange line plots estimates where $y_{c,t}$ is the share of employment in labor-intensive industries, while the blue line plots estimates where $y_{c,t}$ is the share of employment in R&D-intensive industries. All regressions weight commuting zones by working-age population in year t. Standard errors are clustered by commuting zone.

for the same variables.

Figure 9 provides complementary evidence for commuting zones by examining a finer industry partition: R&D-intensive vs. labor-intensive industries (see Section 5 for the construction of these industry groups). Employment in R&D-intensive industries is another measure of high-tech activity, while labor scarcity should also induce reallocation of workers away from labor-intensive industries. This is the pattern we find in Figure 9: lower birth rates are associated with an expansion of employment in R&D-intensive industries and with a notable contraction in labor-intensive industries. Appendix Table A7 (columns 4 and 5) reports long-difference estimates for these outcomes.

In summary, with measures based on high-tech export shares, high-tech employment shares, and the composition of employment between R&D-intensive and labor-intensive industries, we find evidence consistent with our theoretical predictions: lower birth rates and the resulting scarcity of younger labor induce a shift of economic activity toward high-tech sectors and away from labor-intensive sectors.

7.2 TFP and Capital: Country-Level and Industry-Level Evidence

We next explore the response of TFP and capital stocks to birth rates. Recall from Section 3 that the near-zero effects on GDP per worker require a sizable TFP response, which we estimate in this subsection. In addition, our theory clarifies that, contrary to the neoclassical mechanism, labor scarcity will typically be associated with higher investment and higher capital stocks. Evidence from both TFP and capital investment responses broadly supports these expectations, though TFP estimates are computed as residuals. We again report local-projection impulse responses in the text, with TFP regression estimates presented in Appendix Table A7.

Panel A of Figure 10 shows that lower birth rates are associated with significantly faster country-level TFP growth. The estimates imply that a one percentage-point lower $t-20$ (e.g., 1970) birth rate is associated with 15 log points higher TFP by $t+30$ (2020). This exceeds the 9.6 log point response required for a complete offset of the direct effect of labor scarcity in our cross-country sample, even though it excludes the additional boost to GDP per worker coming from the capital response (which we document below). Therefore, we conclude that the TFP response is broadly in line with our theoretical expectations, given roughly zero impact on GDP per worker.[27]

Because TFP and capital data are not available at the commuting-zone level, in Panel B, we turn to industry-level evidence. Specifically, we regress industry TFP growth on "exposure to labor scarcity," defined as

$$\tilde{b}_{i,t-20} = \sum_{c} \frac{\text{EMP}_{i,c,t}}{\sum_k \text{EMP}_{i,k,t}} b_{c,t-20},$$

which is equal to the employment-weighted average birth rate across commuting zones where each industry i operates (lagged by two decades). The idea is that labor scarcity in the commuting zones in which an industry operates should affect its technological trajectory—reflected in TFP growth—whereas labor scarcity in other locations should have a smaller or no effect. Control variables in these specifications (average income and population structure) are constructed analogously, as employment-weighted averages across commuting zones. Panel B of Figure 10 shows that the estimates are consistently negative, indicating that greater exposure to higher birth rates is associated with lower industry TFP. Nevertheless, the standard errors are large and the estimates are not statistically distinguishable from zero.

We next examine capital stocks across countries and industries, estimated using the same specification. Panel C of Figure 10 shows capital stocks rise following a fall in birth rates across countries. Panel D presents similar results across industries, again using the age-exposure design. The estimates are noisy in both cases, so we draw weaker conclusions than from the cross-country analysis.

[27]We report horizons up to only $h = 30$ to ensure comparability with the industry-level data, which begin in 1987.

Figure 10: Effect of Birth Rates on TFP and Capital Stock, 1980–2020

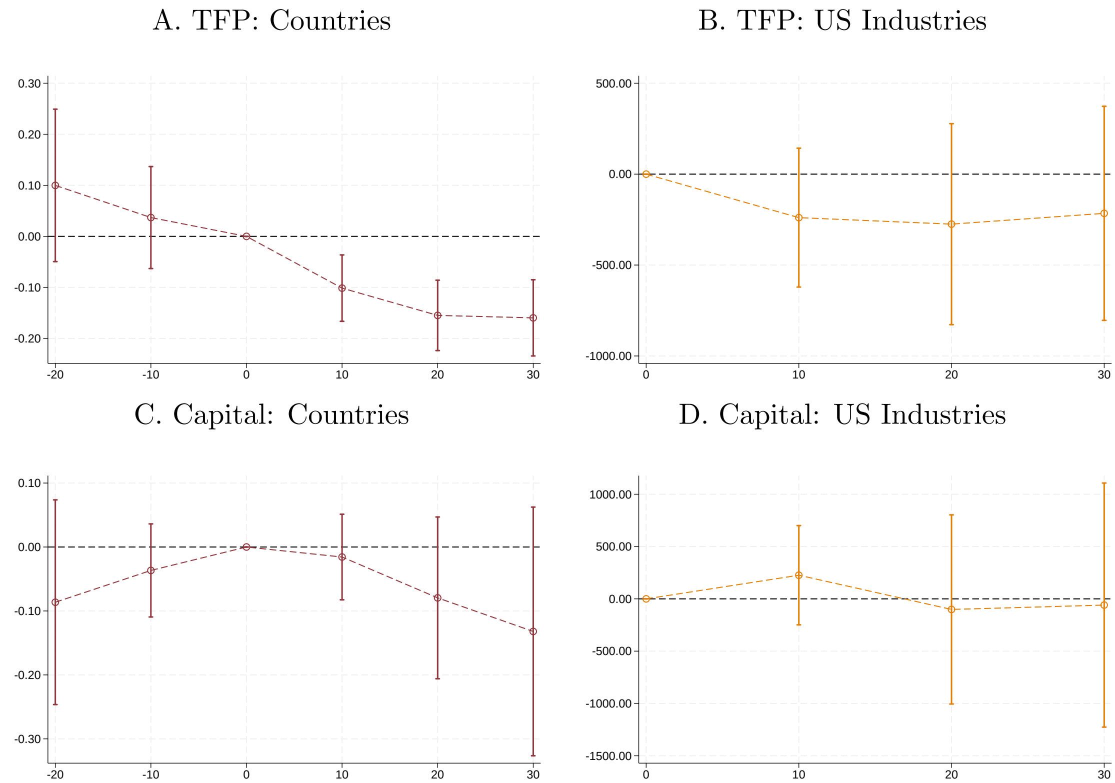

Notes: This figure reports estimates of β^h in equation (10) alongside 95% confidence intervals. Panels A and C report estimates in our cross-country dataset, while Panels B and D report estimates in our cross-industry dataset, which covers 60 US industries. Panels A and B report estimates where the dependent variable $y_{c,t}$ is TFP. Industry TFP is from the BEA, while country TFP is computed as the cumulative Solow residual (see text for details). Panels C and D report estimates where $y_{c,t}$ is the capital stock. The capital stock is from the Penn World Tables for countries; for industries we construct it from the BEA KLEMS accounts. The vector $\boldsymbol{x}_{c,t}$ includes initial average income and labor supply, measured by log population aged 20–45 and log population aged 45–70. Average income is measured as GDP per working-age adult across countries and average wage per worker across industries, whose measurement is described in the text. All regressions weight countries equally and industries by employment in 1990. Standard errors are clustered by geographic unit c.

7.3 The Response of Labor-Saving Technologies

We construct proxies for labor-saving technology from patent data across countries and commuting zones, as described in Section 5. We caution that the primary response should occur at the technology-*adoption* margin, while patents are more informative about technology *development*. Moreover, where patents are developed is only partially informative about where they are used (Kalyani et al., 2025).

We consider three proxies for labor-saving patenting. For the cross-country analysis, we follow Acemoglu and Restrepo (2022a) and focus on foreign (non-US) patents filed with the USPTO, which ensures comparable measurement across countries. For the parallel US analysis, our measure of innovation activity is the flow of new patent applications (that are subsequently granted) by assignees from a given commuting zone and year. Recall also from Section 5 that we consider ten-year averages of patent flows centered on decennial census years, so, for example, our measure of 2010 patenting is the average number of successful patent applications filed between the start of 2005 and the end of 2014. The three proxies are:

1. **Patenting by labor-intensive industries**. We test here whether industries more exposed to labor scarcity exhibit faster patenting growth, complementing the noisier industry-TFP results. We caution, however, that we should not expect all patents in labor-intensive industries to necessarily be labor-saving ex ante.

2. **Information and communication technologies**. Many digital technologies are labor-saving (as noted above), motivating us to examine patenting in ICT categories directly. We similarly caution that not all ICT patents are necessarily labor-saving.

3. **Automation technologies.** Automation is known to be labor-saving ex ante (see Section 3). As discussed in Section 5, we classify patents as automation-related following the approach in Autor et al. (2024).

We report local-projection impulse responses for citation-weighted patents in Figure 11, with the corresponding regression estimates presented in Appendix Table A8.[28] The first two panels show the relationship between birth rates and the share of all newly granted patents originating in a country or commuting zone that are filed by labor-intensive industries. Labor-intensive industries should contract when labor becomes scarce, potentially reducing their innovative activity. At the same time, we expect their production methods to become more labor-saving. Panels A (countries) and B (commuting zones) support this latter prediction: lower birth rates are associated with a higher share of patents in labor-intensive industries, even though the commuting-zone responses are imprecisely estimated. Since commuting zones are likely too small a geographic unit to fully capture innovation responses to labor scarcity, the relatively weak responses at the commuting-zone level are not surprising.

Panels C and D show similar results for the share of patents in ICT classes. These estimates are statistically significant for countries and hover near significance for commuting zones.

Finally, Panels E and F examine automation patents. The point estimates again suggest that labor scarcity is associated with more automation patents, though the results are less precise and are statistically significant at conventional levels only in the first decade of the impulse response.

In sum, our results consistently point in the direction predicted by our theoretical framework: country-level high-tech exports increase; commuting-zone labor reallocates away from labor-intensive and toward R&D-intensive and high-tech industries; country-level TFP growth rises; and patenting shifts toward labor-saving technology categories (ICT and automation) and labor-intensive industries. Where comparisons are feasible, these effects are generally stronger and more precisely estimated in cross-country data. The TFP results also suggest that this response is broadly sufficient to offset the direct effect of the decline in labor input.

[28]Following Acemoglu and Restrepo (2022a) we include the initial share of employment in manufacturing as an additional control in all regressions where patents are an outcome.

Figure 11: Effect of Birth Rates on Patenting: Labor-Intensive, ICT, and Automation Patents

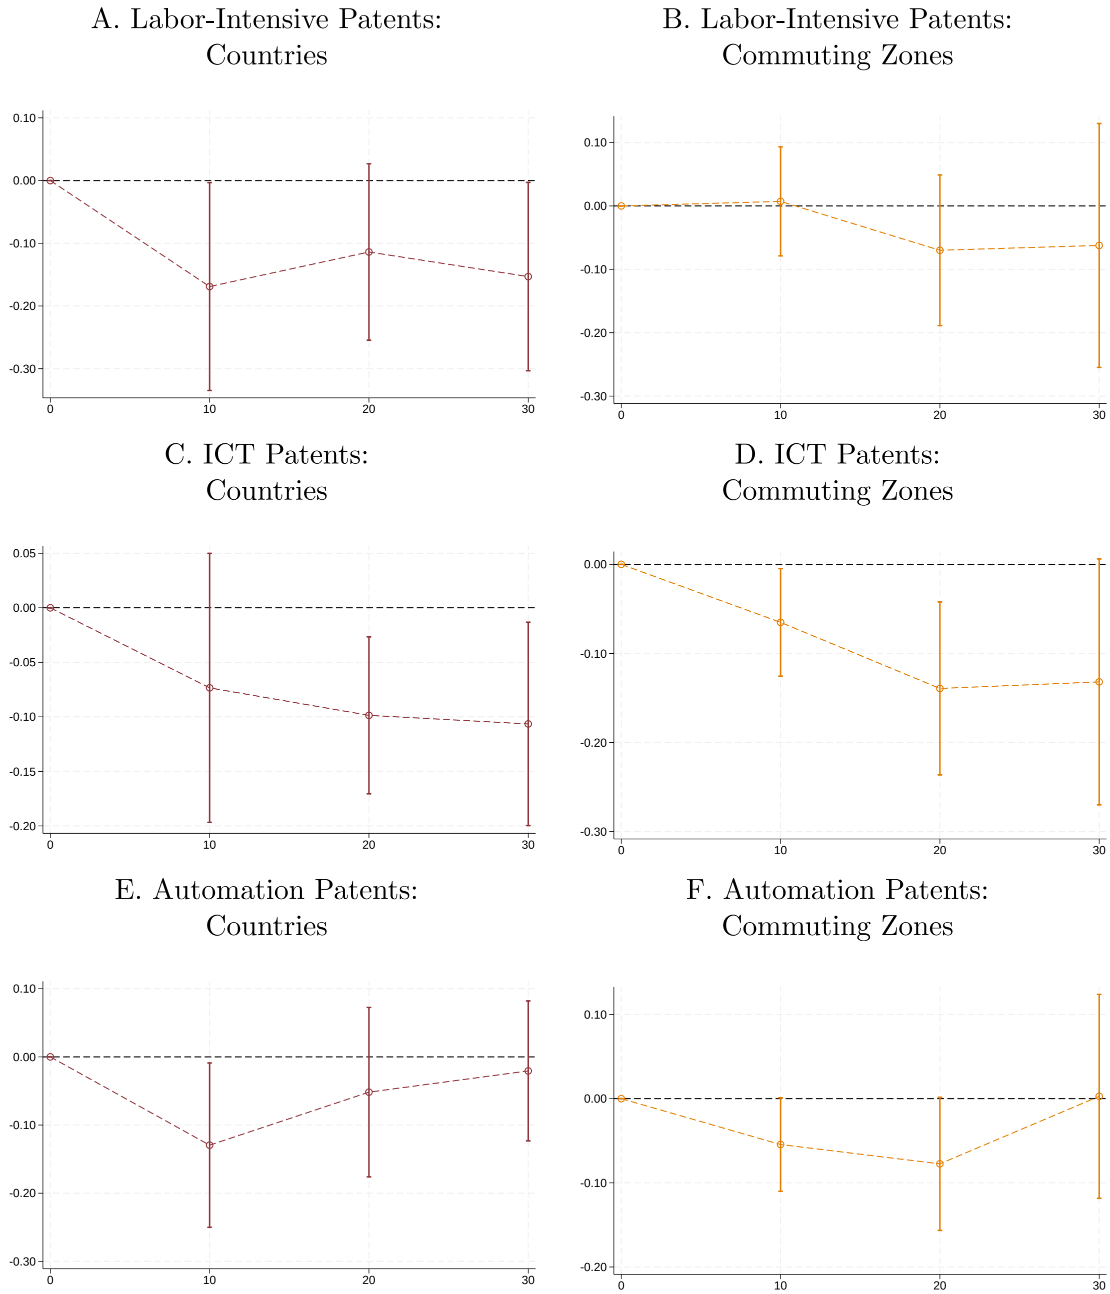

Notes: This figure reports estimates of β^h in equation (10) alongside 95% confidence intervals. Panels A, C, and E report estimates in our cross-country dataset while Panels B, D, and F report the equivalent estimates in our cross-commuting-zone dataset. Panels A and B report estimates where the dependent variable $y_{c,t}$ is the share of new patents (citation-weighted) in labor-intensive industries. Panels C and D report estimates where $y_{c,t}$ is the share of new patents in ICT categories, while Panels E and F report estimates where $y_{c,t}$ is the share of new patents classified as automation. The vector $\boldsymbol{x}_{c,t}$ includes initial average income, manufacturing share of employment, and labor supply, measured by log population aged 20–45 and log population aged 45–70. Average income is measured as GDP per working-age adult across countries and average wage per worker across commuting zones. All regressions weight countries equally and commuting zones by working-age population in year t. Standard errors are clustered by geographic unit c.

8 Alternative Mechanisms

Other channels could link declining birth rates to higher wages and productivity. Recall that the neoclassical-Solow model, in which wages rise because labor scarcity raises capital per worker, is already ruled out by our findings, as it cannot explain the near-zero impact on aggregate GDP or earnings per worker, and it predicts declining investment and capital, the opposite of what we observe. We examine three other channels below: female labor force participation, educational upgrading, and structural transformation from agriculture to industry. We present local-projection impulse responses, with regression results in Appendix Table A9.

8.1 Female Labor Force Participation

Declining birth rates may facilitate, or occur in response to, women's increased participation in the workforce. This might in turn boost output and support further skill acquisition, which could partially explain why lower birth rates predict higher wages and GDP per worker.

The results reported in Panels A and B of Figure 12 do not support this mechanism. Across both countries and commuting zones, higher birth rates are associated with *greater*, rather than lower, female labor force participation, though the commuting-zone estimates are not statistically distinguishable from zero.

8.2 Educational Upgrading

Lower birth rates may facilitate growth by enabling parents to invest more in the human capital of their children. Smaller birth cohorts might also allow schools to devote more resources per student, enabling more or better education (Bound and Turner, 2007).

Panels C and D of Figure 12 examine this mechanism by looking at years of schooling across countries and the share of the population aged 25–45 that has a college degree.[29] At the cross-country level (Panel C), the estimated relationship is positive, suggesting that lower birth rates are associated with *less* schooling. Conversely, across commuting zones, lower birth rates are associated with greater college education among younger adults. But this effect does little to explain the magnitude of our main results. A one percentage-point lower birth rate is associated with a roughly 2 percentage-point increase in the college share versus a more than 10 log point increase in composition-adjusted wages, as reported in Table 2. Since the college premium is approximately 75 log points (Bengali et al., 2025), the 2 percentage-point increase in college education contributes only about a 1.5 ($= 0.02 \times 75$) log point rise in the wage level, which is only a small portion of the 10 log point increase.

Figure 12: Alternative Mechanisms I: Female LFP and Schooling

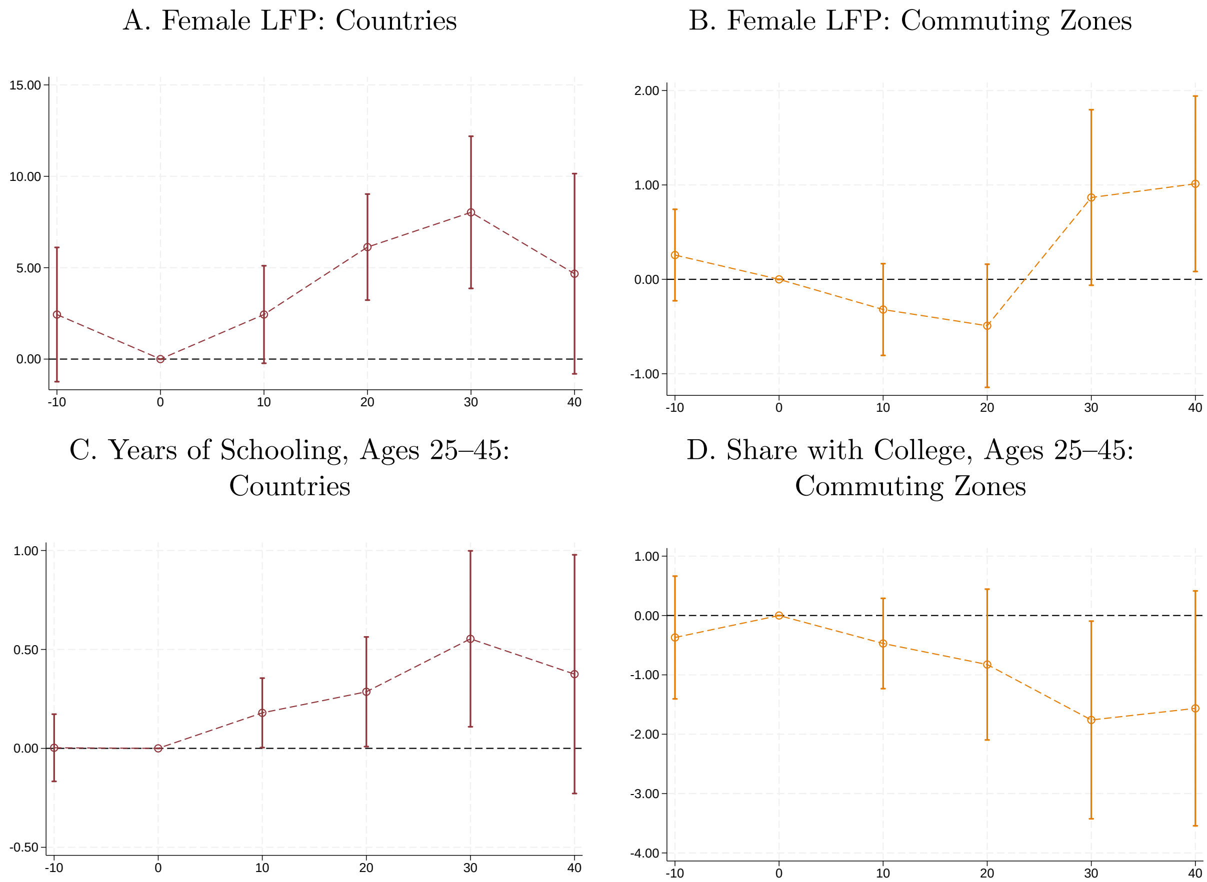

Notes: This figure reports estimates of β^h in equation (10) alongside 95% confidence intervals. Panels A and C report estimates in our cross-country dataset, while Panels B and D report estimates in our cross-commuting-zone dataset. The dependent variable is $\Delta_h y_{c,t+h}$, where $y_{c,t}$ is the female labor force participation rate in Panels A and B, and the education level in Panels C and D. Education is measured by average years of schooling for those aged 25–45 across countries and by the share aged 25–45 with a college degree across commuting zones. The vector $\boldsymbol{x}_{c,t}$ includes initial average income and labor supply, measured by log population aged 20–45 and log population aged 45–70. Average income is measured as GDP per working-age adult across countries and average wage per worker across commuting zones. All regressions weight countries equally and commuting zones by working-age population in year t. Standard errors are clustered by geographic unit c.

8.3 Structural Transformation

Lower birth rates may accelerate structural transformation, enabling workers to move out of agriculture and into manufacturing. We explore these channels in Figure 13 but do not find consistent supporting evidence at the country and commuting-zone levels. Across the four panel estimates, two move in the hypothesized direction—agricultural employment falls and manufacturing employment rises in countries with lower birth rates (Panels A and C)—but these relationships are either statistically insignificant or economically small. The results for commuting zones (Panels B and D) point in the opposite direction.

[29] We do not observe the intensive margin of educational quality.

Figure 13: Alternative Mechanisms II: Structural Change

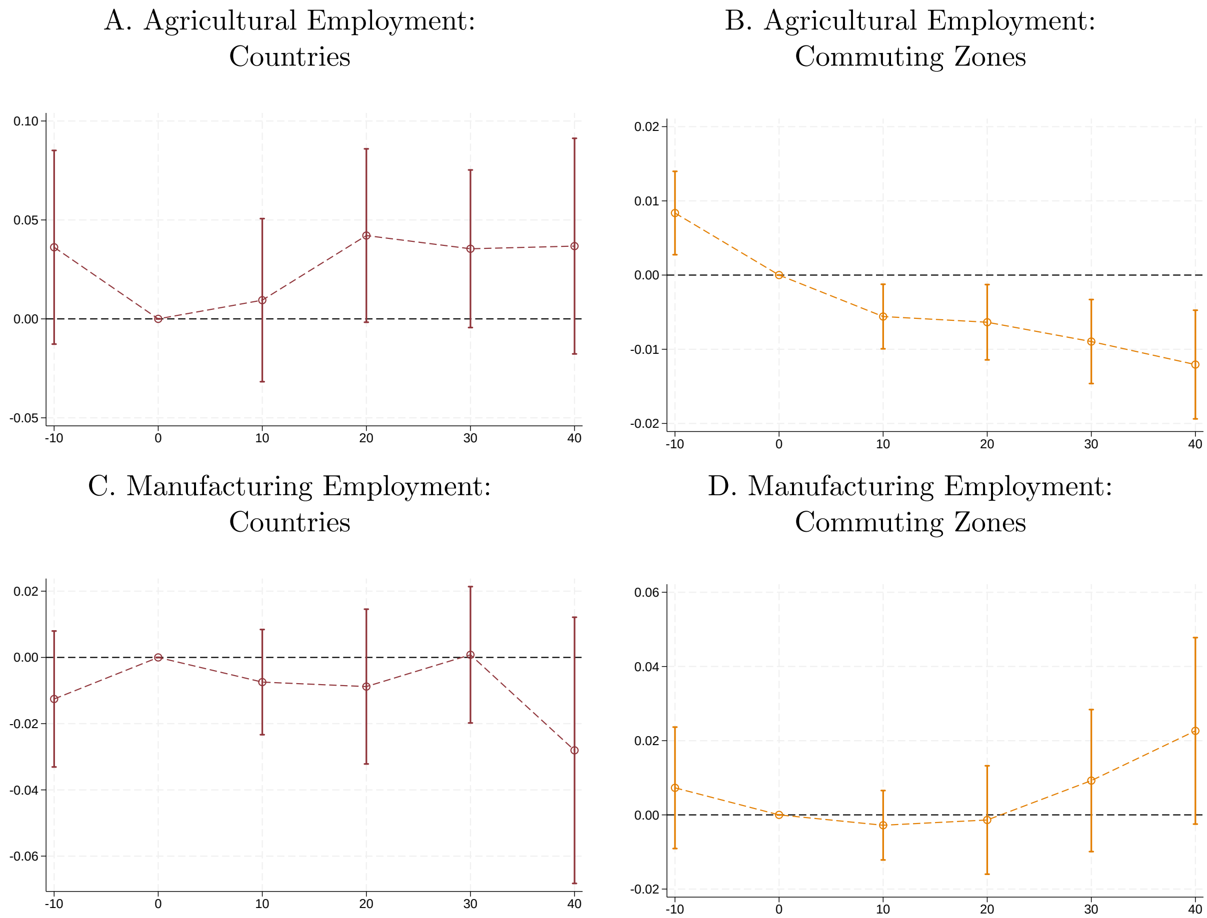

Notes: This figure reports estimates of β^h in equation (10) alongside 95% confidence intervals. Panels A and C report estimates in our cross-country dataset, while Panels B and D report estimates in our cross-commuting-zone dataset. The dependent variable is $\Delta_h y_{c,t+h}$, where $y_{c,t}$ is the share of employment in agriculture in Panels A and B, and the share of employment in manufacturing in Panels C and D. The vector $\boldsymbol{x}_{c,t}$ includes initial average income and labor supply, measured by log population aged 20–45 and log population aged 45–70. Average income is measured as GDP per working-age adult across countries and average wage per worker across commuting zones. All regressions weight countries equally and commuting zones by working-age population in year t. Standard errors are clustered by geographic unit c.

In summary, none of the three alternative mechanisms we have examined (female labor force participation, educational upgrading, or structural transformation) can account for the relationship between birth rates and economic growth.

9 Distinguishing Population and Aging Effects

Across countries and commuting zones, we find robust and quantitatively large positive effects of lower birth rates on income per capita and wage growth. Because falling birth rates simultaneously reduce population growth and increase average age, our evidence does not distinguish the separate contributions of aging and population decline. In this section, we exploit variation across countries in population and demographic structure arising from World War II deaths to disentangle these channels. Military deaths are concentrated among young men, whereas civilian deaths occur across the age distribution.[30] Civilian deaths can therefore be expected to primarily reduce population size without dramatically altering the age distribution, while military deaths both reduce population and shift the age distribution toward older groups.

As detailed in Section 5, we measure WWII civilian and military deaths using data compiled by Kesternich et al. (2014), supplemented with data from the Encyclopedia Britannica and Wikipedia for countries not covered in this source. We combine these data with GDP and population data from the Maddison Project database (Bolt and van Zanden, 2025) and data on population by age from the International Historical Statistics database (Mitchell, 2013). Because these sources do not enable us to compute reliable numbers for working-age population, we use GDP per capita instead of GDP per worker as our outcome variable. Our main sample contains data on population, war deaths, and GDP for 31 countries for which we have well-sourced data on war deaths. The decomposition of population by age in 1935 is available for only 18 of these countries, which yields our restricted sample.

We estimate the following model for the effect of WWII deaths on subsequent country-level log GDP per capita:

$$\ln y_{c,1985} = \alpha^h + \beta_M^h d_c^M + \beta_C^h d_c^C + \rho^h \ln y_{c,1935} + \varepsilon_c^h$$

where $y_{c,1985}$ is GDP per capita in country c in year 1985, d_c^M and d_c^C are military and civilian death rates (measured as a share of 1935 population), and the coefficients β_M^h and β_C^h can be interpreted as the impulse responses to war deaths over horizon h. We use 1935 as the baseline (pre-war) year and control for initial income at that time.

Table 3 reports estimates of the impact of war deaths on 1955 populations and 1985 incomes. Columns 1 and 2 report the impact of military and civilian deaths on log total population in 1955 for the full sample of 31 countries and the restricted sample of 18 countries, respectively. Column 3 reports the impact on the log share of younger adults (aged 20–45) for the 18 countries for which age composition data are available.

Consistent with expectations, both military and civilian WWII deaths predict reduced population levels in 1955 (Columns 1 and 2). In response to a one percentage-point rise in military deaths, 1955 population is estimated to be roughly 2.2 log points lower, though this effect is not statistically significant. The corresponding response to civilian deaths is about 3.6 log points and highly significant. Results are similar in the restricted 18-country subsample.

Column 3 focuses on the share of population aged 20–45 rather than total population. This estimate indicates that military deaths reduce the share of younger adults, whereas civilian deaths show no significant association with age composition.

Columns 4 and 5 report the effects of military and civilian WWII deaths on GDP per capita in 1985, four decades after the war's end, for the full and restricted samples, respectively. Overall,

[30] Jdanov et al. (2005) find that mortality rates spike for men aged 18–25 during World War II in Britain, while the concurrent mortality increase among civilians was roughly uniform across the population.

Table 3: Effect of World War II Deaths on Population, Age Composition, and GDP

	(1) Pop, 1955 (N=31)	(2) Pop, 1955 (N=18)	(3) Young Share, 1955 (N=18)	(4) GDP/cap, 1985 (N=31)	(5) GDP/cap, 1985 (N=18)
Military deaths	-0.022	-0.010	-0.028	0.091	0.045
	(0.018)	(0.012)	(0.011)	(0.037)	(0.025)
Civilian deaths	-0.036	-0.024	-0.003	-0.036	-0.046
	(0.011)	(0.008)	(0.006)	(0.027)	(0.018)
R-sq	0.389	0.522	0.397	0.099	0.449

Notes: This table reports estimates of β_M^h and β_C^h from the regression $z_{c,1935+h} = \alpha^h + \beta_M^h d_c^M + \beta_C^h d_c^C + \rho^h \ln y_{c,1935} + \varepsilon_c^h$ where d_c^M and d_c^C denote military and civilian deaths respectively. Column 1 reports estimates where $z_{c,t}$ is log population and $h = 20$. Column 2 reports the same estimates on the subsample of 18 countries for which age composition data are available. Column 3 reports estimates where $z_{c,t}$ is the log share of population aged 20–45 while again $h = 20$. Column 4 reports estimates where $z_{c,t}$ is log GDP per capita and $h = 50$. Column 5 reports the estimates from the same regression as Column 4, on the 18 countries for which age composition data are available. Heteroskedasticity-robust standard errors are reported in parentheses.

military deaths predict higher GDP per capita. Civilian deaths, by contrast, predict a reduction in future GDP per capita.

Figure 14 depicts local-projection impulse responses for GDP per capita to wartime deaths over the half-century following 1945 for our full sample of 31 countries. This figure shows no evidence of pre-trends—no significant relationship between civilian or military WWII deaths and GDP per capita growth between 1905 and 1935. In the decades after WWII, however, the contrasting effects of civilian and military deaths on GDP per capita growth accumulate visibly. Consistent with the estimates reported in Table 3, we find that a 1% population loss (relative to 1935) due to civilian deaths predicts 3.6 log points lower GDP per capita five decades later, while a 1% loss due to military deaths predicts 9.1 log points higher GDP per capita over the same period. The military death effects become significant at 5% from 1985 onwards.[31]

Although these estimates are somewhat imprecise, reflecting in part our limited sample of 31 countries, they are still informative about the channels through which falling birth rates affect incomes. Civilian wartime deaths reduce population broadly without altering age composition (Column 3), and they are associated with lower GDP per capita. Military deaths, by contrast, disproportionately remove young men from the population, shifting the age distribution toward older workers and creating scarcity of younger workers. This age-composition shift predicts higher GDP per capita, consistent with our conceptual framework: scarcity of younger workers should spur adoption of labor-saving technologies whose productivity gains more than offset the direct loss of workers. Lower birth rates operate through both channels simultaneously, reducing total

[31] Note that both death and birth rates are calculated as a share of 1935 population, and so, a one percentage-point change in death rates mechanically reduces population growth starting from 1935 by one percentage point.

Figure 14: Effect of World War II Deaths on GDP Growth: Distinguishing Civilian vs. Military Deaths

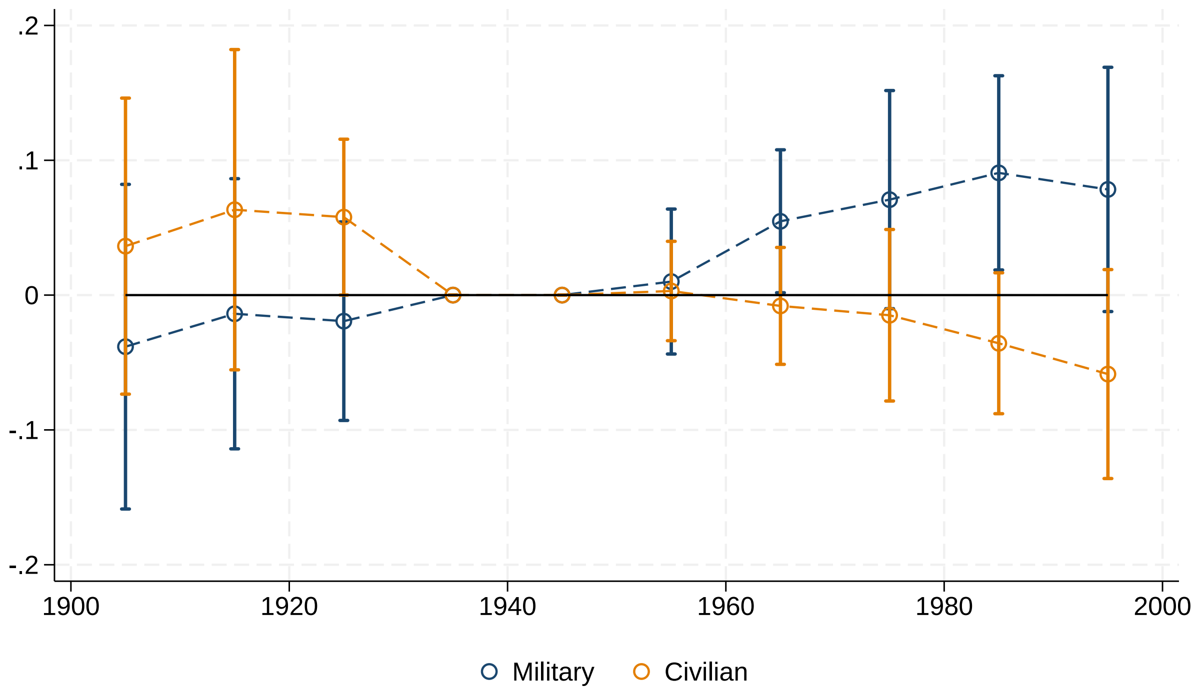

Notes: This figure plots estimates β^h_M and β^h_C from the regressions $\ln y_{c,1935+h} = \alpha^h + \beta^h_M d^M_c + \beta^h_C d^C_c + \rho^h \ln y_{c,1935} + \varepsilon^h_c$ alongside 95% confidence intervals, where $y_{c,t}$ denotes GDP per capita in country c at time t, and d^M_c and d^C_c denote military and civilian deaths respectively. Standard errors are heteroskedasticity robust.

population while also shifting the age distribution rightward. On net, the age-composition effect dominates. This is evident in Table 3, and it is consistent with our main results across countries and commuting zones in Tables 1 and 2 using data from more recent decades.

This WWII-era evidence is an out-of-sample test of our causal mechanism that points to the same conclusion as our main analysis: the primary channel by which lower birth rates raise GDP per capita (and wages across local labor markets) is scarcity of younger workers. Our findings are corroborated by those of Bergeaud et al. (2025), who document that French regions that experienced higher WWI military deaths saw a subsequent rise in patenting in labor-saving technologies.

10 Conclusion

The world has witnessed a remarkable and nearly universal decline in birth rates over the last several decades, with fertility falling to or below replacement level across essentially all industrialized countries and in much of East and Southeast Asia. Many developing countries are on the same trajectory. A broadly shared expectation of academics and policymakers is that the resulting population declines and workforce aging will create strong adverse macroeconomic headwinds that slow productivity and economic growth. These concerns rest on forward projections, but seven decades of rising and falling birth rates and accompanying demographic shifts offer historical evidence with which to assess their plausibility.

We exploit this historical record across a large set of countries and across US local labor markets to evaluate the macroeconomic and labor market consequences of falling birth rates. In contrast to the prevailing wisdom, we find that lower birth rates so far have led to higher growth in GDP per worker across countries and higher wage growth across local labor markets in the United States. Strikingly, despite lower labor input resulting from lower birth rates and aging workforces, GDP across countries and aggregate earnings across commuting zones have not declined.

Our results are inconsistent with neoclassical-Solow mechanisms, which predict lower aggregate output and a dissipation of the positive per-worker effects over time, neither of which we observe. They are also inconsistent with models in which larger populations drive short-run or long-run economic growth. Although these patterns could in principle be mediated by higher female labor force participation, educational upgrading, or accelerated structural transformation accompanying lower birth rates, we find little evidence supporting these channels.

We offer a conceptual framework that shows how the development and adoption of labor-saving technologies may respond elastically to labor scarcity—specifically, the scarcity of younger workers—which can in turn fully offset or even reverse the negative direct effects of lower labor inputs. Multiple threads of evidence support this mechanism. We find that lower birth rates are associated with more labor-saving patents and a growing share of high-tech industries across countries and US commuting zones. They are also predictive of higher TFP growth across countries and US industries, and increased patenting in ICT and broader automation technology classes. Using cross-country variation in WWII-era military and civilian deaths, we present suggestive evidence that the positive growth effects of falling birth rates are primarily driven by the scarcity of younger workers rather than by reductions in aggregate population per se. Although these technology responses—labor-saving patenting and sectoral shifts toward high-tech industries—do not directly translate into compensating GDP effects, the country-level TFP response is large enough to fully offset the direct negative impact of labor scarcity on GDP per worker.

Our findings describe the past; whether they provide a reliable guide to the economic consequences of coming demographic transitions is not yet established. First, there will be much faster aging and larger population declines across many countries in the decades ahead, particularly in China. Changes this rapid are outside the support of the historical evidence, and could theoretically yield different adjustment dynamics from those we document. Second, we can offer only indirect evidence that endogenous technology responses are responsible for the positive effects of falling birth rates on economic growth. Third, technology and investment responses may differ between periods, making extrapolation into the future hazardous. At the same time, the demographic changes currently underway are accompanied by increased life expectancy (Scott, 2021, 2024), which may spur institutional changes, policies, and additional human capital investments that complement longer lifespans. These forces may counteract any negative effects of aging and population declines.

References

Acemoglu, D. (2007). Equilibrium bias of technology. *Econometrica*, 75(5):1371–1409.

Acemoglu, D. (2010). When does labor scarcity encourage innovation? *Journal of Political Economy*, 118(6):1037–1078.

Acemoglu, D., Aghion, P., Barrage, L., and Hémous, D. (2023). Climate change, directed innovation, and energy transition: The long-run consequences of the shale gas revolution. Technical report, National Bureau of Economic Research.

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics*, volume 4, pages 1043–1171. Elsevier.

Acemoglu, D. and Johnson, S. (2007). Disease and development: the effect of life expectancy on economic growth. *Journal of political Economy*, 115(6):925–985.

Acemoglu, D., Kong, F., and Restrepo, P. (2025). Tasks at work: comparative advantage, technology and labor demand. *Handbook of Labor Economics*, 6:1–114.

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

Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American economic review*, 108(6):1488–1542.

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

Acemoglu, D. and Restrepo, P. (2022b). Tasks, automation, and the rise in us wage inequality. *Econometrica*, 90(5):1973–2016.

Aksoy, Y., Basso, H. S., Smith, R. P., and Grasl, T. (2019). Demographic structure and macroeconomic trends. *American Economic Journal: Macroeconomics*, 11(1):193–222.

Allen, R. C. (2009). *The British industrial revolution in global perspective.* Cambridge University Press.

Ashraf, Q. H., Weil, D. N., and Wilde, J. (2013). The effect of fertility reduction on economic growth. *Population and development review*, 39(1):97–130.

Autor, D., Chin, C., Salomons, A., and Seegmiller, B. (2024). New frontiers: The origins and content of new work, 1940–2018. *The Quarterly Journal of Economics*, 139(3):1399–1465.

Autor, D. and Kausik, B. (2026). Resolving the automation paradox: falling labor share, rising wages. *arXiv preprint arXiv:2601.06343*.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). Skill demand, inequality, and computerization: Connecting the dots. In *Technology, growth, and the labor market*, pages 107–129. Springer.

Baldwin, R. and Teulings, C. (2014). Secular stagnation: facts, causes and cures. *London: Centre for Economic Policy Research-CEPR*.

Barro, R. J. and Lee, J. W. (2013). A new data set of educational attainment in the world, 1950–2010. *Journal of development economics*, 104:184–198.

Bengali, L., Valletta, R. G., and Zhao, C. (2025). Explaining stagnation in the college wage premium. Working Paper 2025-01, Federal Reserve Bank of San Francisco.

Bergeaud, A., Chaniot, J., and Malgouyres, C. (2025). Escaping labor scarcity: Innovation and human capital after WW1 in France. CEPR Discussion Paper 20492, Centre for Economic Policy Research.

Berman, E., Bound, J., and Griliches, Z. (1994). Changes in the demand for skilled labor within u.s. manufacturing: Evidence from the annual survey of manufactures. *The Quarterly Journal of Economics*, 109(2):367–397.

Bloom, D., Canning, D., and Sevilla, J. (2003). *The demographic dividend: A new perspective on the economic consequences of population change*. Rand Corporation.

Bojinov, I., Rambachan, A., and Shephard, N. (2021). Panel experiments and dynamic causal effects: A finite population perspective. *Quantitative Economics*, 12(4):1171–1196.

Bolt, J. and van Zanden, J. L. (2025). Maddison-style estimates of the evolution of the world economy: A new 2023 update. *Journal of Economic Surveys*, 39(2):631–671. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/joes.12618.

Boserup, E. (2014). *The conditions of agricultural growth: The economics of agrarian change under population pressure*. Routledge.

Bound, J. and Turner, S. (2007). Cohort crowding: How resources affect collegiate attainment. *Journal of public Economics*, 91(5-6):877–899.

Clemens, M. A., Lewis, E. G., and Postel, H. M. (2018). Immigration restrictions as active labor market policy: Evidence from the mexican bracero exclusion. *American Economic Review*, 108(6):1468–1487.

Dechezleprêtre, A., Hémous, D., Olsen, M., and Zanella, C. (2021). Induced automation: evidence from firm-level patent data. *University of Zurich, Department of Economics, Working Paper*, (384).

Dube, A., Girardi, D., Jorda, O., and Taylor, A. M. (2025). A local projections approach to difference-in-differences. *Journal of Applied Econometrics*, 40(7):741–758.

Eckert, F., Fort, T. C., Schott, P. K., and Yang, N. J. (2021). Imputing missing values in the US census bureau's county business patterns. Working Paper 26632, National Bureau of Economic Research. Revised January 2021.

Eden, M. and Kuruc, K. (2023). The long-run relationship between per capita incomes and population size. Technical report, CEPR Discussion Papers.

European Bank for Reconstruction and Development (2025). Transition report.

Feenstra, R. C., Inklaar, R., and Timmer, M. P. (2015). The Next Generation of the Penn World Table. *American Economic Review*, 105(10):3150–3182.

Fernández-Villaverde, J., Ventura, G., and Yao, W. (2025). The wealth of working nations. *European Economic Review*, 173:104962.

Goos, M., Konings, J., and Vandeweyer, M. (2015). Employment growth in europe: The roles of innovation, local job multipliers and institutions. Technical report, USE Discussion Papers Series.

Gordon, R. (2017). *The rise and fall of American growth: The US standard of living since the civil war*. Princeton university press.

Graetz, G. and Michaels, G. (2018). Robots at work. *The Review of Economics and Statistics*, 100(5):753–768.

Habakkuk, H. J. (1962). *American and British technology in the nineteenth century*, volume 6. Cambridge.

Hall, B., Jaffe, A., and Trajtenberg, M. (2001). The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools. Technical Report w8498, National Bureau of Economic Research, Cambridge, MA.

Hall, R. E. and Jones, C. I. (1999). Why do some countries produce so much more output per worker than others? *The quarterly journal of economics*, 114(1):83–116.

Hanlon, W. W. (2015). Necessity is the mother of invention: Input supplies and directed technical change. *Econometrica*, 83(1):67–100.

Hansen, A. H. (1939). Economic progress and declining population growth. *The American Economic Review*, 29(1):1–15.

Hayashi, F. (2025). The root cause of japan's 30-year stagnation: Implications for china, the us, and other mature economies. *Available at SSRN 4553699*.

Hicks, J. R. (1932). *The Theory of Wages*. Macmillan, London.

Hopenhayn, H., Neira, J., and Singhania, R. (2022). From population growth to firm demographics: Implications for concentration, entrepreneurship and the labor share. *Econometrica*, 90(4):1879–1914.

International Monetary Fund (2025). World economic outlook.

Jdanov, D., Andreev, E., Jasilionis, D., and Shkolnikov, V. M. (2005). Estimates of mortality and population changes in england and wales over the two world wars. *Demographic Research*, 13:389–414.

Jones, C. I. (2022). The end of economic growth? unintended consequences of a declining population. *American Economic Review*, 112(11):3489–3527.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American economic review*, 95(1):161–182.

Kalyani, A., Bloom, N., Carvalho, M., Hassan, T., Lerner, J., and Tahoun, A. (2025). The diffusion of new technologies. *The Quarterly Journal of Economics*, 140(2):1299–1365.

Karahan, F., Pugsley, B., and Şahin, A. (2024). Demographic origins of the start-up deficit. *American Economic Review*, 114(7):1986–2023.

Karra, M., Canning, D., and Wilde, J. (2017). The effect of fertility decline on economic growth in africa: A macrosimulation model. *Population and development review*, 43:237–263.

Kaufmann, D. and Kraay, A. (2010). The worldwide governance indicators.

Kesternich, I., Siflinger, B., Smith, J. P., and Winter, J. K. (2014). The Effects of World War II on Economic and Health Outcomes across Europe. *The Review of Economics and Statistics*, 96(1):103–118.

Kogan, L., Papanikolaou, D., Schmidt, L. D. W., and Seegmiller, B. (2023). Technology and labor displacement: Evidence from linking patents with worker-level data. Technical Report 31846, National Bureau of Economic Research.

Kremer, M. (1993). Population growth and technological change: One million bc to 1990. *The quarterly journal of economics*, 108(3):681–716.

Lee, R. D. (1988). Induced population growth and induced technological progress: Their interaction in the accelerating stage. *Mathematical Population Studies*, 1(3):265–288.

Lewis, E. (2011). Immigration, skill mix, and capital skill complementarity. *The Quarterly Journal of Economics*, 126(2):1029–1069.

Li, H. and Zhang, J. (2007). Do high birth rates hamper economic growth? *The Review of Economics and Statistics*, 89(1):110–117.

Liang, J., Wang, H., and Lazear, E. P. (2018). Demographics and entrepreneurship. *Journal of Political Economy*, 126(S1):S140–S196.

Lybbert, T. J. and Zolas, N. J. (2014). Getting patents and economic data to speak to each other: An 'algorithmic links with probabilities' approach for joint analyses of patenting and economic activity. *Research Policy*, 43(3):530–542.

MacDonald, G. and Weisbach, M. S. (2004). The economics of has-beens. *Journal of Political Economy*, 112(S1):S289–S310.

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

McKendrick, A. G. (1926). Applications of mathematics to medical problems. *Proceedings of the Edinburgh Mathematical Society*, 44:98–130.

Mitchell, B. (2013). *International Historical Statistics*. Palgrave Macmillan, London.

Moscona, J. and Sastry, K. A. (2023). Does directed innovation mitigate climate damage? evidence from us agriculture. *The Quarterly Journal of Economics*, 138(2):637–701.

Murphy, K. M. and Welch, F. (1990). Empirical age-earnings profiles. *Journal of Labor economics*, 8(2):202–229.

Newell, R. G., Jaffe, A. B., and Stavins, R. N. (2010). The induced innovation hypothesis and energy-saving technological change. In *Technological change and the environment*, pages 97–126. Routledge.

Organisation for Economic Co-operation and Development (2025). Oecd economic outlook.

Peters, M. and Walsh, C. (2021). Population growth and firm dynamics. Technical report, National Bureau of Economic Research.

Popp, D. (2002). Induced innovation and energy prices. *American economic review*, 92(1):160–180.

Ruggles, S., Flood, S., Sobek, M., Backman, D., Cooper, G., Drew, J. A. R., Richards, S., Rodgers, R., Schroeder, J., and Williams, K. C. (2025). IPUMS USA: Version 16.0 [dataset].

Schroeder, J., Van Riper, D., Manson, S., Knowles, K., Kugler, T., Roberts, F., and Ruggles, S. (2025). IPUMS National Historical Geographic Information System: Version 20.0 [dataset].

Scott, A. J. (2021). The longevity economy. *The Lancet Healthy Longevity*, 2(12):e828–e835.

Scott, A. J. (2024). *The longevity imperative: how to build a healthier and more productive society to support our longer lives*. Hachette UK.

Solow, R. M. (1956). A contribution to the theory of economic growth. *The quarterly journal of economics*, 70(1):65–94.

Summers, L. H. (2016). The age of secular stagnation: What it is and what to do about it. *Foreign affairs*, 95(2):2–9.

Tolbert, C. M. and Sizer, M. (1996). Us commuting zones and labor market areas: A 1990 update.

United Nations, Department of Economic and Social Affairs, Population Division (2024). World population prospects 2024: Online edition. Accessed: 2026-02-16.

von Foerster, H. (1959). Some remarks on changing populations. In Stohlman, Jr., F., editor, *The Kinetics of Cellular Proliferation*, pages 382–407. Grune & Stratton, New York.

Yoon, M. J.-W., Kim, M. J., and Lee, J. (2014). *Impact of demographic changes on inflation and the macroeconomy*. International Monetary Fund.

Young, A. (2005). The gift of the dying: The tragedy of aids and the welfare of future african generations. *The Quarterly Journal of Economics*, 120(2):423–466.

Online Appendix (Not for Publication)

A Additional Empirical Results

Pre-1950 Cross-Country Birth Rates

Our primary cross-national data source is the UN Population Prospects, which compiles comprehensive demographic data across countries since 1950. Data on birth rates for earlier years are available from the UN Demographic Yearbooks. These data are taken directly from national vital statistics based on birth registrations, which introduces two limitations: first, births were systematically under-registered during this period; second, the data are not harmonized across countries and may reflect differences in how countries register births. The UN Population Prospects, by contrast, is harmonized across countries and reconciled against the population age structure observed in later years (for example, from full-count censuses). This information helps correct for under-reporting, particularly because death registrations were generally more comprehensive than birth registrations.

We explore the consequences of these data-construction issues in Figure A1, using UN Demographic Yearbooks data for 1930 and 1940, initially compiled by Acemoglu and Johnson (2007). Panel A plots the global aggregate birth rate (equivalent to the population weighted average) over the extended sample. We see a large jump in measured birth rates when moving from the UN Demographic Yearbooks to the comprehensive UN Population Prospects data, which strongly suggests under-counting. In Panel B of Figure A1 we report estimates of the ten-year persistence in birth rates by decade, estimated from the cross-sectional regressions

$$b_{c,t+10} = \delta_t + \rho^t b_{c,t} + \varepsilon_{c,t}, this$$

for $t \in \{1930, 1940, ..., 2010\}$, where we combine the two sources into a single series. We see that the persistence in birth rates spikes between 1940 and 1950, when the switch between data sources occurs. This spike suggests that births are under-counted on average and that this under-counting is more prevalent for countries with higher birth rates.

These caveats aside, Figure A2 shows that including the 1930 and 1940 birth rates yields results very similar to our baseline, though the extended series shows a pre-trend, which may partly reflect the data-construction issues mentioned above.

Additional Specifications and Robustness Checks

Table A1 shows cross-country and commuting-zone results with a 20-year horizon, starting from 1950 for countries and 1940 for commuting zones. The results are broadly similar regardless of starting point, confirming that birth-rate effects are stable across the baby-boom and baby-bust eras. Table A2 documents that our results are robust to the inclusion of lagged birth rates.

Tables A7 and A8 provide estimates for the technology responses presented as figures in the main text. Table A9 provides estimates for the alternative mechanisms discussed in Section 8.

Figure A4 depicts the response of (counterfactual) native-born working-age population to birth-rate differences. As described in the text, this quantity is computed by applying national age-specific death rates to (actual) births in each commuting zone.

Table A4 estimates our main specifications, except that instead of per-worker outcomes, we now have aggregate GDP (cross-country) and total labor income (commuting zones) as dependent variables. Panel A presents cross-country results, while Panel B is for commuting zones. Odd-numbered columns present long-differences estimates, while even-numbered columns use stacked differences. We include the same basic controls as in Tables 1 and 2: education, urbanization, and region fixed effects. Consistent with Figure 4, we find no major impact on aggregate GDP. Indeed, a lower birth rate is associated with slightly higher aggregate GDP, marginally significant at the cross-country level in some specifications. This motivates our conclusion that reduced labor input and aging driven by lower birth rates do not substantially change GDP or aggregate earnings.

Tables A3 and A5 document that our results are very similar when we use log birth rates, while Table A6 shows that our commuting-zone estimates are robust to alternative measures of wages.

Figure A3 further explores weighting for commuting zones by reporting unweighted estimates for the largest 137 commuting zones (with population greater than 300,000 in 1970). These estimates are very similar to our baseline weighted estimates, confirming that the slight discrepancy between weighted and unweighted results in Table 2 is driven by smaller commuting zones.

Table A1: Effect of Birth Rates on Growth, by Birth Year

	(1) 1940	(2) 1950	(3) 1960	(4) 1970	(5) 1980
			Panel A: Countries		
birth rate (t-20)		-0.13	-0.26	-0.32	-0.26
		(0.10)	(0.09)	(0.06)	(0.05)
N	.	120	120	142	142
R-sq	.	0.283	0.342	0.366	0.502
			Panel B: Commuting Zones		
birth rate (t-20)	-6.30	-2.84	-7.88	-11.06	-5.76
	(2.17)	(1.80)	(2.21)	(3.14)	(1.63)
N	722	722	720	716	720
R-sq	0.396	0.291	0.306	0.116	0.438

Notes: This table reports estimates of β^h in equation (10). Panel A reports estimates in our cross-country dataset, while Panel B reports estimates in our cross-commuting-zone dataset. The dependent variable is $\Delta_h y_{c,t+h}$, where $y_{c,t}$ is log GDP per worker in Panel A and log composition-adjusted weekly wage in Panel B. Each column corresponds to estimates using a single t, with the heading indicating the starting year for $\Delta_h y_{c,t+h}$ while the birth rate corresponds to $t-20$. Columns 1–4 use $h = 30$; Column 5 corresponds to $t = 2000$ and uses $h = 20$, the longest horizon available given the 2020 data endpoint. The vector $\boldsymbol{x}_{c,t}$ includes initial average income and labor supply, measured by log population aged 20–45 and log population aged 45–70. Average income is measured as GDP per working-age adult across countries and average wage per worker across commuting zones. Countries are weighted equally, while commuting zones are weighted by working-age population in year t. Heteroskedasticity-robust standard errors are reported in parentheses.

Table A2: Effect of Birth Rates on GDP Growth, Controlling for Lagged Birth Rates

	(1) Base	(2) + Lag birth	(3) All controls	(4) + Lag birth rates
Panel A: Countries				
birth rate (t-20)	-0.60	-0.62	-0.42	-0.37
	(0.19)	(0.20)	(0.20)	(0.19)
N	45	45	45	45
R-sq	0.650	0.656	0.757	0.762
Panel B: Commuting Zones				
birth rate (t-20)	-0.22	-0.19	-0.07	-0.04
	(0.08)	(0.07)	(0.03)	(0.03)
N	335	335	335	335
R-sq	0.232	0.248	0.813	0.823

Notes: This table reports estimates of β^h from the regression $\Delta_h \ln y_{c,t+h} = \alpha^h + \beta^h b_{c,t-20} + \rho^h b_{c,t-40} + \boldsymbol{x}'_{c,t}\boldsymbol{\gamma}^h + \varepsilon^h_{c,t}$. Panel A reports estimates in our cross-country sample, setting $h = 50$ and $t = 1970$. Panel B reports estimates across commuting zones, setting $h = 60$ and $t = 1960$. The dependent variable $y_{c,t}$ is log GDP per worker in Panel A, and log composition-adjusted weekly wage in Panel B. In all regressions, the vector $\boldsymbol{x}_{c,t}$ includes initial average income and labor supply, measured by log population aged 20–45 and log population aged 45–70. Average income is measured as GDP per working-age adult across countries and average wage per worker across commuting zones. Column 1 reports our baseline specification estimated on the set of observations with data on $t-20$ birth rates, while Column 2 adds birth rates in $t-40$ as a control. Column 3 reports our baseline estimates including all controls used in Tables 1 and 2, corresponding to columns 5 and 7 in these tables, respectively. For countries, the vector of controls includes education levels, urbanization, and region dummies. For our commuting-zone estimates, it also includes initial manufacturing and agricultural shares of employment. Column 4 adds lagged $t-40$ birth rates as a control in addition to these variables. All regressions weight countries equally and commuting zones by working-age population in year t. Heteroskedasticity-robust standard errors are reported in parentheses.

Table A3: Cross-Country Estimates Using Log $t-20$ Birth Rates

	(1)) Base	(2) Educ	(3) Urban	(4) Region	(5) All	(6) FE
	Panel A: Long Differences					
birth rate (t-20, log)	-0.51	-0.28	-0.47	-0.41	0.14	
	(0.34)	(0.32)	(0.32)	(0.38)	(0.37)	
N	107	107	107	107	107	
R-sq	0.483	0.510	0.496	0.571	0.622	
	Panel B: Stacked Differences					
birth rate (t-20, log)	-0.48	-0.35	-0.47	-0.31	-0.12	0.06
	(0.15)	(0.15)	(0.15)	(0.16)	(0.17)	(0.21)
N	229	229	229	229	229	240
R-sq	0.380	0.394	0.384	0.430	0.452	0.838

Notes: This table reports estimates of β^h from the regression $\Delta_h y_{c,t+h} = \alpha^h + \beta^h \ln b_{c,t-20} + \boldsymbol{x}'_{c,t}\boldsymbol{\gamma}^h + \varepsilon^h_{c,t}$ in our cross-country dataset. The dependent variable is $\Delta_h y_{c,t+h}$, the change in log GDP per worker over horizon h. Panel A reports estimates setting $h = 50$ and $t = 1970$. Panel B reports estimates setting $h = 25$ and $t \in \{1970, 1995\}$. The vector $\boldsymbol{x}_{c,t}$ includes GDP per worker and labor supply, measured by log population aged 20–45 and log population aged 45–70. Column 1 reports our baseline specification with no additional controls beyond $\boldsymbol{x}_{c,t}$. Column 2 additionally controls for average years of schooling among the population aged 45–55 in year $t-20$. Column 3 controls for urbanization, measured by the share of people residing in urban areas in t. Column 4 controls for continent fixed effects, while Column 5 adds all these controls simultaneously. All regressions weight countries equally. Standard errors, clustered by country, are reported in parentheses.

Table A4: Effect of $t-20$ Birth Rates on Aggregate GDP and Earnings

	Base		Educ		Rural		Region	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Panel A: Cross Country							
birth rate (t-20)	0.02	-0.13	0.05	-0.10	0.03	-0.12	0.05	-0.07
	(0.14)	(0.06)	(0.14)	(0.07)	(0.14)	(0.06)	(0.15)	(0.07)
N	107	229	107	229	107	229	107	229
R-sq	0.259	0.206	0.263	0.211	0.265	0.212	0.421	0.272
	Panel B: Commuting Zones							
birth rate (t-20)	0.03	0.10	0.03	0.11	-0.06	0.13	-0.09	0.13
	(0.13)	(0.14)	(0.11)	(0.14)	(0.12)	(0.14)	(0.11)	(0.16)
N	722	1,438	722	1,438	722	1,438	722	1,438
R-sq	0.314	0.298	0.416	0.300	0.367	0.312	0.445	0.319

Notes: This table reports estimates of β^h in equation (10). Panel A reports estimates in our cross-country dataset, while Panel B reports estimates in our cross-commuting-zone dataset. In Panel A the dependent variable $y_{c,t}$ is aggregate real GDP, while in Panel B $y_{c,t}$ is the aggregate earnings of households in a commuting zone. Odd-numbered columns report estimates using long-difference specifications, which set $h = 50$ and $t = 1970$ in Panel A, and set $h = 60$ and $t = 1960$ in Panel B. Even-numbered columns report estimates from stacked-differences specifications, which set $h = 25$ and $t \in \{1970, 1995\}$ in Panel A and $h = 30$ and $t \in \{1960, 1990\}$ in Panel B. The vector $\boldsymbol{x}_{c,t}$ includes initial average income, manufacturing share of employment, and labor supply, measured by log population aged 20–45 and log population aged 45–70. Average income is measured as GDP per working-age adult across countries and average wage per worker across commuting zones. All regressions weight countries equally and commuting zones by working-age population in year t. Standard errors, clustered by geographic unit c, are reported in parentheses.

Table A5: Commuting-Zone Estimates Using Log $t-20$ Birth Rates

	(1)) Base	(2) Educ	(3) Urban	(4) Region	(5) Manuf	(6) Agri	(7) All	(8) FE
				Panel A: Long Differences				
birth rate (t-20, log)	-0.34	-0.28	-0.35	-0.30	-0.38	-0.33	-0.14	
	(0.10)	(0.07)	(0.10)	(0.07)	(0.08)	(0.09)	(0.04)	
N	722	722	722	722	722	722	722	
R-sq	0.187	0.448	0.190	0.465	0.364	0.223	0.695	
				Panel B: Stacked Differences				
birth rate (t-20, log)	-0.23	-0.23	-0.23	-0.19	-0.25	-0.24	-0.16	-0.16
	(0.04)	(0.03)	(0.04)	(0.03)	(0.04)	(0.04)	(0.02)	(0.04)
N	1,438	1,438	1,438	1,438	1,438	1,438	1,438	1,432
R-sq	0.957	0.965	0.957	0.963	0.960	0.957	0.973	0.986

Notes: This table reports estimates of β^h from the regression $\Delta_h y_{c,t+h} = \alpha^h + \beta^h \ln b_{c,t-20} + \boldsymbol{x}'_{c,t}\boldsymbol{\gamma}^h + \varepsilon^h_{c,t}$ in our cross-commuting-zone dataset. The dependent variable is $\Delta_h y_{c,t+h}$, the change in log composition-adjusted weekly wage over horizon h. Panel A reports estimates setting $h = 60$ and $t = 1960$. Panel B reports estimates setting $h = 30$ and $t \in \{1960, 1990\}$. The vector $\boldsymbol{x}_{c,t}$ includes log average wage per worker and labor supply, measured by log population aged 20–45 and log population aged 45–70. Column 1 reports our baseline specification with no additional controls beyond $\boldsymbol{x}_{c,t}$. Column 2 controls for the share of adults with a college education in $t-20$. Column 3 controls for urbanization, measured by population density in year t. Column 4 controls for Census Region fixed effects. Column 5 controls for the initial manufacturing share of employment in t. Column 6 controls for the initial agricultural share of employment. Column 7 includes all the aforementioned controls simultaneously. Column 8 includes commuting-zone fixed effects. All regressions weight commuting zones by working-age population in year t. Standard errors, clustered by commuting zone, are reported in parentheses.

Table A6: Commuting-Zone Estimates Using Alternative Wage Measures, 1980–2020

	(1) Composition Adj	(2) Avg Annual Wage	(3) Wage per Working-Age	(4) Income per Capita	(5) Aggregate Earnings
birth rate (t-20)	-0.113	-0.105	-0.102	-0.053	-0.066
	(0.029)	(0.029)	(0.035)	(0.029)	(0.075)
N	720	720	719	719	720
R-sq	0.209	0.215	0.166	0.119	0.274

Notes: This table reports estimates of β^h from equation (10) in our cross-commuting-zone dataset. The dependent variable is $\Delta_h y_{c,t+h}$, where $y_{c,t}$ varies by column. Column 1 reports estimates where $y_{c,t}$ is our baseline measure, the log composition-adjusted wage rate. Column 2 reports estimates where $y_{c,t}$ is the log average annual wage per employed worker as observed in the PUMS samples. Column 3 reports estimates where $y_{c,t}$ is the log average wage and salary income per working-age adult, while Column 4 reports estimates where $y_{c,t}$ is the log average income per capita in a commuting zone. Column 5 reports estimates where $y_{c,t}$ is log aggregate (composition-adjusted) earnings. The column 3 and column 4 measures are available in the STF only after 1980, so all regressions set $h = 40$ and $t = 1980$. The vector $\boldsymbol{x}_{c,t}$ includes log average wage per worker and labor supply, measured by log population aged 20–45 and log population aged 45–70. All regressions weight commuting zones by working-age population in year t. Heteroskedasticity-robust standard errors are reported in parentheses.

Table A7: Mechanisms I: Effect of Birth Rates on Sectoral Composition, Trade, and Productivity

	Country data		CZ data			Industry data
	Hightech exports	TFP	Hightech emp	Labor-int emp	RD-int emp	TFP
birth rate (t-20)	-0.123	-0.239	-0.006	-0.014	0.020	-215.549
	(0.024)	(0.074)	(0.006)	(0.014)	(0.008)	(300.299)
Observations	74	96	716	716	716	59

Notes: This table reports estimates of β^h in equation (10). In Column 1, we set $t = 1995$ and $h = 25$. All other columns set $t = 1990$ and $h = 30$. Columns 1 and 2 report estimates where the dependent variable $y_{c,t}$ is the share of exports in high-tech products and TFP, both measured across countries. Columns 3, 4, and 5 report estimates where $y_{c,t}$ is the commuting-zone-level employment share in high-tech, labor-intensive, and R&D-intensive industries, respectively. Column 6 reports estimates where $y_{c,t}$ is TFP across industries. The vector $\boldsymbol{x}_{c,t}$ includes initial average income and labor supply, measured by log population aged 20–45 and log population aged 45–70. Average income is measured as GDP per working-age adult across countries and average wage per worker across commuting zones. All regressions weight countries equally, commuting zones by working-age population in year t, and industries by employment in 1990. Heteroskedasticity-robust standard errors are reported in parentheses.

Table A8: Mechanisms II: Effect of Birth Rates on Labor-Saving Innovation

	(1)	(2)	(3)	(4)	(5)	(6)
	Countries			Commuting Zones		
	Labor intensive industries	ICT patents	Automation patents	Labor intensive industries	ICT patents	Automation patents
birth rate (t-20)	-0.11	-0.10	-0.07	-0.07	-0.14	-0.07
	(0.07)	(0.04)	(0.09)	(0.06)	(0.05)	(0.05)
N	230	230	230	680	680	680
R-sq	0.056	0.145	0.087	0.016	0.127	0.037

Notes: This table reports estimates of β^h in equation (10), where $h = 20$ and $t = 1990$. Columns (1)–(3) report estimates in our cross-country dataset, while Columns (4)–(6) report estimates in our cross-commuting-zone dataset. Columns 1 and 4 report estimates where the dependent variable $y_{c,t}$ is the share of new patents (citation-weighted) filed in labor-intensive industries. Columns 2 and 5 report estimates where $y_{c,t}$ is the share of new patents filed in the ICT category. Columns 3 and 6 report estimates where $y_{c,t}$ is the share of new patents that are automation patents, where automation patents are defined from the metrics created by Autor et al. (2024). The vector $\boldsymbol{x}_{c,t}$ includes initial average income, manufacturing share of employment, and labor supply, measured by log population aged 20–45 and log population aged 45–70. Average income is measured as GDP per working-age adult across countries and average wage per worker across commuting zones. All regressions weight countries equally and commuting zones by working-age population in year t. Heteroskedasticity-robust standard errors are reported in parentheses.

Table A9: Alternative Mechanisms: Effect of Birth Rates on Female LFP, Education, and Sectoral Composition

	Countries				Commuting Zones			
	(1) Female LFP	(2) Yrs Schooling	(3) Agri Emp	(4) Manuf Emp	(5) Female LFP	(6) College Share	(7) Agri Emp	(8) Manuf Emp
birth rate (t-20)	11.129	0.767	0.063	0.024	-1.305	-2.074	0.001	-0.016
	(4.941)	(0.241)	(0.061)	(0.020)	(0.736)	(1.149)	(0.004)	(0.010)
Observations	38	108	42	42	715	715	716	716

Notes: This table reports estimates of β^h in equation (10), for $h = 30$. Columns (1)–(4) report estimates in our cross-country dataset, while Columns (5)–(8) report estimates in our cross-commuting-zone dataset. For Column 2 we set $t = 1980$; for all other columns we set $t = 1990$. The earlier t for Column 2 reflects that cross-country years-of-schooling by age cohort is available only up to 2015. The dependent variable is $\Delta_h y_{c,t+h}$, where $y_{c,t}$ is the female labor force participation rate in Columns 1 and 5; average years of schooling for those aged 25–45 in Column 2; the share aged 25–45 with a college degree in Column 6; the share of employment in agriculture in Columns 3 and 7; and the share of employment in manufacturing in Columns 4 and 8. The vector $\boldsymbol{x}_{c,t}$ includes initial average income, manufacturing share of employment, and labor supply, measured by log population aged 20–45 and log population aged 45–70. Average income is measured as GDP per working-age adult across countries and average wage per worker across commuting zones. All regressions weight countries equally and commuting zones by working-age population in year t. Heteroskedasticity-robust standard errors are reported in parentheses.

Figure A1: Birth Rates and Birth Rate Persistence across Countries, 1930–2020

A. Global birth rates, 1930–2020

B. Birth rate persistence, 1930–2010

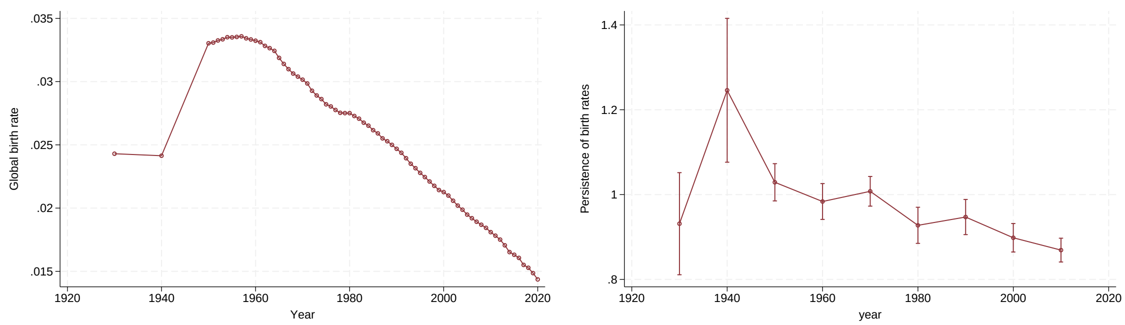

Notes: Panel A reports the global aggregate birth rate (equivalent to the population weighted average) across 46 countries with available births data in 1940. 1930 and 1940 birth rates are from the UN Demographic Yearbooks (Acemoglu and Johnson, 2007), while post-1950 data are from the UN Population Prospects (United Nations, Department of Economic and Social Affairs, Population Division, 2024). Panel B reports estimates of ρ^t from the equation $b_{c,t+10} = \delta_t + \rho^t b_{c,t} + \varepsilon_{c,t}$ in a given cross section t, for decades $t = 1930$ through $t = 2010$, the latest t for which $b_{c,t+10}$ is observed. 95% confidence intervals are also reported, where standard errors are robust to heteroskedasticity.

Figure A2: Cross-Country Impulse Response of Income to Birth Rates in $t-20$, Including 1930 and 1940

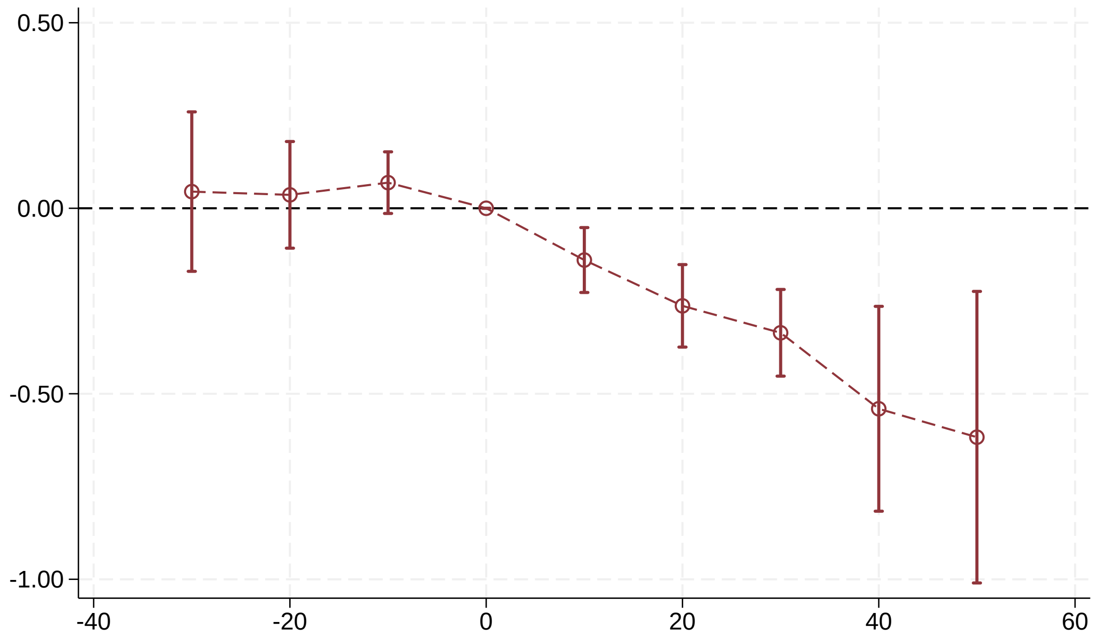

Notes: This figure reports estimates of β^h in equation (10) alongside 95% confidence intervals, estimated in our cross-country dataset. The dependent variable is $\Delta_h y_{c,t+h}$, the change in log GDP per worker over horizon h. The vector $\boldsymbol{x}_{c,t}$ includes initial GDP per working-age adult and labor supply, measured by log population aged 20–45 and log population aged 45–70. Birth rate data come from the UN Population Prospects after 1950 and from the UN Demographic Yearbooks for 1930 and 1940. All regressions weight countries equally. Standard errors are clustered by country.

Figure A3: Impulse Response of Wages to Birth Rates in $t-20$, Large Commuting Zones

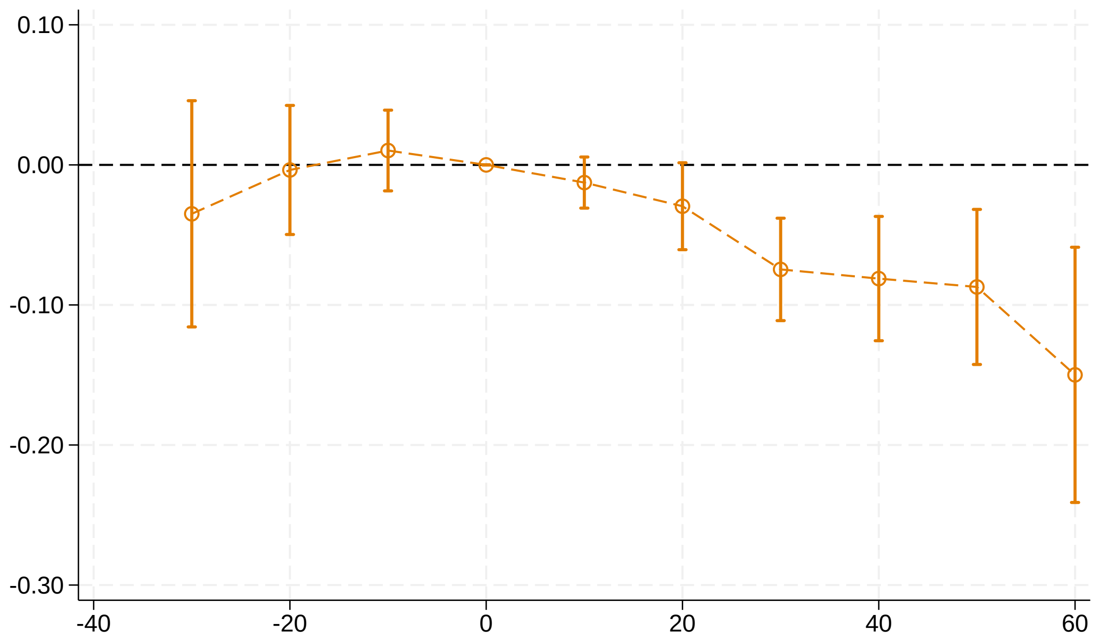

Notes: This figure reports estimates of β^h in equation (10) alongside 95% confidence intervals. All regressions are estimated on the sample of 137 commuting zones with population greater than 300,000 in 1970, and weight these commuting zones equally. The dependent variable $y_{c,t}$ is the log composition-adjusted weekly wage, while the vector of controls $\boldsymbol{x}_{c,t}$ includes initial log average wages and labor supply, measured by log population aged 20–45 and log population aged 45–70. Standard errors are clustered by commuting zone.

Figure A4: Effect of Birth Rates on Native-Born Populations in Commuting Zones

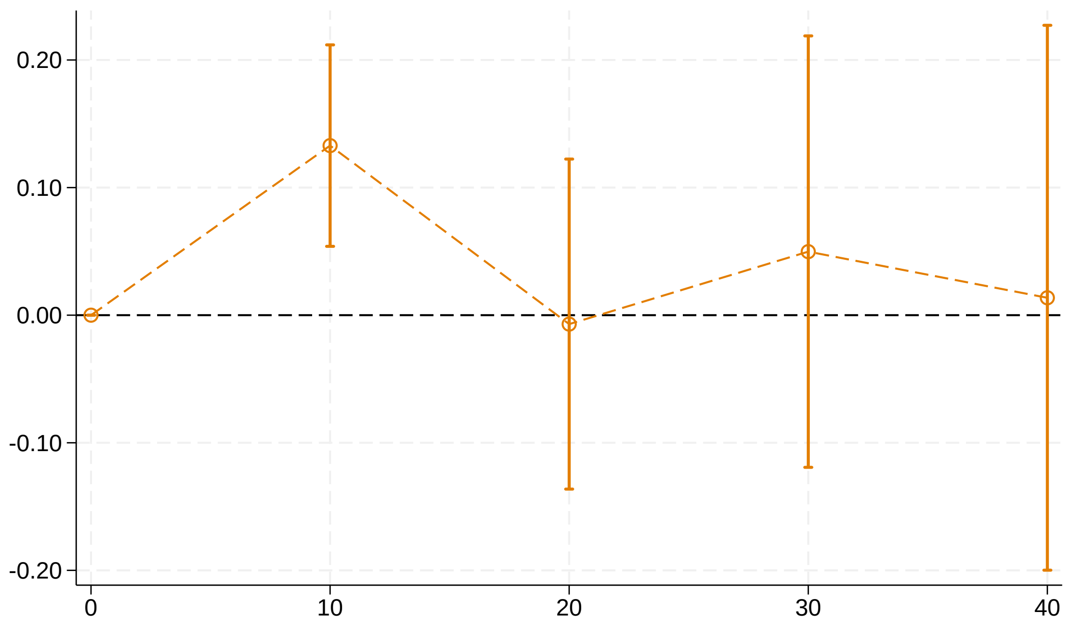

Notes: This figure reports estimates of β^h in equation (10) with 95% confidence intervals. The dependent variable $y_{c,t}$ is our estimate of the native-born working-age population in a commuting zone. We estimate the number of native-born residents by applying national age-specific death rates to births in each commuting zone over time. The vector $\boldsymbol{x}_{c,t}$ includes average wage per worker and labor supply, measured by log population aged 20–45 and log population aged 45–70. We weight commuting zones by working-age population in year t. Standard errors are clustered by commuting zone.

B Derivations and Proofs

In this part of the Appendix, we provide proofs and derivations omitted from the text.

Omitted Proofs

Proof of Lemma 1 Aggregate output (GDP) in this economy is given by Y, while net output, which is also equal to consumption, subtracts the costs of machine production and development, and is thus $NY = Y - \gamma q(\theta) - \Gamma\theta$. Writing this as a function of θ for emphasis, we have

$$\begin{aligned} NY(\theta) &= Y(\theta) - \gamma q(\theta) - \Gamma\theta \\ &= (1-\alpha^2)Y(\theta) - \Gamma\theta. \end{aligned} \tag{B-1}$$

We first prove that an increase in θ increases GDP and net output. From the first-order condition (5), we have $\alpha(1-\alpha)\frac{\partial Y(\theta^*)}{\partial\theta} = \Gamma > 0$, so $\frac{\partial Y(\theta^*)}{\partial\theta} > 0$. We next show that net output also increases. From (B-1), the derivative of net output with respect to θ is $\frac{\partial NY(\theta)}{\partial\theta} = (1-\alpha^2)\frac{\partial Y(\theta)}{\partial\theta} - \Gamma$. From the first-order condition, $\alpha(1-\alpha)\frac{\partial Y(\theta^*)}{\partial\theta} = \Gamma$ implies $(1-\alpha)\frac{\partial Y(\theta^*)}{\partial\theta} = \frac{\Gamma}{\alpha}$. Therefore,

$$\frac{\partial NY(\theta^*)}{\partial\theta} = (1+\alpha)(1-\alpha)\frac{\partial Y(\theta)}{\partial\theta} - \Gamma = (1+\alpha)\frac{\Gamma}{\alpha} - \Gamma = \frac{\Gamma}{\alpha} > 0,$$

proving the lemma.

Proof of Lemma 2 An application of the implicit function theorem to (5) implies:

$$\frac{d\theta^*}{dL} = -\frac{\frac{\partial^2 F}{\partial\theta\partial L}}{\frac{\partial^2 F}{\partial\theta^2}} < 0. \tag{B-2}$$

Concavity ensures that $\frac{\partial^2 F}{\partial\theta^2} < 0$ and recalling the definition of $\varepsilon_{L\theta}$ establishes the desired result.

Automation as Labor-Saving Technology

In this section, we show that the conditions for Lemma 2 are satisfied in a task model with automation (Autor et al., 2003; Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018, 2022a,b). Suppose output is produced using the following constant returns to scale production technology:

$$Y = \left(\int_0^1 x(i)^\beta di\right) H^{1-\beta}. \tag{B-3}$$

Here $x(i)$, for $i \in [0,1]$, denotes the quantity of task i, which can be performed by either capital or unskilled younger labor L, and H is older/skilled labor (used for management or entrepreneurship). The parameter $\beta \in (0,1)$ determines the degree of substitution between different tasks within the aggregator in parentheses in (B-3).

The state of technology is represented by θ and corresponds to the degree of automation in this model. Specifically, let $k(i)$ and $l(i)$ denote the amount of capital and labor allocated to task i. Then, tasks $i \leq \theta$ are automated and are produced using capital with the production function $x(i) = k(i)$, while tasks $i > \theta$ must be produced using labor, so that $x(i) = l(i)$.[32] Market-clearing

[32]Note that we are assuming that θ fully specifies which tasks have to be produced by capital, rather than allowing

for capital and younger workers imposes:

$$\int_0^{\theta} k(i)di \leq K \text{ and } \int_{\theta}^{1} l(i)di \leq L.$$

Market-clearing for older workers is imposed by the production function (B-3), which was written in terms of the full supply of H. Our formulation implies that greater θ corresponds to more automation (more labor-saving technology).

Given our specification and the symmetry of tasks in (B-3), market-clearing imposes:

$$x(i) = \begin{cases} \frac{K}{\theta} & \text{for } i \leq \theta \\ \frac{L}{1-\theta} & \text{for } i > \theta \end{cases}.$$

Substituting these task quantities, the value added production function is

$$F(L, K, H, \theta) = \left[\theta\left(\frac{K}{\theta}\right)^{\beta} + (1-\theta)\left(\frac{L}{1-\theta}\right)^{\beta}\right] H^{1-\beta}$$

Differentiating this expression with respect to L, we obtain

$$\frac{\partial F}{\partial L} = \beta\left(\frac{L}{1-\theta}\right)^{\beta-1} H^{1-\beta}$$

Because $\beta < 1$, it is immediate that $\frac{\partial F}{\partial L}$ is decreasing in L (the marginal product of labor is decreasing) and is also decreasing in θ, as desired. We therefore conclude that $\frac{\partial^2 F}{\partial \theta \partial L} < 0$, and so automation represents labor-saving technical change.

B.1 Deriving Our Empirical Specification

As stated in the text, our main specification separates the working-age population into younger and older workers. A direct generalization of (6) is then

$$\begin{aligned} \Delta_h y_{c,t+h} = & \text{time effect}_{t+h}^h + \beta_L^h \Delta_h \log L_{c,t+h} + \beta_H^h \Delta_h \log H_{c,t+h} \\ & + \mathbf{x}_{c,t}' \gamma^h + \rho^h y_{c,t} + \pi^h \log L_{c,t} + \mu^h \log H_{c,t} + \eta_{c,t+h}^h, \end{aligned} \tag{B-4}$$

where $L_{c,t}$ is younger population in country c at time t, while $H_{c,t}$ denotes older population. Changes in both populations can now impact GDP per worker, and as in the previous subsection, their initial levels are included to allow for error-correction dynamics.

The evolution of the younger population can be written in a similar form to (7):

$$L_{c,t+h} - L_{c,t} = \text{surviving births}_{c,t-n}^h - \text{transitions}_{c,t}^h,$$

where transitions replace exits (most departures from the younger group are entries into the older group). Dividing by $L_{c,t}$ as before, this yields $\Delta_h \log L_{c,t+h} = \text{birth rate}_{c,t-n} - \mathrm{j}_{t-1}^h$, where j_{t-1}^h is

these tasks to be produced by capital or labor. This assumption is adopted for expositional simplicity and is relaxed in Acemoglu and Restrepo (2018, 2022a) and Acemoglu et al. (2025). The condition $\theta/(1-\theta) < K/L$ is sufficient for this allocation to be the unique equilibrium and is required for more automation to increase output (see below). Autor and Kausik (2026) prove the general version of this result: a falling labor share raises output when the capital-labor ratio is sufficiently high.

the transition rate into the older-age group. The dynamics of the older-age population are given by $H_{c,t+h} - H_{c,t} = \text{transitions}^h_{c,t} - \text{exits}^h_{c,t}$,where transitions out of the younger group are approximately equal to entries into the older group, since deaths at these ages are small. Therefore, we have:

$$\Delta_h \log H_{c,t+h} \simeq \text{transition rate}^h_{c,t} - \mathsf{d}^h_{c,t},$$

where $\text{transition rate}^h_{c,t} = \text{transitions}^h_{c,t}/H_{c,t}$ and $\mathsf{d}^h_{c,t} = \text{exits}^h_{c,t}/H_{c,t}$. The transition rate is in turn related to surviving births from m periods ago. Substituting these relationships into equation (B-4), we obtain:[33]

$$\begin{aligned}\Delta_h y_{c,t+h} =& \text{time effect}^h_{t+h} + \beta^h_L \text{ birth rate}_{c,t-n} + \beta^h_H \text{ birth rate}_{c,t-m} \\ &+ \mathbf{x}'_{c,t}\gamma^h + \rho^h\, y_{c,t} + \pi^h \log L_{c,t} + \mu^h \log H_{c,t} + \varepsilon^h_{c,t+h}.\end{aligned}$$

When birth rate changes are persistent, we can further simplify by imposing $\text{birth rate}_{c,t-n} \simeq \text{birth rate}_{c,t-m}$, which gives us our baseline estimating equation, equation (9) in the text.

This equation excludes additional lags of birth rates from the left-hand side. Including them would shorten the panel dimension of our estimation sample, and we test robustness to this restriction in two ways; first, we estimate pre-trend coefficients for negative h and find no predictability of treatment for past outcomes; second, we estimate this equation adding birth rates as a control. Both checks suggest our exclusion of birth rates from the left-hand side is justified.

Why Birth Rates Can Be Treated as Exogenous

This section discusses why the causal effect of birth rates on subsequent wages is identified even when fertility decisions are endogenous. We consider a simple economic environment for fertility decisions that motivates this argument.

Suppose that fertility decisions depend on expectations of future wages and other economic variables. For the first period of life, newborns do not work. Suppose also that wages are a function of the current working-age population and a productivity term:

$$w_{c,t} = W\left(P_{c,t}, z_{c,t}\right) \tag{B-5}$$

where c refers to country or commuting zone, but for concreteness, we focus on commuting zones in this part of the Appendix. We refer to $z_{c,t}$ as "productivity" and assume that it follows a Markov process. In addition, $P_{c,t} = [p^a_{c,t}]_{a\geq 1}$ denotes the working-age population distribution by age.

Population dynamics can be summarized by the McKendrick–von Foerster equation (McKendrick, 1926; von Foerster, 1959):

$$\frac{\partial}{\partial t} p^a_{c,t} + \frac{\partial}{\partial a} p^a_{c,t} = -d^a_{c,t} p^a_{c,t}. \tag{B-6}$$

Births enter as a boundary condition, $p^0_{c,t} = b_{c,t}$.[34] Given an initial age distribution and a sequence of birth rates, this recursion pins down the entire evolution of population.

The partial differential equation in (B-6) enables us to write population as a function of initial

[33] An additional approximation used here is that $\text{birth rate}_{c,t-m} = \text{births}_{c,t-m}/P_{c,t-m} \simeq \text{constant}\times\text{births}_{c,t-m}/H_{c,t}$. Without this approximation, one might need to include additional lags of older-age population or further past birth rates on the right-hand side.

[34] In discrete-time form, this rate-of-change equation corresponds to the recursion $p^{a+1}_{c,t+1} = (1 - d^a_{c,t})\, p^a_{c,t}$: the population aged $a+1$ at time $t+1$ equals the survivors from the cohort that was aged a at time t.

population and the subsequent history of birth rates and productivity shocks as:

$$P_{c,t+h} = \mathcal{P}\left(\boldsymbol{b}_{c,t-1}, \boldsymbol{z}_{c,t}, P_{c,t-1}\right)$$

Substituting this into (B-5), wages can be written as a function of initial working-age population, the birth rates of cohorts who enter the labor force, and productivity:

$$w_{c,t+h} = \mathcal{W}\left(\boldsymbol{b}_{c,t-1}, \boldsymbol{z}_{c,t}, P_{c,t-1}\right)$$

Following Bojinov et al. (2021), we can define the average structural function as

$$\Psi^h(b_{c,t-1}) = \mathbb{E}_{c,t}\left[\mathcal{W}\left(\boldsymbol{b}_{c,t-1}, \boldsymbol{z}_{c,t}, P_{c,t-1}\right)|\Omega_{c,t-1}\right]$$

where $\Omega_{c,t-1}$ is the information set at time $t-1$ (which includes $b_{c,t-1}$), and $\mathbb{E}_{c,t}[\cdot]$ averages across cross-sectional units c and time periods t.[35] Because wages follow a Markov process conditional on $z_{c,t-1}$ and $P_{c,t-1}$, and fertility decisions are made on the basis of the information set $\Omega_{c,t-1}$, future productivity shocks $z_{c,t+h}$ are orthogonal to birth rates at $t-1$. Hence birth rates satisfy the orthogonality condition $b_{c,t-1} \perp z_{c,t+h} \big| w_{c,t-1}, P_{c,t-1}$. By the Markov property, conditioning on $(b_{c,t-1}, w_{c,t-1}, P_{c,t-1})$ summarizes the relevant information in $\Omega_{c,t-1}$ for predicting $w_{c,t+h}$. The conditional mean we observe in the data is thus:

$$\mathbb{E}\left[w_{c,t+h} \big| b_{c,t-1}, w_{c,t-1}, P_{c,t-1}\right] = \mathbb{E}\left[w_{c,t+h} \big| \Omega_{c,t-1}\right] = \Psi^h(b_{c,t-1}). \tag{B-7}$$

This implies that the conditional mean function of wages identifies causal effects of birth rates on wages, even when fertility decisions are endogenous.

[35] The expectation integrates out future birth rates and productivity shocks, so Ψ^h depends only on the scalar $b_{c,t-1}$ rather than the full sequence $\boldsymbol{b}_{c,t-1}$.

C Consistency with Other Estimates

This section considers the findings in Maestas et al. (2023) and Hayashi (2025), two recent studies presenting evidence that aging is associated with slower growth across US states and countries, respectively. These results are, at first blush, difficult to reconcile with our findings that lower birth rates—which reduce population growth and cause aging—are associated with more rapid economic growth. We show below that the results reported in both papers can be reconciled with our findings, after accounting for issues of robustness and sensitivity to specification choices.

Revisiting Maestas et al. (2023)

Maestas et al. (2023) report state-level regressions of ten-year economic growth on aging, measured by the log share of individuals aged 60+ in the population. Because of the endogeneity of age composition in the presence of between-state migration, they instrument this measure using lagged demographic-based predicted aging. Identifying the implications of aging and changes in population size associated with aging is challenging using data from only 50 states (plus DC). For this reason, much recent literature, including the current paper, focuses on variation across 722 commuting zones that arguably better approximate local US labor markets than US states.

Maestas et al. (2023) use all states, including Alaska and Hawaii, plus DC, and stack three ten-year panels. They focus throughout on regressions that control for a full set of industry-times-year controls.

In Appendix Table C-1 we revisit their results. Panel A replicates their findings exactly. Column 1 reports the OLS estimate. Column 2 reports the reduced-form regression of ten-year economic growth on their predicted aging instrument. Column 3 reports the first-stage relationship between the predicted aging instrument and the log share of the population aged 60+. Column 4 reports the full two-stage least squares (2SLS) estimate. The reduced-form coefficient in column 2 and the 2SLS estimate in column 4 are both large and statistically significant, indicating that aging is associated with slower growth.

Panel B explores the same relationship, but without including the industry-times-year controls, since it is useful to start with a parsimonious specification and because such controls may be excessive in a sample with 153 observations (50 states plus DC times three ten-year periods). Dropping the industry-times-year controls substantially affects the Maestas et al. estimates: both the reduced-form and the 2SLS estimates are close to zero with t-statistics well below conventional significance thresholds (the first stage, unsurprisingly, is still strong, since it is based on demographic accounting). These estimates indicate that their results depend on the inclusion of industry-times-year controls.

Panel C reverts to their specification with industry-times-year controls, but drops Alaska, Hawaii, and DC, consistent with much literature that conducts state-level analysis using only the 48 contiguous states. The impact of this change in sample is less dramatic but also substantial: both reduced-form and 2SLS estimates become statistically indistinguishable from zero. Finally, Panel D drops the industry-times-year controls in the sample of 48 contiguous states. Here, the reduced-form and the 2SLS estimates have the opposite sign from the three earlier panels and are far from statistical significance. The non-robustness of the Maestas et al. (2023) results may be due to limited power stemming from a sample of only 50 US states and DC.

A complementary question is whether our main results are confounded by or sensitive to the inclusion of Maestas et al.'s key variables: aging, predicted aging, and industry-times-year controls. We explore this question in Appendix Table C-2.

Panel A of the table reports our main stacked-differences specification for commuting zones, repeated from Panel B of Table 2.

Panel B then includes industry-times-year controls, following Maestas et al. (2023). These controls cause some attenuation in the coefficients of interest, but the results remain economically large and statistically significant. The only exception is column 8 with commuting-zone fixed effects, which, as noted in the text, is a demanding specification even without these additional controls. Here, the point estimate remains negative but loses significance.

Panel C adds Maestas et al.'s aging instrument to Panel A (without the industry-times-year controls).[36] In this specification, the coefficient on the birth-rate variable is essentially identical to the corresponding estimates in Panel A, while the predicted aging variable has the opposite sign from Maestas et al.'s finding (suggesting that faster aging is associated with greater economic growth).

Finally, Panel D adds industry-times-year controls to the Panel C estimates. The coefficients for the birth-rate variable are almost identical to those in Panel B, while the predicted aging variable continues to be associated with faster economic growth, opposite to the implications of Maestas et al.'s state-level analysis.

These findings show that our central results, driven by variation in birth rates, are robust to accounting for the key variables in Maestas et al.'s analysis.

Revisiting Hayashi (2025)

Using cross-country data, Hayashi (2025) finds that population aging decreases GDP per capita growth through its negative effect on TFP growth. Hayashi focuses on countries that have completed the demographic transition, defined as those where life expectancy is high and fertility rates are low. These are effectively the countries where the share of older people in the adult population exceeds 0.4. The cross-country growth regressions relate growth in GDP per capita to the *log older-to-younger ratio* in the population, while also including several controls on the right-hand side (in particular initial GDP per capita, average rule of law, latitude, a measure of openness, and the child dependency ratio).[37] Within this specification and the selected subset of countries, the estimated impact of the share of older people on GDP per capita growth is negative.

Panel A of Appendix Table C-3 presents our replication of these results; its first three columns parallel Table 1 of Hayashi (2025). Our coefficients differ slightly from the published values, but the significance and sign patterns remain consistent. For example, in column 2, the log older-to-younger ratio variable has a coefficient of -3.68 (standard error = 1.08), while in the original paper, this coefficient is -3.77 (standard error = 1.04). This small discrepancy may reflect data revisions in different vintages of the sources we are using. Our replication reproduces the central result: columns 2 and 3 show that for the 52 countries that completed their demographic transition by 1990, aging is associated with lower GDP per capita growth between 1990 and 2020, especially when focusing on only OECD and Asian economies in column 3. However, as Hayashi notes, this relationship becomes weaker in the full sample of countries (columns 4 and 5).

Panel B examines the robustness of Hayashi's results to the choice of specification. As explained above, Hayashi regresses a change on a level: the dependent variable is the change in (log) GDP per capita, while the regressor is the level of the log older-to-younger ratio. As in our Section 4, a more

[36]We focus on the instrument because, as Maestas et al.'s empirical design emphasizes, observed aging is endogenous. This is the logical comparison since predicted aging and birth rates are the primary identifying variables in these two papers.

[37]The child dependency ratio is itself correlated with the aging variable. We nonetheless retain Hayashi's full specification, including this and the other controls, so that our results remain directly comparable to the original.

standard approach relates, instead, the level of log GDP per capita to the level of this variable, or the change to the change.[38] Panel B shows that the negative relationship between this aging variable and GDP per capita disappears under the change-on-change specification. For example, the coefficient in column 2, which was -3.68 in Panel A, is now 0.44 (standard error = 0.69). In fact, all of the coefficient estimates for the change-in-aging variable are now positive. They are statistically insignificant in columns 1–3, but positive and significant in both full-sample columns: 1.96 (standard error = 0.82) in column 4 (1960–1990) and 1.46 (standard error = 0.5) in column 5 (1990–2020). We therefore conclude that Hayashi's evidence of sizable negative effects of aging is not fully robust.

As in the previous section of this Appendix, we next check whether our results are robust to including these variables. Appendix Table C-4 includes the older-to-younger ratio in our baseline stacked-differences specifications (as in Panel B of Table 1 in the main text), with the dependent variable now switching from GDP per capita to GDP per worker as in the rest of our paper. Panel A displays our main results for reference, Panel B adds the log older-to-younger ratio variable, and finally Panel C repeats the analysis of Panel B, restricting the sample to the post-demographic-transition countries. This table shows that including these variables, irrespective of the exact sample, has essentially no impact on our results. If anything, the specifications with all covariates and with country fixed effects (columns 5 and 6) are more precisely estimated in this post-demographic transition sample. We conclude that our results are robust to controlling for Hayashi's aging variable.

[38] Recall that our specification where the change in log GDP per worker is regressed on the birth rate was derived from an underlying specification in which the log of GDP per worker was a function of log younger and log older population.

Hayashi (2025) motivates this specification by arguing that younger workers are engaged in R&D, so that their numbers directly affect TFP growth. Evidence in Acemoglu and Restrepo (2022a) suggests the reverse, however: younger workers are more likely to be employed in manual, physical tasks that compete directly with automation technologies.

Table C-1: Replication of Maestas et al. (2023): State-Level Estimates

	(1) OLS	(2) RF	(3) FS	(4) IV
		Panel A: Replication		
Growth of share 60+ (aging)	-0.826			-0.545
	(0.140)			(0.173)
Predicted aging		-0.390	0.716	
		(0.134)	(0.054)	
N	153	153	153	153
F-stat (1st stage)			174.24	
Industry-times-year controls	X	X	X	X
		Panel B: No industry-times-year controls		
Growth of share 60+ (aging)	-0.799			-0.029
	(0.171)			(0.250)
Predicted aging		-0.015	0.502	
		(0.126)	(0.061)	
N	153	153	153	153
F-stat (1st stage)			67.52	
Industry-times-year controls				
		Panel C: Drop states AK, DC, and HI		
Growth of share 60+ (aging)	-0.786			-0.353
	(0.171)			(0.251)
Predicted aging		-0.226	0.642	
		(0.170)	(0.064)	
N	144	144	144	144
F-stat (1st stage)			100.80	
Industry-times-year controls	X	X	X	X
		Panel D: No industry-times-year controls, drop states AK, DC, HI		
Growth of share 60+ (aging)	-0.834			0.090
	(0.195)			(0.293)
Predicted aging		0.042	0.465	
		(0.134)	(0.067)	
N	144	144	144	144
F-stat (1st stage)			48.86	
Industry-times-year controls				

Notes: This table replicates Table 1 of Maestas et al. (2023) using state-level data. The outcome in columns (1), (2), and (4) is the ten-year log change in GDP per capita. Column (3) reports the first stage, regressing the change in log share aged 60+ on the predicted aging instrument. All regressions include year fixed effects and are weighted by initial-period state population. Panels A and C include initial-period industry employment shares interacted with year dummies. Panels C and D exclude Alaska, Hawaii, and DC. The F-statistic is the robust first-stage F-statistic on the excluded instrument, reported for Column 3 only. Standard errors, clustered by state, are reported in parentheses.

Table C-2: Commuting-Zone Estimates: Birth Rates and the MMP Instrument

	(1) Base	(2) Educ	(3) Density	(4) Region	(5) Manuf	(6) Agri	(7) All	(8) FE
	Panel A: Birth rate, no industry-times-year controls							
Birth rate (t-20)	-0.114	-0.115	-0.112	-0.097	-0.126	-0.117	-0.078	-0.089
	(0.019)	(0.016)	(0.019)	(0.016)	(0.019)	(0.019)	(0.012)	(0.023)
N	1438	1438	1438	1438	1438	1438	1438	1432
R-sq	0.956	0.965	0.956	0.962	0.959	0.957	0.973	0.986
Industry-times-year controls								
	Panel B: Birth rate, with industry-times-year controls							
Birth rate (t-20)	-0.071	-0.066	-0.068	-0.059	-0.072	-0.069	-0.027	-0.019
	(0.015)	(0.013)	(0.014)	(0.012)	(0.015)	(0.015)	(0.012)	(0.017)
N	1434	1434	1434	1434	1434	1434	1434	1426
R-sq	0.969	0.973	0.970	0.974	0.969	0.970	0.979	0.990
Industry-times-year controls	X	X	X	X	X	X	X	X
	Panel C: Birth rate and predicted aging, no industry-year controls							
Birth rate (t-20)	-0.116	-0.117	-0.115	-0.090	-0.128	-0.120	-0.076	-0.082
	(0.019)	(0.016)	(0.019)	(0.015)	(0.018)	(0.019)	(0.012)	(0.021)
Predicted aging	0.208	0.160	0.205	0.224	0.206	0.229	0.177	0.170
	(0.049)	(0.038)	(0.047)	(0.047)	(0.045)	(0.051)	(0.032)	(0.055)
N	1438	1438	1438	1438	1438	1438	1438	1432
R-sq	0.959	0.966	0.959	0.965	0.962	0.960	0.974	0.987
Industry-times-year controls								
	Panel D: Birth rate and predicted aging, with industry-year controls							
Birth rate (t-20)	-0.071	-0.067	-0.066	-0.057	-0.072	-0.068	-0.028	-0.021
	(0.014)	(0.013)	(0.013)	(0.011)	(0.014)	(0.014)	(0.011)	(0.017)
Predicted aging	0.231	0.161	0.244	0.196	0.232	0.228	0.130	0.035
	(0.044)	(0.037)	(0.044)	(0.035)	(0.043)	(0.042)	(0.032)	(0.061)
N	1434	1434	1434	1434	1434	1434	1434	1426
R-sq	0.971	0.974	0.972	0.975	0.972	0.972	0.980	0.990
Industry-times-year controls	X	X	X	X	X	X	X	X

Notes: This table reports estimates of β^h in equation (10), estimated in our cross-commuting-zone dataset. The dependent variable is $\Delta_h y_{c,t+h}$, the change in log composition-adjusted weekly wage over horizon h. All panels set $h = 30$ and restrict the sample to two non-overlapping cross sections, where $t \in \{1960, 1990\}$. All panels report birth-rate coefficients. Panels B and D include initial industry employment shares interacted with year dummies. Panels C and D include Maestas et al.'s predicted aging instrument as a regressor. All regressions weight commuting zones by working-age population in year t. The vector $\boldsymbol{x}_{c,t}$ includes log average wages and labor supply in year t, measured by log population aged 20–45 and log population aged 45–70. Column 2 controls for the share of adults with a college education in $t-20$. Column 3 controls for initial population density in year t. Column 4 controls for Census Region fixed effects. Column 5 controls for the initial manufacturing share of employment in t, while Column 6 controls for the initial agricultural share of employment. Column 7 includes all of the aforementioned controls simultaneously. Column 8 includes commuting-zone fixed effects. Standard errors, clustered by commuting zone, are reported in parentheses.

Table C-3: Cross-Country Growth Regressions

	(1) Early-transitioners 1960-1990	(2) Post-transitioners 1990-2020	(3) OECD, Asian Tigers, & China 1990-2020	(4) Full 1960-1990	(5) Full 1990-2020
	Panel A: Average of Log Older-to-younger Ratio				
Log GDP/capita (initial)	-2.34	-2.03	-1.29	-0.89	-1.36
	(0.51)	(0.29)	(0.39)	(0.20)	(0.18)
Log Older-to-younger Ratio (avg)	-1.85	-3.68	-4.42	-2.02	-1.34
	(1.52)	(1.08)	(1.34)	(1.17)	(0.52)
Rule of Law (avg)	1.10	0.86	0.16	1.06	0.74
	(0.62)	(0.19)	(0.37)	(0.25)	(0.18)
Log Openness (avg)	0.09	0.17	0.12	0.60	-0.18
	(0.36)	(0.23)	(0.23)	(0.19)	(0.21)
Latitude	1.13	2.34	2.43	0.38	0.98
	(2.15)	(1.03)	(1.34)	(1.29)	(0.91)
Log Child Dependency Ratio (avg)	-2.57	-2.22	-0.86	-2.60	-2.54
	(1.15)	(1.09)	(1.47)	(1.11)	(0.59)
N	25	52	37	110	165
R-sq	0.669	0.724	0.812	0.468	0.299
	Panel B: Change in Log Old-to-young Ratio				
Log GDP/capita (initial)	-2.26	-1.91	-1.02	-0.78	-1.16
	(0.64)	(0.34)	(0.48)	(0.20)	(0.16)
Change in Log Older-to-younger Ratio	0.42	0.44	1.15	1.96	1.46
	(1.11)	(0.69)	(0.78)	(0.82)	(0.50)
Rule of Law (avg)	1.14	0.69	-0.23	1.09	0.69
	(0.65)	(0.22)	(0.43)	(0.25)	(0.18)
Log Openness (avg)	0.07	0.58	0.62	0.63	-0.14
	(0.37)	(0.22)	(0.20)	(0.19)	(0.21)
Latitude	0.71	0.04	-0.63	0.19	1.04
	(2.34)	(0.90)	(0.92)	(1.25)	(0.90)
Log Child Dependency Ratio (avg)	-1.62	0.60	3.29	-0.62	-0.74
	(0.99)	(0.79)	(0.99)	(0.74)	(0.51)
N	25	52	37	110	165
R-sq	0.644	0.657	0.761	0.481	0.307

Notes: The dependent variable is annualized per capita GDP growth, in percent, over the 30-year period with data from Penn World Table (PWT), version 10.01 (Feenstra et al., 2015). Column (1) covers 1960–1990 and includes the 25 early-transitioners (post-transitioners as of 1960), and column (2) covers 1990–2020 and includes all 52 post-transitioners as of 1990 (both as listed in Hayashi (2025), Appendix 2). Column (3) covers 1990–2020 and restricts to OECD countries, Asian Tigers (Hong Kong, Singapore, South Korea, and Taiwan), and China ($N = 37$ as listed in Hayashi (2025), Appendix 6). Columns (4) and (5) use the full sample of countries with available data, for the periods 1960–1990 and 1990–2020, respectively. Log initial GDP per capita is log GDP per capita at the start of the period (1960 for columns (1) and (4); 1990 for columns (2), (3), and (5)). Log old-to-young ratio (Panel A) is the mean of log(population aged 50–79/population aged 20–49) over the 30-year period; change in log old-to-young ratio (Panel B) is the 30-year log difference in that ratio between the end and start of the period. The population data come from 2024 UN Population Prospects (United Nations, Department of Economic and Social Affairs, Population Division, 2024) extended to 2020 using the method in Hayashi (2025), Appendix 1. Rule of law is the Kaufmann and Kraay (2010) (World Governance Indicators, 2023 update) measure of institutional quality, averaged over 1996–2019 for the 1990–2020 period; its 1996 value (the first year of the series) is used for the 1960–1990 period. Log openness is the log of the ratio of imports plus exports to GDP from PWT, averaged over the 30-year period. Latitude is absolute latitude divided by 90 from Hall and Jones (1999). Log child dependency ratio is the log of the ratio of the population aged 0–19 to those aged 20–79, averaged over the 30-year period from UN Population Prospects. Following Hayashi (2025), conventional OLS standard errors are reported in parentheses.

Table C-4: Effect of Birth Rates and Older-to-younger Ratio on Growth in GDP per Worker

	(1) Base	(2) Educ	(3) Urban	(4) Region	(5) All	(6) FE
			Panel A: Birth Rate			
Birth Rate (t-20)	-0.22	-0.18	-0.21	-0.15	-0.10	-0.04
	(0.05)	(0.05)	(0.05)	(0.06)	(0.06)	(0.07)
N	229	229	229	229	229	240
R-sq	0.406	0.416	0.408	0.443	0.460	0.838
			Panel B: Birth Rate + Log Older-to-younger Ratio			
Birth Rate (t-20)	-0.22	-0.19	-0.21	-0.14	-0.09	0.00
	(0.06)	(0.06)	(0.06)	(0.06)	(0.07)	(0.11)
Log Older-to-younger Ratio	-0.01	-0.08	0.00	0.12	0.11	0.19
	(0.24)	(0.24)	(0.23)	(0.26)	(0.25)	(0.42)
N	229	229	229	229	229	240
R-sq	0.406	0.416	0.408	0.444	0.461	0.838
			Panel C: Post-Transitioners Sample			
Birth Rate (t-20)	-0.21	-0.17	-0.20	-0.27	-0.18	-0.49
	(0.10)	(0.10)	(0.10)	(0.11)	(0.11)	(0.19)
Log Older-to-younger Ratio	-0.12	0.05	-0.11	-0.38	-0.19	-1.74
	(0.52)	(0.51)	(0.52)	(0.52)	(0.49)	(0.91)
N	75	75	75	75	75	66
R-sq	0.577	0.588	0.577	0.623	0.664	0.885

Notes: This table reports estimates of β^h in equation (10), estimated in our cross-country dataset. The dependent variable is $\Delta_h y_{c,t+h}$, the change in log GDP per worker over horizon h. All panels set $h = 25$ and restrict the sample to two non-overlapping cross sections, where $t \in \{1970, 1995\}$. All panels weight countries by working-age population in year t. Panels A and B use the full country sample; Panel C restricts to post-transitioner countries as of 1990 (as listed by Hayashi (2025), Appendix 2). The vector $\boldsymbol{x}_{c,t}$ includes initial log GDP per worker and labor supply in t, measured by log population aged 20–45 and log population aged 45–70. Panel A includes the birth rate in $t-20$ as the sole key regressor. Panel B includes both the birth rate in $t-20$ and the average log old-to-young ratio over the h-year horizon, where the log old-to-young ratio is defined as log(population aged 50–79/population aged 20–49). Panel C replicates Panel B's specification restricted to post-transitioner countries. Column 2 additionally controls for average years of schooling among the population aged 45–55 in year $t-20$. Column 3 controls for the initial urbanization share in t. Column 4 controls for continent fixed effects, while Column 5 adds all controls simultaneously. Column 6 includes country fixed effects. Standard errors, clustered by country, are reported in parentheses.