INTERNATIONAL MONETARY FUND

Aging Economies and AI Adoption: Firm-Level Evidence from the World Bank Enterprise Surveys

Ha Minh Nguyen

WP/26/176

2026
AUG

WORKING PAPER

 WP/26/176

IMF Working Paper
Institute for Capacity Development

Aging Economies and AI Adoption: Firm-Level Evidence from the World Bank Enterprise Surveys

Prepared by Ha Minh Nguyen*

Authorized for distribution by Mercedes Garcia-Escribano
August 2026

ABSTRACT: How do demographic trends shape the adoption of AI and automation technologies? This paper provides the first large-scale cross-country firm-level test of the demographic–automation hypothesis using World Bank Enterprise Surveys data covering 89,380 firms across 144 countries from 2022 to 2025. I classify adopters by applying a large language model to firms' open-ended process innovation descriptions, identifying 1,656 AI and automation adopters (1.9 percent of the sample). A ten-percentage-point increase in the old-age dependency ratio raises process adoption probability by approximately 0.6 percentage points, after accounting for countries' income levels, digital infrastructure, firm size and sector, and broad regional and time differences. The result is robust across specifications and supported by an instrumental variable strategy based on predetermined demographic cohort structure. Heterogeneity analysis shows the effect concentrates in manufacturing, large firms, and developing economies for the broad adoption measure; restricting to firms with explicit references to AI reverses the sector pattern, with services firms significantly more likely to adopt than manufacturing firms, pointing to distinct sectoral profiles for software-based AI and hardware-based automation. Aging also predicts firms' development of AI-enabled products across both manufacturing and services. The results indicate that demographic aging shapes AI and automation adoption through both process and product innovation channels: firms substitute technology for increasingly scarce and costly labor in production, and separately develop AI-enabled products for labor-constrained customers.

RECOMMENDED CITATION: Ha Minh Nguyen (2026) "Aging Economies and AI Adoption: Firm-Level Evidence from the World Bank Enterprise Surveys", IMF Working Paper 26/176

JEL Classification Numbers:	J11, O33, L60, F63, J24
Keywords:	aging; automation; artificial intelligence; firm-level; technology adoption; labor-saving technology; World Bank Enterprise Surveys; demographics; labor substitution
Author's E-Mail Address:	hnguyen7@imf.org

Aging Economies and AI Adoption: Firm-Level Evidence from the World Bank Enterprise Surveys

Ha Nguyen*

August 11, 2026

Abstract

How do demographic trends shape the adoption of AI and automation technologies? This paper provides the first large-scale cross-country firm-level test of the demographic–automation hypothesis using World Bank Enterprise Surveys data covering 89,380 firms across 144 countries from 2022 to 2025. I classify adopters by applying a large language model to firms' open-ended process innovation descriptions, identifying 1,656 AI and automation adopters (1.9 percent of the sample). A ten-percentage-point increase in the old-age dependency ratio raises process adoption probability by approximately 0.6 percentage points, after accounting for countries' income levels, digital infrastructure, firm size and sector, and broad regional and time differences. The result is robust across specifications and supported by an instrumental variable strategy based on predetermined demographic cohort structure. Heterogeneity analysis shows the effect concentrates in manufacturing, large firms, and developing economies for the broad adoption measure; restricting to firms with explicit references to AI reverses the sector pattern, with services firms significantly more likely to adopt than manufacturing firms, pointing to distinct sectoral profiles for software-based AI and hardware-based automation. Aging also predicts firms' development of AI-enabled products across both manufacturing and services. The results indicate that demographic aging shapes AI and automation adoption through both process and product innovation channels: firms substitute technology for increasingly scarce and costly labor in production, and separately develop AI-enabled products for labor-constrained customers.

Keywords: aging, automation, artificial intelligence, firm-level, technology adoption, labor-saving technology, World Bank Enterprise Surveys, demographics, labor substitution

JEL codes: J11, O33, L60, F63, J24

*I thank Karim Barhoumi, Karla Chaman, Habtamu Edjigu, Mercedes Garcia-Escribano, Rishi Goyal, Gon Huertas, Asif Islam, Anna Ivanova, Yosuke Kido, Rose Kouwenhoven, Nils Haakon Lehr, Lamin Leigh, Vivian Malta, Albert Nyikuli, Carlo Pizzinelli, Mehdi Raissi, Emma Rockall, Baobing Shang, Nikola Spatafora, and Alice Zhang, and participants at the IMF SPR seminar for helpful comments and suggestions. All views expressed in this paper are my own and do not necessarily reflect the position of the IMF, its Executive Board, or IMF management.

1 Introduction

The global workforce is aging at a rapid pace. Old-age dependency ratios have risen sharply across Europe, East Asia, and increasingly middle-income economies over the past three decades, and projections suggest this trend will intensify through 2050. Germany, Japan, South Korea, and Italy are prominent examples, but the phenomenon is spreading: China's working-age population peaked around 2015, and countries from Thailand to Brazil are aging faster, at earlier stages of economic development, than today's advanced economies did when they reached comparable old-age dependency levels. Simultaneously, a wave of artificial intelligence and automation technologies has become increasingly accessible to firms of all sizes, from robotic process automation and industrial robots to machine learning platforms and generative AI tools that became commercially available following the launch of large language models in late 2022. A natural and policy-relevant question arises: are these two trends connected? Do firms facing shrinking domestic labor supplies adopt labor-saving technologies more intensively?

The theoretical case is straightforward. When the price or scarcity of a factor of production rises, rational firms have an incentive to substitute toward technologies that reduce their dependence on that factor. Aging and low fertility, by reducing the working-age population relative to dependents, raise the effective cost and scarcity of labor over time. The canonical prediction is that this should accelerate adoption of automation, robotics, and AI among firms operating in demographically older economies relative to their counterparts in younger economies. This is the labor-substitution hypothesis, and it is the organizing framework of this paper.

Despite its intuitive appeal, the systematic firm-level evidence on this hypothesis remains limited. The existing literature has primarily operated at the country or industry level, documenting correlations between population aging and aggregate measures of technology investment or robot adoption (e.g., Acemoglu and Restrepo, 2022). What is missing is direct evidence from firms themselves. Aggregate analysis cannot resolve whether the demographic-adoption link concentrates in large firms, which can spread the fixed costs of automation across a larger production base, or whether it also reaches smaller firms facing tighter financing constraints. And it collapses two conceptually distinct channels into a single aggregate: firms adopting AI and robotics in their own production processes, the direct labor-substitution response to aging, and firms developing or offering AI-powered products and services for labor-constrained customers, a demand-side complement. These channels have different implications for labor markets, industrial structure, and the distribution of productivity gains. Only firm-level data with rich descriptions of actual innovation activity can disentangle them.

This paper fills that gap. I use the World Bank Enterprise Surveys (WBES), a harmonized firm-level survey covering 89,380 firms across 144 countries surveyed between 2022 and 2025. The WBES asks firms to describe in open-ended text their recent product and process innovations, providing a unique window into the nature of what firms are actually doing. I apply a large language model (Claude) to classify firms' process innovation descriptions, identifying those that describe adopting artificial intelligence, automation, or robotics as adopters of labor-saving technology. This yields 1,656 adopting firms (1.9 percent of the full sample), a measure grounded in firms' own descriptions of their process innovations rather than in pre-defined survey categories. Process innovation is the natural baseline for testing the labor-substitution hypothesis: aging raises the effective cost of labor as a factor of production, and the most direct firm-level response is to adopt technologies that substitute for labor in production processes. Product innovation, by contrast, concerns what firms produce rather than how they produce it, and may reflect demand-side incentives distinct from labor scarcity; I treat it as a complementary channel in Section 7.

I merge this firm-level adoption indicator with country-level demographic and economic variables from the World Bank's World Development Indicators (WDI), matching each firm to the conditions prevailing in its country at the time of the survey. My primary demographic variable is the old-age dependency ratio; I also examine the total fertility rate and the share of the population aged 65 and above as alternatives. My empirical strategy is a probit regression with standard errors clustered at the country level, controlling for log GDP per capita, internet penetration, firm size, sector, region fixed effects, and survey-year fixed effects. The year fixed effects are particularly important in this period: global awareness and accessibility of AI technologies rose sharply after 2022, and without year controls the demographic coefficient would partly reflect the spurious correlation between demographic maturity and the timing of country surveys.

My main finding is that the old-age dependency ratio is positively and significantly associated with firm-level process AI adoption, that is, firms adopting AI, automation, or robotics to improve their production processes. A ten-percentage-point increase in the old-age dependency ratio raises the process adoption probability by approximately 0.6 percentage points, conditional on the full set of controls. Given a baseline process adoption rate of 1.9 percent, this represents a roughly 30 percent increase in the adoption probability relative to the mean. The coefficient is stable across four successive specifications, declining by only about 22 percent as the full set of controls is added, suggesting it is not a proxy for income or digital readiness. The result is robust to seven alternative estimations, including a keyword classifier, lagged demographics, restriction to innovator firms, exclusion of high-income

OECD countries, linear probability models, and country-equal-weighted least squares. Separately, replicating the analysis with a strict AI-only measure (firms with explicit references to artificial intelligence or machine learning, 221 firms) yields consistent results, confirming that the demographic signal is present in explicit AI process adoption and not only in the broader automation and robotics component.

To address the endogeneity of the old-age dependency ratio, I construct an instrument following Acemoglu and Restrepo (2022): predicted aging is computed by projecting each country's year-2000 demographic cohort structure forward using global average survival rates, so that the instrument varies only through initial cohort composition and not through differential mortality improvements correlated with development. The instrument is strong and the IV estimate exceeds the OLS estimate.

Three heterogeneity findings sharpen the substantive interpretation. First, the demographic effect on process adoption is significant in manufacturing but not services, consistent with the greater scope for labor substitution in physical production; restricting to a strict AI-only measure reverses this pattern, with services firms significantly more likely to adopt explicit AI tools than manufacturing firms, pointing to distinct sectoral profiles for software-based AI and hardware-based automation (Section 5.4). Second, large firms drive the process adoption result, reflecting the substantial fixed costs of automation investment that favor firms with sufficient scale. Third, the effect is present across both low- and middle-income developing economies, with imprecise estimates for advanced economies where within-group variation in old-age dependency is limited. Section 7 extends the analysis to product AI adoption (firms developing or offering AI-powered products), finding a similarly positive and significant demographic gradient, but with a strikingly different sector pattern: the effect spans both manufacturing and services, consistent with a demand-side channel complementary to the process-level labor-substitution motive. Finally, to illuminate the mechanism, I test whether demographic aging predicts firm-level reporting of labor constraints and whether the adoption response concentrates among labor-intensive firms. Firms in older economies are significantly more likely to cite an inadequately educated workforce as their primary business obstacle and to report higher severity of workforce skill gaps. Interacting the old-age dependency ratio with firm-level labor intensity – total labor cost as a share of sales – reveals that the aging effect is larger among firms for which labor is the dominant production input, consistent with the substitution mechanism operating most forcefully where labor costs are greatest.

The paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the data and classification methodology. Section 4 presents the empirical specification. Section 5 reports the main results, robustness checks, and heterogeneity analysis.

Section 6 presents instrumental variable estimates using predicted demographic aging constructed from UN World Population Prospects cohort data. Section 7 extends the analysis to product AI adoption. Section 8 discusses the findings and concludes.

2 Related Literature

This paper sits at the intersection of the economics of aging, the adoption of automation and artificial intelligence, and firm-level innovation. A central idea linking these literatures is that changes in labor supply conditions shape the direction of technological change. Population aging carries well-documented macroeconomic costs, lowering productivity and GDP growth by shrinking the middle-aged share of the workforce, with estimated effects on the order of a few tenths of a percentage point of annual growth (Feyrer, 2007; Aiyar et al., 2016; Maestas et al., 2023; International Monetary Fund, 2025). Automation is one channel through which firms and economies might offset these headwinds, motivating the question this paper addresses: does aging itself accelerate technology adoption?

The idea that labor scarcity induces labor-saving innovation dates to Habakkuk (1962) and Margo (2000), who link higher nineteenth-century American wages to faster mechanization relative to Britain. In modern settings, Acemoglu and Restrepo (2022) show that aging economies adopt industrial robots more intensively, the key aggregate-level precedent for this paper. A related literature documents that aging affects firm dynamics and innovation more broadly: it slows firm entry and worker reallocation (Engbom, 2024), reduces R&D and patenting through weaker demand for new inventions (Lehr, 2023), and shapes individual technology adoption near retirement (Friedberg, 2003). These findings raise the possibility that aging dampens innovation overall, which would complicate the interpretation of a positive correlation between aging and the automation share of innovation; I address this directly in Section 5.3 by testing whether aging predicts the extensive margin of innovation.

A large body of work documents how automation reshapes labor demand and firm outcomes. At the aggregate level, industrial robots raise productivity and wages while displacing routine-task and non-college workers (Graetz and Michaels, 2018; Autor, 2015; Autor et al., 2015; Charles et al., 2019; Lordan and Neumark, 2018). At the firm level, robot and AI adoption raise productivity and shift labor demand toward higher-skilled workers, concentrated among larger, more skill-intensive firms in more developed economies (Koch et al., 2021; Bonfiglioli et al., 2024; Brynjolfsson et al., 2025; Acemoglu and Johnson, 2023; Calvino et al., 2022; Calvino and Fontanelli, 2023). These studies establish the mechanisms through which technology affects labor demand but largely take the direction of adoption as given.

My classification approach also draws on a growing methodological literature applying large language models to measurement problems in economics, which emphasizes validating LLM output against human-coded benchmarks and documents strong agreement between LLM and human classification (Ludwig et al., 2026; Asirvatham et al., 2026; Ash et al., 2026). I follow this guidance by validating my classifications against both a keyword-based benchmark and my own hand coding of a sample of responses (Section 3.2).

Despite this progress, direct firm-level evidence on how demographic forces shape technology adoption remains limited. Existing work relies on country- or industry-level data that cannot distinguish adoption decisions across firm size, sector, or income group, and conflates distinct channels such as firms automating their own production versus developing AI-enabled products. This gap is compounded by the fact that technology diffusion across countries is uneven not only in timing but also in intensity of use, with poorer countries historically adopting new technologies more slowly and using them less intensively (Comin and Mestieri, 2018), relevant background for the income-group heterogeneity documented in Section 5.5. This paper contributes to these literatures by shifting the unit of analysis to individual firms, introducing the demographic driver to the firm-level AI adoption literature, measuring adoption from firms' own open-ended descriptions rather than pre-defined categories, and separately identifying process and product adoption channels with distinct sector footprints. Consistent with the literature on firm size and innovation (Cohen, 1995; Akcigit and Kerr, 2018), I find the demographic effect concentrated among larger firms and in more routine-task-intensive sectors, connecting demographic pressures to firm heterogeneity in technology diffusion.

3 Data

3.1 World Bank Enterprise Surveys

The World Bank Enterprise Surveys (WBES) is a firm-level survey designed to assess the business environment in formal private sector establishments. It is conducted by the World Bank in collaboration with national statistical offices and development banks, covering manufacturing and services firms with five or more employees. The survey is stratified by firm size, sector, and geographic region within each country, with sampling weights (denoted w_{ic}) designed to make each country's sample representative of its formal private sector.

I use the most recent round of surveys, conducted between 2022 and 2025. This wave is particularly relevant for studying AI and automation adoption because it encompasses the commercial breakthrough of generative AI following the launch of ChatGPT in late 2022. The adoption of large language models and related automation tools accelerated rapidly during this period, making the 2022-2025 WBES window an unusually informative snapshot of firms' responses to newly accessible AI technologies.

After restricting to observations with non-missing innovation text fields and merging with country-level demographic data (described in Section 3.3), my working sample comprises 89,380 firms across 144 countries. The survey covers all major world regions: Europe and Central Asia (ECA, 26,021 firms), Sub-Saharan Africa (AFR, 16,389 firms), East Asia and Pacific (EAP, 14,707 firms), South Asia (SAR, 11,607 firms), Middle East and North Africa (MNA, 10,072 firms), Latin America and the Caribbean (LAC, 6,872 firms), and high-income OECD economies (HIC, 3,712 firms).[1] Among these, 22.4 percent of firms report introducing a new or significantly improved product or service, 15.5 percent report a new or significantly improved process, and 28.4 percent report at least one of the two.

3.2 Measuring Technology Adoption from Open-Ended Text

The WBES includes open-ended text fields that are central to my identification strategy. For firms that report process innovation, question h6x asks them to describe the new or improved process, question h7x asks what specifically was improved in that process, and question h62x provides an additional description of the process change. For firms that report product innovation, questions h3x and h4x capture descriptions of the new or improved product or service and what specifically was improved, with h32x providing an additional product description. These fields capture what firms are actually doing in their own words, providing a richer and more credible basis for identifying technology adoption than pre-defined survey categories. The process innovation fields (h6x, h7x, h62x) are used in the main analysis; the product innovation fields (h3x, h4x, h32x) are used in Section 7.

Applying a large language model to these text responses constitutes a methodological contribution, with three properties that distinguish it from prior approaches. First, no additional data collection is required: the open-ended fields exist in every WBES wave and have previously been used primarily for broad innovation counts rather than for identifying specific technologies. Classifying them for AI and automation content adds a new measurement layer at near-zero marginal cost, a property that generalizes immediately to future waves and to other large-scale firm surveys that include similar open-ended innovation questions. Second, the LLM reads *semantic context* rather than matching fixed strings, allowing it to identify technology adoption expressed through synonyms, adjective forms, paraphrases, and non-standard terminology that a keyword classifier cannot reach. Third, setting the generation temperature to zero makes the classification *fully deterministic*: identical input text

[1] Regional classifications follow those used by the World Bank Enterprise Surveys.

always produces the same output, so the measure is replicable by any researcher with access to the same model version.

I classify process AI adoption using Claude (Anthropic's Claude-Opus-4-6 large language model) applied to the concatenated process innovation text (h6x, h7x, and h62x) for each firm. Among the 14,001 firms with non-empty responses in at least one of these process innovation fields, each response is submitted to Claude with a structured classification prompt (reproduced in Appendix A.1) instructing the model to assign the firm to one of three tiers.

Tier 1 (explicit AI/ML) captures unambiguous artificial intelligence and machine learning terminology: artificial intelligence, machine learning, deep learning, neural networks, large language models, generative AI, named AI products and platforms (ChatGPT, Copilot, Gemini, Databricks, Salesforce Einstein), computer vision, natural language processing, sentiment analysis, image and facial recognition, voice and speech recognition, anomaly detection, robotic process automation (RPA), predictive models and analytics, and AI-powered or AI-enabled systems. Representative responses include: *"I designed and developed an AI control system through which I control the mounting of the components on the board"* (Romania, manufacturing, 2023) and *"AI hardware automation in the technological process of assembling toys, previously this process was performed manually"* (Moldova, manufacturing, 2024).

Tier 2 (automation and robotics) captures direct labor-substituting technologies: automation and its morphological variants (automated, automating, automatic), robots and robotics (including robotization), cobots (collaborative robots), and descriptions of machines explicitly replacing manual labor processes. Other dimensions of digital transformation that appear in the literature, such as ERP systems, cloud computing, IoT sensors, and 3D printing, are classified as Tier 0 because they do not directly substitute for labor in the manner required by the demographic hypothesis. Representative Tier 2 responses include: *"Introduced new automation equipment on the production line, before done by hand"* (Sierra Leone, manufacturing, 2023) and *"Automatic production line replacing fully manual process"* (Vietnam, manufacturing, 2024).

Tier 0 (non-adopter) covers all other responses. For contrast, firms reporting genuine innovation that does not involve labor-substituting technology are not classified as adopters. Examples include: *"Updated the reward programs for loyal customers"* (United States, services, 2024) and *"Dehydration of the quinoa grain to speed up cooking"* (Bolivia, manufacturing, 2025), neither of which describes automating a production task. Table 1 presents representative examples from each classification category drawn from the actual WBES responses, illustrating the distinction between AI process adoption, AI product adoption, and conventional innovation that the classifier must draw.

Table 1: Representative WBES Responses by Classification Category

Firm response (verbatim)	Country	Year	Sector
Panel A: *AI process adoption (ai_process_broad = 1)*			
"I designed and developed an AI control system through which I control the mounting of components on the board. Until now, there was an operator who did this. Now, there is software that checks if the assembly is correct." [Tier 1]	Romania	2023	Services
"AI hardware automation in the technological process of assembling toys, previously this process was performed manually." [Tier 1]	Moldova	2024	Manufacturing
"The sheets, which were previously welded by hand, were henceforth manufactured by robots." [Tier 2]	Germany	2022	Manufacturing
Panel B: *AI product adoption (ai_product_broad = 1)*			
"AI software for the optimization of plastics production. The use of AI in the plastics sector is new." [Tier 1]	Germany	2022	Manufacturing
"A new AI solutions system which can calculate the demand of a shop based on previous data, with fewer workers than before." [Tier 1]	Hungary	2024	Services
"Self-service cash register – completely new product." [Tier 2]	Georgia	2023	Services
Panel C: *Non-adopting innovation (Tier 0 – excluded from both measures)*			
"A new software system that enables more efficient ordering, helping to avoid food waste."	Denmark	2025	Services
"I revised the supply of raw material on the machine and the redistribution of workers in order to cut down downtime."	Italy	2024	Manufacturing
"I implemented a virtual store through which sales are direct."	Peru	2023	Services

Notes: Responses are drawn from WBES process innovation text fields (h6x, h7x, h62x) for Panels A and C, and product innovation text fields (h3x, h4x, h32x) for Panel B. Tier 1 denotes explicit AI or machine learning; Tier 2 denotes automation or robotics without explicit AI terminology. Panel C responses describe genuine innovations – digital ordering systems, production reorganization, e-commerce – that do not involve direct labor substitution and are therefore classified Tier 0.

The LLM temperature is set to zero, as noted above. My main process adoption outcome variable, *ai_process_broad*, equals one if a firm's process innovation text is classified as Tier 1 or Tier 2, and zero otherwise. This yields 1,656 process adopters, or 1.9 percent of the full sample. As an extension, *ai_process_strict* equals one for Tier 1 classifications only (explicit AI/ML), capturing 221 firms (0.2 percent of the sample).

To validate the LLM measure and quantify what it adds relative to existing approaches, I compare it against a traditional keyword-based classifier using 33 fixed regular expressions (exact patterns listed in Appendix A.2).[2] The keyword classifier identifies 1,104 adopters; the LLM identifies 1,656, approximately 550 more. The two classifiers agree on 96.0 percent of text-response firms, confirming that the additional adopters represent a targeted recovery of genuine cases rather than a broadening of the concept. The disagreements are systematic: the LLM recovers firms whose responses describe the same underlying activity, adopting robots, automating production, implementing AI-assisted processes, but express it in forms that fixed strings cannot reach. These include adjective and morphological variants ("automatic", "automatically", "robotization"), paraphrases ("the process is now done by machine rather than by hand"), and minor misspellings ("Artificial Inteligence", "mchine learning"). Representative examples of firms recovered by the LLM but missed by keywords include: *"All production is now fully automatic, no manual steps remain"* (Indonesia, manufacturing, 2024) and *"Robotization of the assembly line eliminated the need for manual labor at three stages"* (Poland, manufacturing, 2023).

From a measurement standpoint, keyword false negatives, real adopters coded as zeros, create attenuation bias in OLS: if some adopters are misclassified as non-adopters, the estimated relationship between aging and adoption is biased toward zero. The LLM's higher recall reduces this source of measurement error. Consistent with this, the LLM-based coefficient (0.0151) modestly exceeds the keyword-based coefficient (0.014; column 2 of Table 6), as expected under reduced attenuation. The main findings are robust to using either measure.

As a further check on classification accuracy, I validate the LLM measure against my own hand coding on a stratified random sample of 50 firms (evenly split, subject to availability, across process and product innovation text and across the three tiers, so that both adopter tiers are well represented despite their low base rate). I classified each response using the same tier definitions given to the LLM (Appendix A.1), without consulting the LLM's classification while coding. Raw agreement is 90.0 percent and Cohen's kappa is 0.85, conventionally interpreted as "almost perfect" agreement; agreement is similar for process text (92.3 percent, $\kappa = 0.885$, $n = 26$) and product text (87.5 percent, $\kappa = 0.812$, $n = 24$). All five disagreements involve borderline cases rather than systematic errors.

Weighing these results against the properties described above, the more fundamental limitation is one shared by any LLM-based measurement strategy: the classification is a model-based inference rather than a direct observation, and absent a validation sample, researchers cannot assess the magnitude or pattern of LLM errors, while seemingly innocuous choices such as model version or prompt wording can materially change downstream estimates (Ludwig et al., 2026). LLMs are also complex and lacking in transparency relative

[2] Keyword-based classification of WBES innovation text is also used by Udomsaph (2025), who applies a similar approach to approximately 72,000 firms in the WBES.

to simpler text-analysis methods (Ash et al., 2026), and can in principle infer a label from correlated textual cues rather than the construct actually being measured, a failure mode sometimes called "shortcut inference" (Asirvatham et al., 2026). The validation exercise above is my empirical response to this concern, but it mitigates rather than eliminates it.

3.3 Country-Level Demographic and Economic Variables

I merge the firm-level data with country-level variables from the World Bank's World Development Indicators (WDI), matching each firm to the conditions prevailing in its country in the calendar year of the WBES. Of the 144 countries, five lack WDI data on internet penetration or GDP per capita (Central African Republic, Ethiopia, Kosovo, Somalia, and South Sudan), so the preferred specification uses up to 139 countries depending on the set of controls included. Specifications without internet penetration use up to 144 countries.

My primary demographic variable is the **old-age dependency ratio**, the number of people aged 65 and above per 100 people of working age (15 to 64), expressed as a percentage. This variable directly captures the stock of demographic aging in the workforce. A higher ratio implies more retirees relative to active workers, tightening the effective labor market and raising wage pressure. Across the 144 countries in my sample, the old-age dependency ratio ranges from below 3 in the youngest Sub-Saharan African economies to above 38 in Germany and Japan.

I also examine two alternative demographic measures. The **total fertility rate** (TFR) captures the flow of future labor supply: lower fertility implies fewer young workers entering the labor force in the future, creating a forward-looking incentive to automate. The **share of population aged 65 and above** is an alternative stock measure of aging highly correlated with the dependency ratio ($r = 0.997$ across countries in my sample).

Economic control variables include: (i) log GDP per capita, measured in PPP-adjusted constant 2017 international dollars, controlling for the level of economic development and the affordability of technology investment; and (ii) internet users as a share of the population, capturing digital readiness, which conditions the feasibility of AI and automation adoption regardless of demographic incentives. Population growth is available but not included in the preferred specification as it is collinear with the demographic measures included.

Table 2 reports the country-level correlation matrix for the key variables. As the table shows, the old-age dependency ratio is moderately correlated with log GDP per capita ($r = 0.72$), reflecting the tendency for richer countries to have older populations. Importantly, this correlation is far from perfect, so the two variables identify different sources of variation in my regressions. The correlation between the old-age dependency ratio and the share of population aged 65 and above is 0.997, confirming they should not be included simultaneously (as Table 5 ensures). The total fertility rate (TFR) is highly correlated with both income ($r = -0.84$) and internet penetration ($r = -0.84$), which helps explain its lack of independent identification in the regressions.

Table 2: Country-Level Correlation Matrix

	(1) Old-age dep.	(2) TFR	(3) ln(GDPpc)	(4) Internet (%)
(1) Old-age dep. ratio	1.000			
(2) Total fertility rate	−0.690	1.000		
(3) ln(GDP per capita)	0.720	−0.840	1.000	
(4) Internet users (%)	0.620	−0.840	0.880	1.000

Notes: Pairwise Pearson correlations across countries (138–144 depending on variable coverage). Each country is one observation, equal-weighted.

3.4 Descriptive Statistics

Table 3 reports descriptive statistics for the full sample and separately for process AI/automation adopters (Tiers 1 and 2) and non-adopters. The raw differences between the two groups are large and consistent with the hypothesis. Process adopters are located in countries with an average old-age dependency ratio of 23.3, compared with 16.1 for non-adopters, a gap of 7.2 percentage points or roughly 0.7 standard deviations. Adopters also face lower fertility rates (mean TFR 1.75 versus 2.28) and a higher share of elderly population (15.1 versus 10.5 percent). These differences persist along economic dimensions: adopting firms operate in richer countries (average GDP per capita $41,400 versus $27,600) with higher internet penetration (85.4 versus 72.2 percent).

At the firm level, process AI/automation adopters are concentrated in larger firms: 38.7 percent of adopters are in the large-firm category (100 or more employees), compared with 19.9 percent of non-adopters. Adopters are also disproportionately in manufacturing: 64.8 percent of adopters are manufacturing firms, compared with 40.7 percent of non-adopters, suggesting that process AI adoption concentrates in manufacturing rather than services.

Table 3, Panel B reports process AI/automation adoption rates and demographic averages by world region. The pattern is strongly consistent with the hypothesis. Europe and Central Asia, the most demographically mature developing region in the sample (average old-age dependency ratio 28.3, average TFR 1.64), has the highest adoption rate at 3.29 percent. High-income OECD economies have similarly high old-age dependency (28.4) and an adoption rate of 1.83 percent. At the other extreme, South Asia, with an average old-age dependency ratio of 10.2 and TFR of 2.00, has an adoption rate of just 0.16 percent, and Sub-Saharan Africa (old-age dependency ratio 5.9, TFR 4.15) has an adoption rate of 0.92 percent.

Turning to innovation activity more broadly, 22.4 percent of the full sample report product innovation, 15.5 percent report process innovation, and 28.4 percent report at least one type; roughly 9.5 percent report both. Among the 14,001 firms that provided non-empty process innovation text (h6x, h7x, or h62x), process AI adoption is still rare: only 1,656 firms (11.8 percent of text-response firms, 1.9 percent of the full sample) qualify as process adopters. Similarly, among the 20,127 firms with non-empty product innovation text (h3x, h4x, or h32x), only 589 firms (2.9 percent of text-response firms, 0.7 percent of the full sample) qualify as product AI adopters. This underscores that while innovation broadly is relatively common, the specific adoption of AI and automation in production processes, or the development of AI-enabled products, remains uncommon even among firms that explicitly describe their innovations.

The country-level picture in Appendix Table 19 reinforces the narrative. The top adopters are almost exclusively aging European economies: Denmark leads with 15.5 percent adoption, followed by Switzerland (10.6 percent), Belgium (9.4 percent), Czechia (9.3 percent), Moldova (8.6 percent), and Slovenia (8.2 percent). China (7.9 percent) also features prominently. The bottom of the distribution is dominated by Sub-Saharan African and South Asian countries with young populations and low old-age dependency ratios. Several countries at the top of the table have unusually small samples: Somalia (5.0 percent, 241 firms) and Zimbabwe (4.5 percent, 359 firms) rank higher than their demographic profile would predict, so these and the other small-sample countries flagged in Table 19 should be interpreted with caution.

Figure 1 plots the country-level relationship directly. Each bubble represents one country, with size proportional to the log of the number of surveyed firms, and the dashed line is an unweighted OLS fit. The positive slope is clearly visible across the full range of the old-age dependency ratio, from the youngest Sub-Saharan African economies clustered near zero adoption to the oldest European economies concentrated in the upper right. The R^2 of 0.18 indicates that demographic aging alone explains a meaningful share of the cross-country variation in adoption rates, even before controlling for income, digital infrastructure, or any other country characteristics.

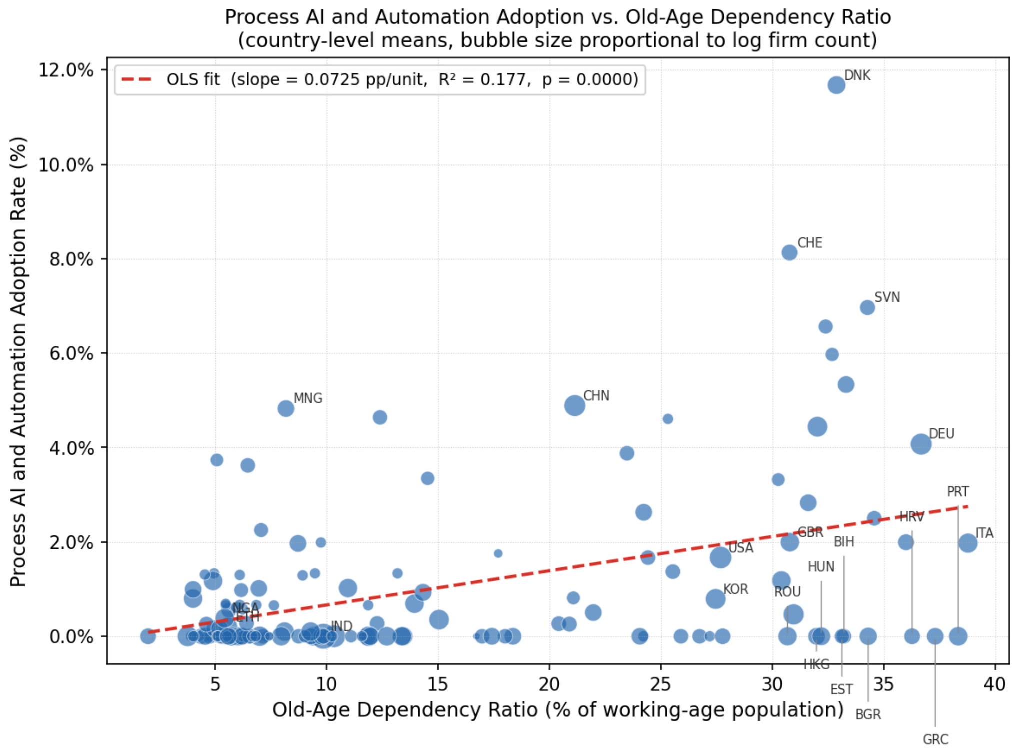

Figure 1: Process AI and Automation Adoption Rate and Old-Age Dependency Ratio

Notes: Each bubble is one country. Bubble size is proportional to the log number of surveyed firms. The dashed line is an unweighted OLS fit (slope = 0.073 percentage points per unit of OAD, $R^2 = 0.18$, $p < 0.01$). Adoption rates are country-level means computed from the WBES 2022–2025 sample (89,380 firms, 144 countries).

Table 3: Descriptive Statistics

	Full sample Mean (SD)	Adopters Mean (SD)	Non-adopters Mean (SD)
Panel A: Full sample			
AI/automation adoption (%)	1.9 (13.5)	100.0 (0.0)	0.0 (0.0)
Old-age dep. ratio	16.3 (11.0)	23.3 (10.8)	16.1 (11.0)
Total fertility rate	2.27 (1.14)	1.75 (0.87)	2.28 (1.14)
Share pop. 65+ (%)	10.6 (7.2)	15.1 (6.9)	10.5 (7.1)
GDP per capita, PPP ($000s)	27.9 (25.7)	41.4 (27.2)	27.6 (25.6)
Internet users (%)	72.4 (24.5)	85.4 (18.4)	72.2 (24.5)
Population growth (%)	1.22 (1.31)	0.76 (1.15)	1.23 (1.32)
Small firm (%)	46.4 (49.9)	22.3 (41.7)	46.8 (49.9)
Medium firm (%)	33.3 (47.1)	38.9 (48.8)	33.2 (47.1)
Large firm (%)	20.3 (40.2)	38.7 (48.7)	19.9 (39.9)
Manufacturing (%)	41.2 (49.2)	64.8 (47.8)	40.7 (49.1)
Services (%)	58.8 (49.2)	35.2 (47.8)	59.3 (49.1)
Observations	89,380	1,656	87,724
Countries	144	123	144

Panel B: Adoption rate by region

Region	Firms	Adopters	Rate (%)	Old-age dep.	TFR
Sub-Saharan Africa	16,389	151	0.92	5.9	4.10
East Asia & Pacific	14,707	324	2.20	15.4	1.71
Europe & Central Asia	26,021	857	3.29	28.3	1.61
Latin America & Carib.	6,872	171	2.49	13.3	1.86
Middle East & N. Africa	10,072	67	0.67	7.6	2.68
South Asia	11,607	18	0.16	10.2	2.00
High income: OECD	3,712	68	1.83	28.4	1.52
Total	89,380	1,656	1.85	16.2	2.27

Notes: Adopters are firms with *ai_process_broad*=1 (LLM classification), i.e., firms that mentioned AI, automation, or robotics in their open-ended process innovation descriptions (WBES text fields h6x, h7x, or h62x). Country-level WDI variables are means across all firm-year observations within each region. Small firms: fewer than 20 employees. Medium: 20–99. Large: 100 or more.

4 Empirical Specification

4.1 Baseline Model

I estimate the following probit model:

$$\Pr(ai\_process_{ic} = 1) = \Phi(\alpha + \beta\, old\_age\_dep_c + \gamma'\mathbf{X}_{ic} + \delta_r + \lambda_t + \varepsilon_{ic}) \tag{1}$$

where $\Phi(\cdot)$ is the standard normal CDF; $ai\_process_{ic}$ equals one if firm i in country c adopted AI, automation, or robotics (the broad measure) and zero otherwise; $old\_age\_dep_c$ is the old-age dependency ratio in country c at the time of the survey; $\mathbf{X}_{ic}$ is a vector of firm-level and country-level controls; δ_r are region fixed effects; and λ_t are survey-year fixed effects. Standard errors are clustered at the country level to account for within-country correlation arising from the fact that both the outcome variable and the key regressor are effectively at the country level.

The parameter β is the coefficient of interest. A positive estimate of β is consistent with the labor-substitution hypothesis: firms in countries with older workforces are more likely to adopt AI, automation, and robotics.[3] I report probit coefficients alongside average marginal effects (AMEs) evaluated at the sample means of all regressors:

$$\widehat{\mathrm{AME}} = \frac{1}{N}\sum_{i=1}^{N} \phi\Big(\hat{\alpha} + \hat{\beta}\, old\_age\_dep_c + \hat{\gamma}'\mathbf{X}_{ic} + \hat{\delta}_r + \hat{\lambda}_t\Big) \cdot \hat{\beta} \tag{2}$$

where $\phi(\cdot)$ is the standard normal PDF. The AME translates the probit coefficient into a percentage-point change in the probability of adoption for a one-unit change in the old-age dependency ratio, averaged across the observed distribution of covariates.

4.2 Control Variables

The vector $\mathbf{X}_{ic}$ includes both firm-level and country-level controls. Firm-level controls consist of dummy variables for firm size: *med* (medium, 20 to 99 workers) and *large* (100 or more workers), with small firms (5 to 20 workers) as the base category; and *services*, an indicator for the services sector with manufacturing as the base. These controls capture the well-documented positive relationship between firm scale and adoption capacity, and the differential propensity to automate across sectors.

Country-level controls include log GDP per capita in PPP-adjusted constant 2017 international dollars and internet users as a share of the population. Log GDP per capita controls for the possibility that richer countries have both older populations and more resources to invest in technology, which could otherwise produce a spurious positive correlation between

[3] This general labor-scarcity framing differs from Acemoglu and Restrepo (2022), whose theoretical mechanism operates specifically through the ratio of middle-aged to older workers and their comparative advantage in manual production tasks; my design does not distinguish this task-specific channel from broader labor-market tightening. A related channel I do not separately identify is capital availability: life-cycle savings tend to rise with age, so aging economies may also have cheaper access to capital independent of labor-cost pressures, which could contribute to the same empirical pattern.

aging and adoption. Internet penetration controls for digital readiness, which conditions the feasibility of AI and software-based automation adoption independently of demographic incentives. The two country-level controls are moderately correlated with the old-age dependency ratio ($r = 0.72$ and $r = 0.62$, respectively) but not collinear, so their inclusion tests whether the demographic effect survives conditioning on these confounders.

4.3 Fixed Effects

Region fixed effects (δ_r) are included for six World Bank regions with Sub-Saharan Africa as the omitted base: East Asia and Pacific, Europe and Central Asia, Latin America and the Caribbean, Middle East and North Africa, South Asia, and high-income OECD economies. These absorb time-invariant differences in technology adoption propensity related to institutional quality, regulatory environment, cultural attitudes toward technology, or other persistent regional characteristics.

Survey-year fixed effects control for four survey years: 2023, 2024, and 2025 (with 2022 as the base). These are particularly important in the 2022-2025 period, which saw a sharp increase in AI awareness and accessibility following the commercial launch of ChatGPT in November 2022 and subsequent generative AI tools. Without year fixed effects, the demographic coefficient would partly capture the spurious correlation between demographic maturity and the timing of country surveys within this window, since wealthier and older countries tend to be surveyed more recently. The year fixed effects ensure that identification comes from cross-country variation within survey years rather than from trends.

4.4 Identification

My identification of β relies on cross-country variation in the old-age dependency ratio within survey years and world regions. The identifying assumption is that, conditional on income, digital readiness, region, and year, the remaining variation in old-age dependency ratios is orthogonal to unobserved determinants of firm-level technology adoption. This assumption could be violated if, for instance, countries with older populations also have stronger innovation policy frameworks, more technology-friendly regulation, or higher quality institutions that independently promote AI adoption.

I address the endogeneity of the old-age dependency ratio directly in Section 6 using an instrumental variable strategy following Acemoglu and Restrepo (2022). The instrument is constructed from each country's year-2000 cohort structure projected forward with global average survival rates, generating variation in current aging that was predetermined by fertility and mortality conditions prevailing decades before AI technologies existed. The instrument is strong (first-stage partial $F \approx 413$), the IV estimate exceeds the OLS estimate, consistent with classical attenuation bias rather than upward confounding. The probit estimates in Section 5 are therefore best understood as economically interpretable conditional correlations, with the IV analysis in Section 6 providing the causal benchmark.

5 Results

5.1 Main Results

Table 4 presents the baseline probit estimates. The old-age dependency ratio is positively and significantly associated with firm-level process AI adoption in all four specifications.

In the most parsimonious specification (column 1), which includes only the old-age dependency ratio alongside region and year fixed effects, the probit coefficient is 0.019, significant at the 1 percent level. The corresponding average marginal effect is 0.0008, implying that a ten-percentage-point increase in the old-age dependency ratio raises the adoption probability by 0.8 percentage points. Adding log GDP per capita (column 2) reduces the coefficient to 0.015, still significant at the 5 percent level: income is associated with adoption on its own, but controlling for it absorbs only part of the demographic effect. Adding internet penetration (column 3) leaves the demographic coefficient essentially unchanged at 0.016, significant at the 5 percent level, while internet penetration enters with the expected positive sign but is statistically insignificant in this column. This pattern suggests that internet infrastructure captures some developmental variation in adoption propensity but is distinct from the demographic channel.

The preferred specification (column 4) adds firm size and sector controls. The demographic coefficient is 0.015, significant at the 5 percent level, with an average marginal effect of 0.0006. A ten-percentage-point increase in the old-age dependency ratio raises the probability of adoption by approximately 0.6 percentage points. At a baseline adoption rate of 1.9 percent, this represents a relative increase of about 30 percent, a substantively meaningful effect. The stability of the coefficient across columns 1 through 4 (declining by approximately 18 percent over the full set of controls) is reassuring: the demographic effect is not a proxy for income or digital readiness.

Among the control variables, firm size is strongly and robustly associated with adoption. Medium-sized firms are 1.20 percentage points more likely to adopt than small firms (AME), and large firms are 1.93 percentage points more likely. These effects are significant at the 1 percent level and are consistent with the substantial fixed costs of automation and AI adoption favoring firms with sufficient scale. Services firms are 1.60 percentage points less likely to adopt than manufacturing firms, a finding I examine more carefully in the heterogeneity analysis. Log GDP per capita and internet penetration lose statistical significance in the full specification, reflecting multicollinearity with each other ($r = 0.88$) once both are included alongside the demographic variable.

Table 4: Main Results: Old-Age Dependency and Technology Adoption

	(1)	(2)	(3)	(4)
Old-age dep. ratio	0.0185***	0.0148**	0.0171**	0.0151**
	(0.0064)	(0.0066)	(0.0068)	(0.0071)
	[0.0008]	[0.0006]	[0.0007]	[0.0006]
ln(GDP per capita, PPP)		0.0805	−0.0148	0.0195
		(0.0591)	(0.0984)	(0.1080)
		[0.0034]	[−0.0006]	[0.0008]
Internet users (%)			0.0076	0.0075
			(0.0049)	(0.0050)
			[0.0003]	[0.0003]
Medium firm				0.2950***
				(0.0266)
				[0.0120]
Large firm				0.4743***
				(0.0401)
				[0.0193]
Services sector				−0.3936***
				(0.0401)
				[−0.0160]
Region FEs	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes
Observations	89,380	89,224	88,376	88,376
Countries	144	143	139	139
Pseudo R^2	0.072	0.073	0.076	0.123

Notes: Probit estimates. Dependent variable: *ai_process_broad* (=1 if firm classified as AI/automation adopter by LLM; in process innovation text). Standard errors clustered at the country level in parentheses ( ). Average marginal effects in brackets [ ]. Base categories: Small firm, Manufacturing sector, Sub-Saharan Africa, survey year 2022. Columns (3) and (4) exclude 5 countries lacking internet penetration data in the WDI relative to column (2). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

To convey the economic magnitude, consider a simple projection. United Nations, Department of Economic and Social Affairs, Population Division (2024) project the global average old-age dependency ratio to rise by approximately 10 percentage points between 2025 and 2050 under the medium fertility scenario. Applying the preferred-specification average marginal effect of 0.06 percentage points per unit of OAD (Table 4, column 4), this demographic shift alone implies an increase in firm-level process AI adoption probability of roughly 0.6 percentage points, a 30 percent increase relative to the current baseline adoption rate of 1.9 percent, holding all other factors constant. The implied acceleration is largest in currently young economies where projected aging is steepest, suggesting that demographic forces will sustain a long expansion in AI adoption even absent further declines in technology costs.

It is useful to benchmark this magnitude against other predictors in the same specification (Table 4, column 4). The AME of a ten-percentage-point rise in OAD (0.6 percentage points) is modest relative to the gap between small and large firms (AME of 1.93 percentage points) or between small and medium firms (AME of 1.20 percentage points), consistent with firm size being the single strongest predictor of adoption in this sample. At the same time, the demographic effect is estimated considerably more precisely than the effect of income: ln(GDP per capita, PPP) is statistically insignificant in the preferred specification (AME of 0.08 percentage points), indicating that the demographic channel is identified with more confidence than, and operates separately from, cross-country income differences once the full set of controls is included.

5.2 Alternative Demographic Measures

Table 5 examines whether the result is robust to alternative demographic measures. Each column replaces old-age dependency with a single alternative measure in the preferred specification, entering the three measures separately given their high mutual correlations.

The total fertility rate enters with the expected negative sign (lower fertility is associated with more adoption, coefficient -0.066) but is statistically insignificant. This null result likely reflects the stock-versus-flow distinction: TFR captures expected future labor supply conditions, while the old-age dependency ratio captures the current stock of aging already present in the workforce, which is more directly relevant to today's firm hiring and investment decisions. The high correlation between TFR and income ($r = -0.84$ at the country level) also limits the precision with which TFR can be separately identified in a specification that already controls for log GDP per capita. The share of population aged 65 and above enters positively and is marginally significant (coefficient 0.022), consistent with the old-age dependency result. As a further check on the stock-versus-flow interpretation, I replace contemporaneous TFR with TFR lagged 20 years, approximating the fertility conditions that produced today's new labor-market entrants. The lagged measure remains statistically insignificant and smaller in magnitude (coefficient -0.019) than the contemporaneous measure, and old-age dependency remains significant when both are included jointly. The null for TFR is therefore not explained by comparing a forward-looking measure to a current one at mismatched horizons.

The coherence between old-age dependency and 65 and above share, combined with the null for TFR, supports the interpretation that the relevant demographic channel operates through the *current* age composition of the workforce rather than through expectations about future labor scarcity.

A related question is why TFR tracks internet exposure and other digital readiness measures more closely than old-age dependency does, and what OAD's residual variation captures that TFR does not. At the country level (139 countries), TFR's raw correlation with internet penetration is -0.84, similar in magnitude to its correlation with log GDP per capita (-0.84), while OAD's raw correlation with internet penetration is a weaker 0.62. Once log GDP per capita is partialled out, this gap widens sharply: the residual correlation between OAD and internet penetration is essentially zero (-0.04), while the residual correlation between TFR and internet penetration remains substantial (-0.39). This indicates that TFR's association with digital infrastructure operates largely through broader modernization processes correlated with income, education, and urbanization, of which internet penetration is one marker, whereas old-age dependency's variation is close to orthogonal to digital readiness once income is held fixed. This supports interpreting the OAD coefficient in Table 4 as capturing a demographic channel specifically, rather than picking up residual variation in a country's general state of technological development.

An important caveat is that none of these measures directly captures migration. The World Bank's population estimates underlying old-age dependency, fertility, and the population aged 65 and above already incorporate net migration flows, so a country's measured demographic structure reflects the combined effect of natural change and migration. However, I do not separately test whether migration constitutes an independent margin of labor-supply adjustment that could substitute for automation.

For the remainder of the analysis, I use old-age dependency as my primary measure.[4]

[4] As an alternative measure of labor market tightness, I replace the old-age dependency ratio with the annual growth rate of the working-age population (ages 15–64). Working-age population growth is insignificant when entered alone and takes the wrong sign relative to the labor-scarcity mechanism. When both variables are included jointly, the old-age dependency ratio strengthens and becomes significant at the 1 percent level, while working-age population growth remains inconsistent with the automation-pressure channel. These results suggest that firms respond to the accumulated stock of demographic aging rather than to the recent rate of change in labor supply.

Table 5: Alternative Demographic Measures

	(1) Old-age dep.	(2) TFR	(3) Pop. 65+
Old-age dep. ratio	0.0151**		
	(0.0071)		
	[0.0006]		
Total fertility rate		−0.0655	
		(0.1029)	
		[−0.0035]	
Share population 65+ (%)			0.0218*
			(0.0115)
			[0.0008]
ln(GDP per capita)	Yes	Yes	Yes
Internet users (%)	Yes	Yes	Yes
Size dummies	Yes	Yes	Yes
Services dummy	Yes	Yes	Yes
Region FEs	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes
Observations	88,376	88,376	88,376
Countries	139	139	139

Notes: Probit estimates. Each column enters one demographic measure in the preferred specification (column 4 of Table 4). Old-age dependency, TFR, and pop. 65+ share are not included simultaneously due to high mutual correlations ($r = 0.997$ between old-age dep. and pop. 65+ share; $r = -0.69$ between old-age dep. and TFR). Standard errors clustered at the country level in parentheses. Average marginal effects (AMEs) in brackets. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5.3 Robustness Checks

Table 6 presents seven robustness checks (columns 2–8). In each case, the demographic coefficient on old-age dependency maintains its positive sign and remains statistically significant or close to significance.[5]

Keyword classifier (column 2). Using the traditional keyword-based measure (1,104 adopters, 33 fixed regular expressions) yields a coefficient of 0.014 (AME = 0.0004), signif-

[5]Controlling for the labor force participation rate among the working-age population (ages 15–64) leaves the old-age dependency coefficient essentially unchanged (0.0135, significant at the 5 percent level, against 0.0142 in the baseline). This finding addresses the concern that the old-age dependency ratio may partly reflect healthier, more active older populations rather than genuine labor scarcity: the demographic effect survives even after directly accounting for how much of the working-age population is in the labor force.

icant at the 5 percent level and virtually identical to the LLM baseline. This 96.0 percent agreement across classifiers confirms that the main result does not depend on the more flexible LLM approach.

Linear probability model (column 3). The linear probability model (i.e., OLS) estimate is 0.0009, significant at the 5 percent level and directly interpretable as a percentage-point marginal effect. This is quantitatively consistent with the probit AME of 0.0006, confirming that the probit functional form is not driving the result.

Lagged demographics (column 4). Replacing old-age dependency at the survey year with the value two years prior yields a coefficient of 0.016, significant at the 5 percent level.

Innovators only (column 5). Restricting the sample to firms that report any process or product innovation (28.4 percent of the full sample, 24,969 firms) yields a coefficient of 0.01. The coefficient is smaller in magnitude but remains positive and marginally significant. The AME is 0.0011, slightly larger than in the full sample, reflecting the higher baseline adoption rate among innovators.

This restriction addresses the composition of innovation among firms that already innovate, but a separate question is whether aging affects the propensity to innovate at all. A literature spanning macroeconomics and corporate finance suggests that population aging may depress innovation and firm dynamism more broadly (Jones, 2022; Peters and Walsh, 2026; Derrien et al., 2023; Aksoy et al., 2019), which would complicate interpretation of my composition-based result: if older economies innovate less overall, a positive coefficient on the automation share among innovators could coexist with, or even be a consequence of, a shrinking pool of innovating firms. I test this directly by regressing an indicator equal to one if the firm reports any process or product innovation (the same indicator used to construct the innovators-only subsample above) on old-age dependency and the preferred-specification controls. The coefficient is 0.018 (AME = 0.0055, $N = 88{,}376$), significant at the 10 percent level, and remains positive and significant across all four successive specifications. Aging is therefore associated with more overall innovation, not less, in this sample. This finding indicates that the automation-share result in Table 4 reflects a genuine compositional shift toward labor-saving innovation, rather than a mechanical consequence of a shrinking innovating population.

Excluding high-income OECD (column 6). Removing the 3,712 high-income OECD firms yields a coefficient of 0.015, significant at the 5 percent level and essentially unchanged from the baseline. This rules out the possibility that the result is entirely driven by the richest and oldest corner of the sample.

Weighted least squares specifications (columns 7 and 8). The country-equal-weighted WLS specification (column 8, inverse country sample-size weights) yields a coefficient of 0.0007, marginally significant, confirming the result when every country receives equal weight regardless of the sample size for each country. The survey-weighted WLS specification (column 7, WBES probability weights) attenuates the coefficient to insignificance. This attenuation is informative rather than troubling: WBES upweight underrepresented small firms in low-income countries, which have systematically lower adoption rates. The attenuation is therefore consistent with the heterogeneity results, which show that the demographic effect concentrates in large firms.

Firm age and foreign ownership (untabulated). I additionally check whether the result reflects omitted firm characteristics correlated with both demographic exposure and adoption. Adding ln(firm age), constructed from the survey year and the year the establishment began operations, leaves the OAD coefficient essentially unchanged (0.0142 versus 0.0151 at baseline, significant at the 5 percent level); older firms are themselves somewhat more likely to adopt (coefficient 0.0483), marginally significant. Adding an indicator for foreign ownership of 10 percent or more, the standard threshold used to define multinational-affiliated firms, also leaves the OAD coefficient unchanged (0.0151, significant at the 5 percent level); foreign-owned firms are substantially more likely to adopt (coefficient 0.1244, AME = 0.0051, significant at the 1 percent level), consistent with multinationals' greater access to capital and technology. Including both controls jointly (87,361 firms, 139 countries) leaves the OAD coefficient at 0.0142, significant at the 5 percent level. The demographic result is therefore not explained by older or foreign-affiliated firms being disproportionately located in higher-OAD countries.

Table 6: Robustness Checks

	(1) Baseline	(2) Keyword	(3) LPM	(4) Lag $t-2$	(5) Innovators	(6) Excl. HIC	(7) Survey wts	(8) Country wts
Old-age dep. ratio	0.0151**	0.0142**	0.0009**	0.0169**	0.0104*	0.0153**	0.0003	0.0007**
	(0.0071)	(0.0066)	(0.0004)	(0.0072)	(0.0055)	(0.0071)	(0.0002)	(0.0004)
	[0.0006]	[0.0004]		[0.0007]	[0.0012]	[0.0006]		
Outcome variable	LLM Broad	Kw. Broad	LLM Broad	LLM Broad	LLM Broad	LLM Broad	LLM Broad	LLM Broad
Estimator	Probit	Probit	LPM	Probit	Probit	Probit	WLS	WLS
Observations	88,376	88,376	88,376	88,376	24,969	84,664	88,376	88,376
Countries	139	139	139	139	139	137	139	139

Notes: All columns include the full set of controls from the preferred specification (column 4 of Table 4): ln(GDP per capita), internet users (%), size dummies, services dummy, region FEs, and year FEs. Column (1) is the baseline LLM broad measure (Probit). Column (2) uses the keyword classifier as an alternative measure (33 fixed regular expressions, 1,104 adopters), confirming 96.0% agreement with the LLM measure. Column (3) is a linear probability model (LPM) using the LLM broad outcome. Column (4) uses old-age dependency lagged by two years (old_age_dep at $t-2$). Column (5) restricts to firms reporting product innovation (h1=Yes) or process innovation (h5=Yes). Column (6) excludes high-income OECD countries. Columns (7) and (8) are WLS linear probability models; column (7) uses WBES survey sampling weights; column (8) uses country-equal weights (each firm weighted by the inverse of its country's sample size, normalized to mean 1). Probit AMEs in brackets; no AMEs reported for LPM or WLS (coefficients are already in percentage-point units). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5.4 Strict AI-Only Measure

Table 7 replicates the four successive specifications of Table 4 using *ai_process_strict*, which equals one only for Tier 1 classifications: firms whose innovation text contains explicit references to artificial intelligence or machine learning (ChatGPT, ML models, computer vision, NLP, and related terms). This measure captures 221 firms (0.2 percent of the sample) and is conservative by design, excluding any response that mentions automation or robotics without naming an AI or ML technology.

The demographic coefficient is positive and statistically significant across all four specifications. In the most parsimonious specification (column 1), the coefficient is 0.021, significant at the 1 percent level. In the preferred specification (column 4), it is 0.015, marginally significant, with an AME of 0.0001. The stability across specifications mirrors the pattern in the broad measure and confirms that the demographic-adoption link is not confined to the automation and robotics component; it runs through the adoption of technologies explicitly described as artificial intelligence or machine learning.

A notable contrast with the broad measure concerns the sector coefficient. For the strict AI-only measure, services firms are significantly *more* likely to adopt than manufacturing firms (coefficient 0.260), significant at the 1 percent level. This reversal reflects a compositional difference in the type of AI being adopted: explicit AI tools such as language models, sentiment analysis, and predictive analytics are disproportionately deployed in service-sector activities (customer relations, logistics, finance), whereas the automation and robotics component of the broad measure is concentrated in manufacturing. This pattern is consistent with the view that software-based AI and hardware-based automation have different sectoral profiles, both driven by the same aging-induced labor scarcity incentive.

Table 7: Strict AI-Only Measure: Explicit AI/ML Adoption

	(1)	(2)	(3)	(4)
Old-age dep. ratio	0.0214***	0.0129**	0.0152**	0.0147*
	(0.0057)	(0.0065)	(0.0068)	(0.0076)
	[0.0002]	[0.0001]	[0.0001]	[0.0001]
ln(GDP per capita, PPP)		0.2418**	0.1555	0.1653
		(0.0966)	(0.1258)	(0.1311)
		[0.0017]	[0.0011]	[0.0011]
Internet users (%)			0.0101	0.0094
			(0.0074)	(0.0076)
			[0.0001]	[0.0001]
Medium firm				0.2148***
				(0.0564)
				[0.0015]
Large firm				0.5190***
				(0.0692)
				[0.0035]
Services sector				0.2603***
				(0.0664)
				[0.0018]
Region FEs	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes
Observations	89,380	89,224	88,376	88,376
Countries	144	143	139	139
Pseudo R^2	0.117	0.123	0.128	0.155

Notes: Probit estimates. Dependent variable: *ai_process_strict* (=1 if firm's innovation text was classified as Tier 1 by the LLM: explicit AI or machine learning technology). 221 adopters, 0.2 percent of the full sample. Standard errors clustered at the country level in parentheses. Average marginal effects in brackets. Base categories: Small firm, Manufacturing sector, Sub-Saharan Africa, survey year 2022. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5.5 Heterogeneity Analysis

Table 8 examines heterogeneity along four dimensions. Each panel reports separate probit regressions on subsamples, using the preferred specification throughout.

Panel A: By Sector

For the broad adoption measure (automation and robotics combined with AI), the demographic effect is significant in manufacturing (coefficient 0.021, AME = 0.0012, N = 36,437) but not statistically significant in services (coefficient 0.010, N = 51,939). This asymmetry is consistent with the labor-substitution hypothesis: manufacturing involves repetitive, physical production tasks that are directly amenable to automation and robotics, and aligns with the broader literature (Acemoglu and Restrepo, 2022; Graetz and Michaels, 2018). It is important to note, however, that this pattern reflects the composition of the broad measure. As shown in Table 7, for the strict AI-only measure the sector pattern reverses: services firms are significantly *more* likely to adopt explicit AI tools than manufacturing firms, reflecting the greater applicability of software-based AI (language models, predictive analytics, sentiment analysis) in service-sector activities. The sector asymmetry in the broad measure is therefore driven by hardware automation and robotics in manufacturing.

Panel B: By Firm Size

The demographic effect varies substantially by firm size. The coefficient is strongest for large firms (100 or more employees): 0.024 (AME = 0.0016), significant at the 1 percent level. For small firms, the coefficient is 0.013 (AME = 0.0003), marginally significant. Medium firms show no significant demographic response (coefficient 0.010). The concentration of the effect among large firms is consistent with the literature on innovation and scale: the fixed costs of automation investment in equipment, integration, and worker retraining are more easily borne by firms with sufficient scale to spread these costs across a larger production base.

Panel C: By Country Income Level

I use the official IMF country classification to define three groups: low-income EMDEs, middle-income EMDEs, and advanced economies (AE). The demographic effect is positive and statistically significant in both developing-country groups: coefficient 0.038, marginally significant (AME = 0.0009, N = 20,535, 50 countries) for low-income economies and 0.023, significant at the 1 percent level (AME = 0.0007, N = 46,652, 61 countries) for middle-income economies. In advanced economies the point estimate is smaller in magnitude (0.009) and not statistically significant (N = 20,829, 27 countries). This finding is consistent with a ceiling effect: many advanced economies already have high old-age dependency, compressing the within-group variation that identification requires. The country-level standard deviation of the old-age dependency ratio within advanced economies is 5.4 percentage points, compared with 9.0 percentage points for middle-income EMDEs – a 40 percent narrower spread around a mean of 30 percent. With OAD clustered near the top of the global distribution and little cross-country dispersion remaining, identification of the demographic gradient is mechanically limited. More substantively, in economies with deep capital markets, mature technology ecosystems, and decades of automation experience, the marginal adoption decision is driven more by competition and firm strategy than by demographic pressure at the country level – OAD variation across advanced economies may not be the binding constraint on adoption once these complementary factors are in place.

The finding that the demographic effect is precisely identified across both low- and middle-income EMDE groups is therefore where the mechanism is most economically relevant. Even in low-income countries, where managerial capacity and access to capital are more constrained, aging raises the probability of AI and automation adoption – suggesting that labor scarcity is a sufficiently strong force to overcome these barriers. The concentration of the effect in EMDEs is where the policy stakes are also highest, as these economies face the steepest demographic transitions with the least institutional capacity to manage them.

5.6 Labor Market Obstacles Channel

To probe the labor scarcity channel more directly, I test whether demographic aging predicts the probability that a firm reports labor constraints as its primary business obstacle. The WBES asks each firm to identify the single biggest obstacle to its current operations from a list of fifteen categories: access to finance, access to land, business licensing and permits, corruption, courts, crime and theft, customs and trade regulations, electricity supply, inadequately educated workforce, labor regulations, political instability, practices of informal-sector competitors, tax administration, tax rates, and transport. If aging tightens aggregate labor markets, firms in older economies should be systematically more likely to name workforce skill gaps or labor regulations as their binding constraint.

Table 9 reports results. Column (1) uses as outcome an indicator equal to one if the firm named *inadequately educated workforce* as its biggest obstacle; column (2) broadens this to include firms that named *labor regulations* as their biggest obstacle. Columns (3) and (4) use a separate question in which firms rate how severely workforce skill gaps constrain their operations on a five-point scale (no obstacle, minor, moderate, major, or very severe). Column (3) is a binary indicator for a major or very severe rating; column (4) uses the full 0–4 numeric score estimated by OLS. Coverage of the severity question exceeds 98 percent of the sample.

A ten-percentage-point increase in the old-age dependency ratio raises the probability of naming an inadequately educated workforce as the biggest obstacle by 0.34 percentage

Table 8: Heterogeneity Analysis

Panel A: By sector

	Manufacturing	Services
Old-age dep. ratio	0.0205**	0.0097
	(0.0080)	(0.0060)
	[0.0012]	[0.0003]
Observations	36,437	51,939
Countries	139	139

Panel B: By firm size

	Small	Medium	Large
Old-age dep. ratio	0.0126*	0.0100	0.0237***
	(0.0066)	(0.0079)	(0.0076)
	[0.0003]	[0.0005]	[0.0016]
Observations	40,983	29,400	17,993
Countries	139	139	138

Panel C: By country income level

	Low income	Mid income	Advanced
Old-age dep. ratio	0.0375*	0.0229***	0.0090
	(0.0213)	(0.0083)	(0.0138)
	[0.0009]	[0.0007]	[0.0007]
Observations	20,535	46,652	20,829
Countries	50	61	27

Notes: Probit estimates. Each panel reports separate regressions on the indicated subsamples. All columns include the full preferred specification controls: ln(GDP per capita), internet users (%), region FEs, and year FEs, plus size dummies in Panel A, and a services dummy in Panel B. Panel C income groups follow the official classification: Low income = EMDE economies with low-income status; Middle income = EMDE economies without low-income status; Advanced = advanced economies (AE). Standard errors clustered at country level in parentheses. AMEs in brackets. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

points (column 1), significant at the 1 percent level, and raises the probability of naming either workforce skills or labor regulations by 0.32 percentage points (column 2), significant at the 5 percent level. The severity-scale results confirm the pattern: firms in older economies report significantly higher workforce skill constraints, both on the binary high-severity measure (column 3, AME = 0.69 percentage points) and on the full 0–4 scale estimated by OLS (column 4), both significant at the 1 percent level. These results are suggestive of an intermediate link in the mechanism chain: demographic aging is associated with the kind of labor market tightness that firms perceive and report as a binding constraint, consistent with a channel through which aging raises the incentive to adopt labor-saving technologies. I present this as supportive rather than definitive evidence, since obstacle categories such as "inadequately educated workforce" may also reflect skill mismatches, educational-system quality, or other factors not specific to demographic labor scarcity. Read alongside the main adoption results in Table 4, this finding is consistent with, though does not on its own establish, a two-step process in which aging first tightens labor markets and firms then respond by adopting AI and automation.

To examine whether the aging effect concentrates among firms for which labor substitution pressure is most acute, I interact the old-age dependency ratio with a firm-level measure of labor intensity. Labor intensity is defined as total labor cost (wages, salaries, and bonuses) divided by total annual sales, log-transformed and winsorized at the first and ninety-ninth percentiles. The ratio is dimensionless; 92.3 percent of firms in the estimation sample report both items. Because AI adoption may itself reduce labor costs, I interpret this exercise as documenting heterogeneity consistent with the substitution mechanism rather than as a causal identification of that channel.

Table 10 reports the results. Column (1) replicates the preferred specification (column 4 of Table 4) on the subsample of firms with non-missing labor intensity, confirming that the main OAD coefficient is unchanged at 0.0147, significant at the 5 percent level. Column (2) adds the log labor-intensity ratio as a direct control; the OAD coefficient is virtually unaffected (0.0146), significant at the 5 percent level, and the labor intensity main effect is positive but imprecisely estimated. Column (3) adds the interaction term. The OAD coefficient rises to 0.0197, significant at the 1 percent level, and the interaction OAD $\times$ ln(labor intensity) is 0.0031, also significant at the 1 percent level. This finding is consistent with the aging effect being concentrated among firms whose labor costs represent a larger share of sales – the firms for which demographic tightening of labor markets most directly raises the incentive to automate. Aging economies therefore generate adoption pressure unevenly across firms,

Table 9: Demographic Aging and Labor Constraint Reporting

	(1)	(2)	(3)	(4)
Outcome:	Skill shortage biggest obstacle	Labor obstacle biggest obstacle	High severity (rating $\geq$ major)	Severity score (0–4 scale)
Estimator:	Probit	Probit	Probit	OLS
Old-age dependency ratio	0.0197***	0.0151**	0.0260***	0.0295***
	(0.0068)	(0.0069)	(0.0087)	(0.0104)
Average marginal effect	0.0034	0.0032	0.0069	—
Controls	Yes	Yes	Yes	Yes
Region FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	88,376	88,376	88,376	87,446
Countries	139	139	139	139

Notes: Dependent variables capture firm-level reporting of labor constraints from the WBES obstacle battery. Columns (1) and (2) use the question asking each firm to identify its single biggest obstacle from fifteen categories. Column (1): indicator equal to one if the firm chose "Inadequately Educated Workforce" as its biggest obstacle. Column (2): indicator equal to one if the firm chose either "Inadequately Educated Workforce" or "Labor Regulations." Columns (3) and (4) use a separate question asking firms to rate how severely workforce skill gaps constrain operations on a five-point scale (0 = no obstacle; 1 = minor; 2 = moderate; 3 = major; 4 = very severe); Column (3): indicator equal to one if the firm gave a rating of major or very severe (score ≥ 3). Column (4): the continuous 0–4 score estimated by OLS. All columns include the full preferred-specification controls: log GDP per capita, internet penetration, firm size dummies, services indicator, region fixed effects, and survey-year fixed effects. Standard errors (in parentheses) are clustered at the country level. $^{*}\, p < 0.10$; $^{**}\, p < 0.05$; $^{***}\, p < 0.01$.

Table 10: Labor Intensity and the Aging–Adoption Channel

	(1)	(2)	(3)
Outcome:	Process AI/automation adoption (broad)		
Estimator:	Probit		
Old-age dependency ratio	0.0147**	0.0146**	0.0197***
	(0.0073)	(0.0073)	(0.0074)
Average marginal effect	[0.0006]	[0.0006]	—
ln(labor intensity)		0.0176	−0.0408
		(0.0155)	(0.0286)
OAD × ln(labor intensity)			0.0031***
			(0.0012)
Controls	Yes	Yes	Yes
Region FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Observations	81,577	81,577	81,577
Countries	139	139	139

Notes: Labor intensity is total labor cost (wages, salaries, and bonuses) divided by total annual sales, log-transformed and winsorized at the first and ninety-ninth percentiles of the estimation sample. Column (1) replicates the preferred specification on the 92.3 percent of firms with non-missing labor intensity. Column (2) adds ln(labor intensity) as a control. Column (3) adds the interaction OAD × ln(labor intensity); average marginal effects are omitted for the interaction specification because they depend on the value of the moderating variable. All columns include the full preferred-specification controls: log GDP per capita, internet penetration, firm size dummies, services indicator, region fixed effects, and survey-year fixed effects. Standard errors (in parentheses) are clustered at the country level. $^{*}\,p < 0.10$; $^{**}\,p < 0.05$; $^{***}\,p < 0.01$.

in a pattern consistent with capital–labor substitution.

5.7 Wage and Employment Channel

The labor-substitution mechanism assumes an intermediate step: aging raises the cost or scarcity of labor, which in turn induces firms to substitute toward labor-saving technology. I test this step directly using WBES labor cost and sales data. A natural first measure is unit labor cost, defined as total labor cost (wages, salaries, and bonuses) divided by total annual sales; because both components are in local currency, the ratio is dimensionless and comparable across countries without any exchange-rate adjustment, and it is the same underlying construction as the labor-intensity measure used in Table 10, here used as an outcome rather than a moderator. Regressing this ratio on OAD, the unrestricted correlation is positive and significant at the 5 percent level (coefficient 0.0087), but it collapses to approximately zero and loses significance once income is controlled for. Taken at face value, this looks like a null result for the wage channel. But unit labor cost is algebraically the ratio of wage per worker to labor productivity, so a flat ratio is also consistent with wages and productivity rising together rather than with no wage pressure at all. To distinguish these possibilities, I decompose unit labor cost into its two components, converting firm-level sales and labor cost from local currency into purchasing-power-parity (PPP) adjusted international dollars using the World Bank's PPP conversion factor (the same underlying WDI series used to construct ln(GDP per capita, PPP) elsewhere in the paper), matched at the country-year level.

Table 11 shows that all three components move as the labor-substitution mechanism predicts. A ten-percentage-point increase in the old-age dependency ratio is associated with approximately 16 percent fewer workers for the same level of real output (column 1), approximately 44 percent higher output per worker (column 2), and approximately 41 percent higher wages per worker (column 3), all significant at the 1 percent level. This finding resolves the apparent unit-labor-cost puzzle above: because wage per worker and labor productivity rise by similar magnitudes (0.0406 versus 0.0442), they largely offset each other in the ratio, leaving the wage-bill share of sales roughly flat even as both wages and productivity respond strongly to demographic aging. Firms in older economies do not simply pay more for the same workforce; they employ fewer workers, each of whom produces and earns more, consistent with substitution away from labor and toward capital- and technology-intensive production as labor becomes scarcer and more expensive.

Table 11: Old-Age Dependency, Employment, Productivity, and Wages

Outcome:	(1) ln(employment) holding output fixed	(2) ln(labor productivity) = ln(sales/employment)	(3) ln(wage per worker) = ln(labor cost/employment)
Old-age dependency ratio	-0.0163^{***}	0.0442^{***}	0.0406^{***}
	(0.0045)	(0.0100)	(0.0076)
ln(sales, PPP)	0.4590^{***}		
	(0.0170)		
Controls	Yes	Yes	Yes
Region FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Observations	84,094	84,094	82,848
Countries	139	139	139
R^2	0.581	0.306	0.520

Notes: OLS, country-clustered standard errors in parentheses. Sales and labor cost are converted from local currency to PPP-adjusted international dollars using the WDI PPP conversion factor (PA.NUS.PPP), matched to each firm's country and survey year (falling back up to two prior years if unavailable), then divided by employment (size_num) and winsorized at the first and ninety-ninth percentiles. Column (1) regresses log employment on OAD controlling for log PPP-adjusted sales, so the OAD coefficient captures how many workers a firm needs for a given level of real output. Columns (2) and (3) regress log labor productivity and log wage per worker, respectively, on OAD. All columns include ln(GDP per capita, PPP), internet penetration, a services-sector indicator, region fixed effects, and survey-year fixed effects; firm-size dummies are excluded because they are constructed from employment, the outcome (or a component of the outcome) in every column. $^{*}\,p < 0.10$; $^{**}\,p < 0.05$; $^{***}\,p < 0.01$.

6 Instrumental Variable Estimates

While the baseline results establish a robust correlation between aging and adoption, a key concern is whether this relationship reflects causal effects of demographic structure or omitted country characteristics that independently promote both older populations and technology adoption. I address this using an instrumental variable strategy that exploits predetermined cohort structure to isolate variation in aging that is orthogonal to current economic conditions.

The probit and LPM estimates in Section 5 are conditional correlations. Two concerns motivate a causal identification strategy. First, classical measurement error in the old-age dependency ratio (which is itself constructed from census and survey data) attenuates the OLS coefficient toward zero. Second, unobserved country characteristics that simultaneously drive demographic aging and technology adoption, such as institutional quality or historical human capital accumulation, could bias estimates in either direction. I address both concerns using an instrumental variable strategy following Acemoglu and Restrepo (2022).

6.1 Instrument Construction

The instrument exploits the mechanical link between a country's age structure at a point in the past and its old-age dependency ratio today. People who were 40–44 years old in 2000 will be 65–69 in 2025 regardless of intervening economic conditions. If a country happened to have a large cohort of 40-somethings in 2000, its elderly population in 2025 was largely predetermined by that demographic fact.

Formally, for each country c I construct a predicted old-age dependency ratio using UN World Population Prospects (2024) data on population by five-year age group in BASE_YEAR = 2000. For each seed cohort aged a in 2000, I project the surviving population 25 years forward by multiplying by a globally-averaged survival rate $S(a)$:

$$S(a) = \frac{\text{World population aged } a + 25 \text{ in } 2025}{\text{World population aged } a \text{ in } 2000} \tag{3}$$

Using *global* rather than country-specific survival rates is deliberate: it ensures that $\widehat{\text{OAD}}_c$ varies only through initial cohort sizes, not through differential mortality improvements that may themselves correlate with economic development.[6] The predicted old-age dependency ratio is then:

[6] This does not fully resolve the underlying identification concern, since it might relocate rather than eliminate it: if a country's year-2000 cohort structure is itself correlated with long-run institutional, educational, or technological trajectories that independently affect adoption today, the exclusion restriction is still violated.

$$\widehat{\text{OAD}}_c = \frac{\sum_{a\in\mathcal{E}} N_c(a,\, 2000) \cdot S(a)}{\sum_{a\in\mathcal{W}} N_c(a,\, 2000) \cdot S(a)} \times 100 \tag{4}$$

where $N_c(a, 2000)$ is the population of country c aged a in 2000 (in thousands), $\mathcal{E} = \{40, 45, \ldots, 95\}$ are the seed cohort ages whose members will be aged 65 or above by 2025, and $\mathcal{W} = \{0, 5, \ldots, 35\}$ are the seed cohort ages whose members will be aged 25–64 in 2025. Because people aged 15–24 in 2025 were born after 2000, they are not present in the base-year data; the denominator therefore covers ages 25–64 rather than the full 15–64 working-age range. The instrument is constructed for 144 countries and covers all 139 countries in the preferred specification.

The exclusion restriction requires that $\widehat{\text{OAD}}_c$, conditional on income, digital infrastructure, region, and year fixed effects, affects firm-level AI and automation adoption only through current demographic aging. The cohort sizes determining $\widehat{\text{OAD}}_c$ were fixed by fertility decisions and mortality conditions prevailing in the 1950s and 1960s, long before AI technologies existed. Conditional on current log GDP per capita and internet penetration, there is no plausible direct channel from the size of birth cohorts five decades ago to whether a firm in 2022–2025 adopted AI or automation, other than through the current age structure of the workforce. Importantly, this identification strategy is identical in spirit to that of Acemoglu and Restrepo (2022), who instrument current demographic aging with lagged cohort sizes and face the same exclusion restriction assumption in a peer-reviewed setting. As in their paper, the condition cannot be tested directly; however, the stability of the coefficient across income groups and world regions, documented in Section 5.5, is consistent with the exclusion restriction holding.

6.2 First-Stage Results

Table 12 reports the first-stage regression of the actual old-age dependency ratio on the predicted value, with the full set of controls and fixed effects from the preferred specification. The instrument is very strongly correlated with actual aging: the coefficient on $\widehat{\text{OAD}}_c$ is 1.057, significant at the 1 percent level, implying that a one-percentage-point increase in predicted OAD is associated with a 1.058-percentage-point increase in actual OAD. The first-stage R^2 is 0.957. The cluster-robust F-statistic on the excluded instrument is approximately 413, far above the conventional threshold of 10 for strong instruments. Weak instruments are not a concern here.

Table 12: First Stage: Predicted OAD and Actual OAD

	Dependent variable: Old-age dep. ratio
Predicted OAD ($\widehat{\text{OAD}}_c$)	1.0571***
	(0.0520)
ln(GDP per capita)	Yes
Internet users (%)	Yes
Size dummies	Yes
Services dummy	Yes
Region FEs	Yes
Year FEs	Yes
Observations	88,376
Countries	139
R^2	0.957
First-stage F	413.0

Notes: OLS regression of actual old-age dependency ratio on predicted OAD (constructed from UN WPP 2000 cohort sizes and global average 25-year survival rates, equation (4)) and the full set of controls from the preferred specification. Standard errors clustered at the country level in parentheses. First-stage F is the cluster-robust partial F-statistic on the excluded instrument ($\approx t^2$). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6.3 IV Estimates

Table 13 presents the IV-LPM (2SLS) estimates alongside the OLS-LPM for direct comparison. Because IV probit requires numerical integration that is sensitive to the binary outcome distribution across 138 countries, I follow the standard practice of reporting IV results as a linear probability model (Angrist and Pischke, 2009); the OLS-LPM in column (1) confirms the probit marginal effects reported in Table 4 are quantitatively similar.

The IV-LPM coefficient on old-age dependency is 0.000900, significant at the 1 percent level, compared with the OLS-LPM coefficient of 0.000698, significant at the 5 percent level. The IV estimate is approximately 29 percent larger than OLS. A ten-percentage-point increase in the old-age dependency ratio raises the adoption probability by 0.090 percentage points under IV, compared with 0.070 percentage points under OLS. The direction of the bias is consistent with classical attenuation: measurement error in the old-age dependency ratio (which is itself constructed from imperfect census and survey data, particularly in low-income countries) biases OLS toward zero, and the IV corrects for this. This finding mirrors Acemoglu and Restrepo (2022), who also report that their IV estimates of the effect of aging on robot adoption are slightly larger than the corresponding OLS estimates.

The IV estimate is significant at the 1 percent level and exceeds the OLS estimate, ruling out the most concerning form of omitted variable bias, which would predict an upward-biased OLS coefficient. The evidence points instead toward the demographic-adoption link being, if anything, understated by the probit estimates in Section 5.

Table 13: IV-LPM Estimates: Instrumenting Demographic Aging

	(1) OLS-LPM	(2) IV-LPM (2SLS)
Old-age dep. ratio	0.000698**	0.000900***
	(0.000274)	(0.000270)
ln(GDP per capita)	Yes	Yes
Internet users (%)	Yes	Yes
Size dummies	Yes	Yes
Services dummy	Yes	Yes
Region FEs	Yes	Yes
Year FEs	Yes	Yes
Observations	88,376	88,376
Countries	139	139
First-stage F	—	413.0

Notes: Linear probability model estimates. Dependent variable: `ai_process_broad`. Column (1) is OLS. Column (2) instruments old-age dependency with the predicted OAD constructed from 2000 UN WPP cohort sizes (equation (4)). Standard errors clustered at the country level in parentheses. 10pp increase in OAD raises adoption probability by 0.070 pp (OLS) and 0.090 pp (IV). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

7 Product AI and Automation Adoption

Having established the process adoption result and confirmed it causally in Section 6, this section extends the analysis to a distinct margin of adjustment: whether aging also predicts firms developing or offering AI-powered *products*. Process adoption is driven by the cost and scarcity of labor as an input; product adoption may reflect demand-side incentives, as firms in aging economies develop labor-saving solutions for their own customers, or broader innovation effects associated with technological maturity. I classify product adopters using the same LLM procedure applied to WBES product innovation text (fields h3x, h4x, and h32x; see Section 3).

7.1 Descriptive Statistics

Product AI adoption is rarer than process adoption: 589 firms (0.66 percent of the wave-filtered sample of 89,380) qualify as product AI adopters under the broad measure (Tier 1 or Tier 2), compared with 1,656 firms (1.85 percent) for process adoption. Under the strict measure (Tier 1 AI only), 241 firms are product adopters. The overlap between the two channels is small: only 82 firms report both process and product AI adoption, confirming that these largely represent distinct firm populations responding to different incentives.

The sector distribution of product adopters differs markedly from process adopters. While process adoption is concentrated in manufacturing (64.8 percent of process adopters), product adopters are more evenly split between services (67.1 percent) and manufacturing (32.9 percent). Adoption rates tell the same story: the product adoption rate is 0.75 percent in services and 0.53 percent in manufacturing, compared with rates of 1.11 percent and 2.91 percent respectively for process adoption. This pattern is consistent with the intuition that service firms, which face limited scope for automating their own production, can nonetheless develop and offer AI-powered products or platforms to labor-constrained clients.

By firm size, adoption rates increase with scale: 0.53 percent for small firms, 0.71 percent for medium firms, and 0.87 percent for large firms. The regional distribution mirrors that of process adoption: Europe and Central Asia leads with 1.38 percent, followed by East Asia and Pacific (0.67 percent) and Latin America and the Caribbean (0.76 percent). South Asia has the lowest rate at 0.04 percent.

7.2 Baseline Results

Table 14 presents progressive probit specifications for product AI adoption, directly paralleling Table 4 for process adoption. Column (1) includes only the old-age dependency ratio and region and year fixed effects. Columns (2) through (4) add controls sequentially, with column (4) the preferred specification.

The demographic gradient for product adoption is positive and statistically significant across all specifications. In the preferred specification (column 4), a one-percentage-point increase in the old-age dependency ratio raises the probability of product AI adoption by 0.0128, significant at the 5 percent level, with an average marginal effect of 0.0002. This is somewhat smaller in absolute magnitude than the corresponding process adoption coefficient (0.0151), which is expected given that product adoption is roughly one-third as prevalent.

The result is robust to the inclusion of income, digital infrastructure, size, and sector controls, and the coefficient is stable from column (2) onward.

Table 14: Product AI Adoption: Baseline Results

	(1)	(2)	(3)	(4)
Old-age dep. ratio	0.0183***	0.0128**	0.0130**	0.0128**
	(0.0052)	(0.0056)	(0.0057)	(0.0056)
	[0.0003]	[0.0002]	[0.0002]	[0.0002]
ln(GDP per capita)		Yes	Yes	Yes
Internet users (%)			Yes	Yes
Size dummies				Yes
Services dummy				Yes
Region FEs	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes
Observations	89,380	89,224	88,376	88,376
Countries	144	143	139	139

Notes: Probit estimates. Dependent variable: `ai_product_broad` (= 1 if firm mentioned AI or automation in product innovation text h3x/h4x/h32x, Tier 1 or Tier 2). AMEs (average marginal effects) in brackets. Standard errors clustered at country level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

7.3 Strict AI-Only Measure (Product)

Table 15 replicates the preferred specification using the strict product adoption measure (*ai_product_strict* = 1 for Tier 1 classifications only), which requires an explicit mention of AI or machine learning technology and excludes automation and robotics without an AI label. This captures 241 firms (0.27 percent of the sample). The demographic gradient is positive and significant across all four specifications, and if anything stronger than for the broad measure: the preferred-specification coefficient is 0.0156, significant at the 5 percent level (AME = 0.0001), indicating that aging predicts not only broad automation adoption in products but specifically the development or offering of AI-powered products.

Table 15: Strict AI-Only Measure: Explicit AI/ML Product Adoption

	(1)	(2)	(3)	(4)
Old-age dep. ratio	0.0215***	0.0151**	0.0148**	0.0156**
	(0.0055)	(0.0061)	(0.0061)	(0.0061)
	[0.0002]	[0.0001]	[0.0001]	[0.0001]
ln(GDP per capita)		Yes	Yes	Yes
Internet users (%)			Yes	Yes
Size dummies				Yes
Services dummy				Yes
Region FEs	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes
Observations	89,380	89,224	88,376	88,376
Countries	144	143	139	139
Pseudo R^2	0.105	0.108	0.109	0.138

Notes: Probit estimates. Dependent variable: *ai_product_strict* (= 1 if firm's product innovation text was classified as Tier 1: explicit AI or machine learning technology). 241 adopters, 0.27 percent of the full sample. Standard errors clustered at the country level in parentheses. Average marginal effects in brackets. Base categories: Small firm, Manufacturing sector, Sub-Saharan Africa, survey year 2022. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

7.4 Heterogeneity by Sector and Firm Size

Table 16 presents heterogeneity analysis by sector (Panel A) and firm size (Panel B).

The sector results are striking in their contrast with process adoption. For process adoption (Table 8, Panel A), the demographic effect is significant only in manufacturing and is absent in services. For product adoption, the effect is significant in *both* manufacturing (coefficient 0.0148) and services (coefficient 0.0123), both significant at the 5 percent level. The magnitudes are similar across sectors, and the AMEs are equal at 0.0002 in both cases. This symmetric pattern is consistent with a demand-side or innovation channel: service firms, facing scarcity of their labor inputs due to aging, may be induced to innovate by offering AI-powered services to other similarly constrained firms or consumers.

For firm size, the pattern is also somewhat different from process adoption. Large firms show the strongest effect (coefficient 0.0186, significant at the 5 percent level), consistent with fixed costs of product innovation favoring scale. Medium firms also show a significant positive effect (coefficient 0.0113), significant at the 5 percent level. Small firms have a positive and marginally significant coefficient (0.0111). In contrast, for process adoption, medium firms were insignificant and large firms drove the result almost exclusively.

Table 16: Product AI and Automation Adoption: Heterogeneity by Sector and Firm Size

Panel A: By sector

	Manufacturing	Services
Old-age dep. ratio	0.0143**	0.0123**
	(0.0071)	(0.0059)
	[0.0002]	[0.0002]
Observations	36,437	51,939
Countries	139	139

Panel B: By firm size

	Small	Medium	Large
Old-age dep. ratio	0.0111*	0.0113**	0.0186**
	(0.0065)	(0.0056)	(0.0079)
	[0.0002]	[0.0002]	[0.0004]
Observations	40,983	29,400	17,993
Countries	139	139	138

Notes: Probit estimates. Each panel reports separate regressions on the indicated subsamples. All columns include the full preferred specification controls: ln(GDP per capita), internet users (%), region FEs, and year FEs, plus size dummies in Panel A and a services dummy in Panel B. AMEs in brackets. Standard errors clustered at country level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

7.5 Causal Estimates

Table 17 presents IV-LPM estimates for product adoption, using the same Acemoglu-Restrepo instrument described in Section 6.1. The first stage is identical to that reported in Table 12: the instrument is strongly predictive of the old-age dependency ratio (partial $F = 413$) and the first-stage coefficient is unchanged since I instrument the same endogenous variable in the same sample.

The IV-LPM coefficient on old-age dependency for product adoption is 0.000419 , approximately 22 percent larger than the OLS-LPM coefficient of 0.000342. As with process adoption, the IV exceeds OLS, consistent with classical attenuation bias in the OLS estimate. A ten-percentage-point increase in the old-age dependency ratio raises the probability of product AI adoption by 0.042 percentage points under IV, compared with 0.034 percentage points under OLS. Both estimates are significant at the 5 percent level.

The upward IV correction is proportionally similar to that found for process adoption (29 percent for process versus 22 percent for product), suggesting that measurement error in the old-age dependency ratio attenuates the demographic gradient comparably across both channels.

Table 17: IV-LPM Estimates: Product AI Adoption

	(1) OLS-LPM	(2) IV-LPM (2SLS)
Old-age dep. ratio	0.000342**	0.000419**
	(0.000156)	(0.000168)
ln(GDP per capita)	Yes	Yes
Internet users (%)	Yes	Yes
Size dummies	Yes	Yes
Services dummy	Yes	Yes
Region FEs	Yes	Yes
Year FEs	Yes	Yes
Observations	88,376	88,376
Countries	139	139
First-stage F	—	413.0

Notes: Linear probability model estimates. Dependent variable: `ai_product_broad`. Column (1) is OLS. Column (2) instruments old-age dependency with the predicted OAD constructed from 2000 UN WPP cohort sizes (equation (4)). Standard errors clustered at the country level in parentheses. A 10pp increase in OAD raises product adoption probability by 0.034 pp (OLS) and 0.042 pp (IV). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

8 Discussion and Conclusion

This paper provides firm-level evidence that demographic aging is an important driver of AI and automation adoption. Using data on nearly 90,000 firms across 144 countries, I show that firms operating in older economies are systematically more likely to adopt labor-saving technologies. This relationship is robust across specifications and supported by an instrumental variable strategy exploiting predetermined demographic structure.

The results point to two distinct but complementary channels through which aging shapes technology adoption. The first is a process adoption channel: firms respond to rising labor scarcity by adopting AI, automation, and robotics to reduce labor input in production. A ten-point rise in the old-age dependency ratio raises the probability of process adoption by 0.15 percentage points (Table 4, column 4) – modest in absolute terms but large relative to the 1.9 percent base rate. This channel is concentrated in manufacturing (coefficient 0.021, Table 8, Panel A), with no statistically discernible effect in services, consistent with a substitution margin driven by routinized production tasks where robots and automated equipment can most directly replace manual labor. The second is a product adoption channel: firms in aging economies are also more likely to develop and offer AI-enabled products and services. The baseline coefficient is smaller (0.013, Table 14, column 4), but its sectoral reach is symmetric – manufacturing (0.014) and services (0.012) are indistinguishable in magnitude and both statistically significant (Table 16, Panel A) – consistent with a demand-side mechanism in which firms innovate to serve labor-constrained customers across the full economy. The two channels also identify largely non-overlapping firm populations: only 82 firms appear in both adopter groups, confirming that aging operates through two structurally separate margins rather than a single set of highly innovative firms. Instrumental variable estimates reinforce both results, with the IV coefficient exceeding OLS by 29 percent for process adoption and 23 percent for product adoption, consistent in both cases with attenuation from measurement error in the old-age dependency ratio.

Section 5.6 tests the intermediate link in this mechanism chain directly: demographic aging predicts the probability that a firm reports labor constraints as its primary business obstacle, providing reduced-form evidence that aging tightens labor markets in a way that firms perceive and act upon (Table 9).

These results highlight that demographic change affects not only the level of economic activity, but also the direction of technological change. Aging shifts relative factor scarcities, increasing the incentives for firms to substitute away from labor and toward capital- and technology-intensive production methods, while simultaneously creating new market opportunities for labor-saving products. The result is a broad-based push toward automation that operates through both the production and product margins.

The implications for labor markets and policy are nuanced. On the one hand, automation may partially offset the economic consequences of aging by sustaining productivity growth in the face of declining labor supply. On the other hand, the adjustment is uneven. The concentration of process adoption in large firms and in sectors with routinized tasks suggests that displacement pressures may be localized but intense, while the broader reach of product innovation implies that AI-related job change may extend across a wider set of occupations. Moreover, the weaker demographic response among small firms points to the importance of complementary conditions, such as access to finance, digital infrastructure, and human capital, in determining whether demographic pressures translate into technology adoption.

More broadly, the results suggest that demographics are not merely a consequence of technological change but an active force shaping its direction. As populations continue to age across both advanced and emerging economies, the incentives to adopt AI and automation are likely to intensify. A complementary literature examines the reverse relationship, how AI adoption itself reshapes labor market outcomes for older workers in aging advanced economies (Pizzinelli and Tavares, 2025); integrating both directions of causality is a natural direction for future work. Understanding how these incentives interact with firm capabilities, sectoral characteristics, and policy environments will be central to assessing whether the gains from technological progress are broadly shared or increasingly concentrated.

References

Acemoglu, Daron and Pascual Restrepo, "Demographics and Automation," *Review of Economic Studies*, 2022, *89* (1), 1–44.

_ **and Simon Johnson**, *Power and Progress: Our Thousand-Year Struggle over Technology and Prosperity*, New York: Basic Books, 2023.

Aiyar, Shekhar, Christian Ebeke, and Xiaobo Shao, "The Impact of Workforce Aging on European Productivity," IMF Working Paper WP/16/238, International Monetary Fund 2016.

Akcigit, Ufuk and William R. Kerr, "Growth through Heterogeneous Innovations," *Journal of Political Economy*, 2018, *126* (4), 1374–1443.

Aksoy, Yunus, Henrique S. Basso, Ron P. Smith, and Tobias Grasl, "Demographic Structure and Macroeconomic Trends," *American Economic Journal: Macroeconomics*, 2019, *11* (1), 193–222.

Angrist, Joshua D. and Jörn-Steffen Pischke, *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton, NJ: Princeton University Press, 2009.

Ash, Elliott, Stephen Hansen, Yabra Muvdi, and Claudia Marangon, "Large Language Models in Economics," in Shlomo Weber and Victor Ginsburgh, eds., *The Palgrave Handbook of Economics and Language*, Cham: Palgrave Macmillan, 2026, pp. 191–210.

Asirvatham, Hemanth, Elliott Mokski, and Andrei Shleifer, "GPT as a Measurement Tool," NBER Working Paper 34834, National Bureau of Economic Research 2026.

Autor, David H., "Why Are There Still So Many Jobs? The History and Future of Workplace Automation," *Journal of Economic Perspectives*, 2015, *29* (3), 3–30.

_ **, David Dorn, and Gordon H. Hanson**, "Untangling Trade and Technology: Evidence from Local Labour Markets," *Economic Journal*, 2015, *125* (584), 621–646.

Bonfiglioli, Alessandra, Rosario Crinò, Harald Fadinger, and Gino Gancia, "Robot Imports and Firm-Level Outcomes," *The Economic Journal*, 2024, *134* (664), 3428–3444.

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond, "Generative AI at Work," *Quarterly Journal of Economics*, 2025, *140* (2), 889–941.

Calvino, Flavio and Luca Fontanelli, "A Portrait of AI Adopters Across Countries: Firm Characteristics, Assets Complementarities and Productivity," OECD Science, Technology and Industry Working Papers 2023/02, OECD Publishing, Paris 2023.

_ , Chiara Criscuolo, Luca Marcolin, and Mariagrazia Squicciarini, "Identifying and Characterising AI Adopters: A Novel Approach Based on Big Data," OECD Science, Technology and Industry Working Papers 2022/06, OECD Publishing, Paris 2022.

Charles, Kerwin Kofi, Erik Hurst, and Mariel Schwartz, "The Transformation of Manufacturing and the Decline in U.S. Employment," in Martin Eichenbaum and Jonathan A. Parker, eds., *NBER Macroeconomics Annual 2018*, Vol. 33, University of Chicago Press, 2019, pp. 307–372.

Cohen, Wesley M., "Empirical Studies of Innovative Activity," in Paul Stoneman, ed., *Handbook of the Economics of Innovation and Technological Change*, Oxford: Blackwell, 1995, pp. 182–264.

Comin, Diego A. and Martí Mestieri, "If Technology Has Arrived Everywhere, Why Has Income Diverged?," *American Economic Journal: Macroeconomics*, 2018, *10* (3), 137–178.

Derrien, François, Ambrus Kecskes, and Phuong-Anh Nguyen, "Labor Force Demographics and Corporate Innovation," *Review of Financial Studies*, 2023, *36* (7), 2797–2838.

Engbom, Niklas, "Misallocative Growth," Working Paper, New York University 2024.

Feyrer, James, "Demographics and Productivity," *Review of Economics and Statistics*, 2007, *89* (1), 100–109.

Friedberg, Leora, "The Impact of Technological Change on Older Workers: Evidence from Data on Computer Use," *Industrial and Labor Relations Review*, 2003, *56* (3), 511–529.

Graetz, Georg and Guy Michaels, "Robots at Work," *Review of Economics and Statistics*, 2018, *100* (5), 753–768.

Habakkuk, H. J., *American and British Technology in the Nineteenth Century: The Search for Labour-Saving Inventions*, Cambridge: Cambridge University Press, 1962.

International Monetary Fund, "World Economic Outlook, April 2025: Chapter 2, The Rise of the Silver Economy: Global Implications of Population Aging," Technical Report, International Monetary Fund, Washington, DC 2025.

Jones, Charles I., "The End of Economic Growth? Unintended Consequences of a Declining Population," *American Economic Review*, 2022, *112* (11), 3489–3527.

Koch, Michael, Ilya Manuylov, and Marcel Smolka, "Robots and Firms," *Economic Journal*, 2021, *131* (638), 2553–2584.

Lehr, Nils H., "Innovation in an Aging Economy," 2023. Boston University, mimeo.

Lordan, Grace and David Neumark, "People versus Machines: The Impact of Minimum Wages on Automatable Jobs," *Labour Economics*, 2018, *52*, 40–53.

Ludwig, Jens, Sendhil Mullainathan, and Ashesh Rambachan, "Large Language Models: An Applied Econometric Framework," *Annual Review of Economics*, 2026, *18.* Forthcoming.

Maestas, Nicole, Kathleen J. Mullen, and David Powell, "The Effect of Population Aging on Economic Growth, the Labor Force, and Productivity," *American Economic Journal: Macroeconomics*, 2023, *15* (2), 306–332.

Margo, Robert A., "The Labor Force in the Nineteenth Century," in Stanley L. Engerman and Robert E. Gallman, eds., *The Cambridge Economic History of the United States, Volume II: The Long Nineteenth Century*, Cambridge: Cambridge University Press, 2000, chapter 5, pp. 207–243.

Peters, Michael and Conor Walsh, "Population Growth and Firm Dynamics," *Journal of Political Economy Macroeconomics*, 2026. Forthcoming.

Pizzinelli, Carlo and Marina M. Tavares, "AI and the Future of Work in an Aging Economy," in Satyan Kolluri, Olivia S. Mitchell, and Nikolai Roussanov, eds., *The Future of Healthy Aging and Successful Retirement*, Oxford: Oxford University Press, 2025.

Udomsaph, Charles, "AI Adoption in MENAAP: Evidence from Enterprise Surveys," 2025. Mimeo.

United Nations, Department of Economic and Social Affairs, Population Division, "World Population Prospects 2024," 2024. Online edition. Accessed April 2026.

Appendix: Data and Classification Details

A.1 LLM Classification Prompts

The LLM classifier described in Section 3.2 submits each firm's open-ended innovation text to Claude (Anthropic's Claude-Opus-4-6) with temperature set to zero and the system prompt reproduced below. A separate but structurally identical prompt is used for process and product innovation text, differing only in a one-sentence description of what the text describes. In both cases the model returns a single JSON object with the assigned tier and a one-sentence justification; no other output is permitted.

Process innovation prompt (applied to fields h6x, h7x, h62x):

```
You are a research assistant classifying firm survey responses for an
economics study on AI and automation adoption.

Classify each firm response into one of three tiers:

Tier 1 -- Explicit AI or machine learning technology:  Includes:
artificial intelligence, machine learning, deep learning, neural
networks, large language models, generative AI, ChatGPT, GPT, computer
vision, natural language processing (NLP), sentiment analysis,
image/facial/ speech recognition, anomaly detection, robotic process
automation (RPA), predictive models/analytics, AI-powered/driven tools.

Tier 2 -- Automation or robotics (without explicit AI mention):
Includes:  automation, automated, automating, robots, robotics, cobots.
Excludes:  generic digitization, ERP/CRM software, cloud, IoT, 3D
printing, data analytics unless explicitly paired with
automation/robotics language.

Tier 0 -- Neither:  general innovation, process improvement, new
products, digital tools, or technology that does not involve AI or
automation/ robotics.

Respond with valid JSON only, no extra text:  {"tier":  <0, 1, or 2>,
"reason":  "<one sentence, max 20 words>"}
```

Product innovation prompt (applied to fields h3x, h4x, h32x) is identical except for an added introductory sentence, "The responses describe new or improved products and services that firms have introduced," and Tier 1 is clarified to apply "whether the firm is USING AI to develop the product OR offering an AI-powered product/service to customers."

Language handling. WBES responses are recorded in the language of the original interview and are not machine-translated prior to classification; Claude processes the raw multilingual text directly. I do not separately verify classification accuracy across languages, and this is noted as a limitation in Section 3.2.

A.2 Tier 1 and Tier 2 Keyword Patterns

The following patterns define the two tiers of the keyword classifier described in Section 3.2. All patterns are applied case-insensitively using regular expressions. Tier 2 patterns are only applied to firms that did not match any Tier 1 pattern.

Tier 1: Explicit AI/ML (28 patterns)

<table>
<tr><th>Regex pattern (case-insensitive)</th><th>Matches</th></tr>
<tr><td><code>\bartificial intelligence\b</code></td><td>“artificial intelligence”</td></tr>
<tr><td><code>\bmachine learning\b</code></td><td>“machine learning”</td></tr>
<tr><td><code>\bdeep learning\b</code></td><td>“deep learning”</td></tr>
<tr><td><code>\bneural network</code></td><td>“neural network(s)”</td></tr>
<tr><td><code>\blarge language model</code></td><td>“large language model(s)”</td></tr>
<tr><td><code>\bgenerative ai\b</code></td><td>“generative AI”</td></tr>
<tr><td><code>\bchatgpt\b</code></td><td>“ChatGPT”</td></tr>
<tr><td><code>\bgpt[-\s]?\d</code></td><td>“GPT-4”, “GPT 3”, etc.</td></tr>
<tr><td><code>\bcopilot\b</code></td><td>“Copilot”</td></tr>
<tr><td><code>\bgemini\b</code></td><td>“Gemini”</td></tr>
<tr><td><code>\bperplexity\b</code></td><td>“Perplexity”</td></tr>
<tr><td><code>\bdatabricks\b</code></td><td>“Databricks”</td></tr>
<tr><td><code>\bsalesforce einstein\b</code></td><td>“Salesforce Einstein”</td></tr>
<tr><td><code>\bcomputer vision\b</code></td><td>“computer vision”</td></tr>
<tr><td><code>\bnatural language processing\b</code></td><td>“natural language processing”</td></tr>
<tr><td><code>\bnlp\b</code></td><td>“NLP”</td></tr>
<tr><td><code>\bsentiment analysis\b</code></td><td>“sentiment analysis”</td></tr>
<tr><td><code>\bimage recognition\b</code></td><td>“image recognition”</td></tr>
<tr><td><code>\bfacial recognition\b</code></td><td>“facial recognition”</td></tr>
<tr><td><code>\bvoice recognition\b</code></td><td>“voice recognition”</td></tr>
<tr><td><code>\bspeech recognition\b</code></td><td>“speech recognition”</td></tr>
<tr><td><code>\banomal[yi] detection\b</code></td><td>“anomaly/anomali detection”</td></tr>
<tr><td><code>\brobotic process automation\b</code></td><td>“robotic process automation”</td></tr>
<tr><td><code>\brpa\b</code></td><td>“RPA”</td></tr>
<tr><td><code>\bpredictive (model|analytic|algorithm)\b</code></td><td>“predictive model/analytic/algorithm”</td></tr>
<tr><td><code>\bml[-\s]?(model|algorithm|pipeline)\b</code></td><td>“ML model”, “ML algorithm”, etc.</td></tr>
<tr><td><code>\bai[-\s]?(powered|driven|...)\b</code></td><td>“AI-powered”, “AI-enabled”, etc.</td></tr>
<tr><td><code>\bai\b</code></td><td>standalone “AI”</td></tr>
</table>

Tier 2: Automation and Robotics (5 patterns)

<table>
<tr><th>Regex pattern (case-insensitive)</th><th>Matches</th></tr>
<tr><td><code>\bautomation\b</code></td><td>“automation”</td></tr>
<tr><td><code>\bautomated\b</code></td><td>“automated”</td></tr>
<tr><td><code>\bautomat(ing|ion|ed)\b</code></td><td>“automating”, “automation”, “automated”</td></tr>
<tr><td><code>\brobot(ic|s|ics)?\b</code></td><td>“robot”, “robotic”, “robotics”</td></tr>
<tr><td><code>\bcobots?\b</code></td><td>“cobot”, “cobots”</td></tr>
</table>

Tier 2 patterns are applied only to firms that matched no Tier 1 pattern.

A.2 Country Coverage

The working sample covers 144 countries. Of these, five lack WDI data on internet penetration or GDP (Central African Republic, Ethiopia, Kosovo, Somalia, and South Sudan), leaving 139 countries in the preferred specification. Specifications without internet penetration as a control include up to 144 countries.

A.3 Survey Year Distribution

Survey year	Firms	Countries
2022	14,668	12
2023	24,657	48
2024	23,257	65
2025	26,804	75
Total	89,380	144

Note: Country counts reflect unique countries surveyed in each wave. Countries may appear in multiple waves; the total of 144 unique countries exceeds the sum of per-year country counts because some countries are surveyed in more than one wave.

A.4 Full Country Sample

Table 18 lists all 144 countries in the WBES 2022–2025 sample. For countries surveyed in multiple waves, firm counts are reported separately for each wave in the same order as the years listed. The regression sample includes up to 139 countries depending on controls; see Section 3 for details.

Table 18: All Countries in the WBES 2022–2025 Sample

Country	Year(s)	Firms	Country	Year(s)	Firms
Afghanistan	2025	426	Lao PDR	2024, 2025	303; 60
Albania	2025	365	Latvia	2024, 2025	155; 205
Angola	2024, 2025	280; 150	Lesotho	2023	150
Antigua and Barbuda	2025	81	Liberia	2025	158
Armenia	2024, 2025	360; 15	Madagascar	2022	402
Austria	2025	601	Malawi	2025	153
Azerbaijan	2024, 2025	249; 119	Malaysia	2024, 2025	345; 634
Bahrain	2024, 2025	116; 34	Maldives	2025	154
Bangladesh	2022	998	Mali	2024, 2025	374; 18
Barbados	2023	153	Malta	2024, 2025	170; 71
Belgium	2024, 2025	233; 87	Mauritania	2025	163
Benin	2024, 2025	354; 10	Mauritius	2023, 2024	338; 15
Bhutan	2024	155	Mexico	2023	1,322
Bolivia	2025	610	Moldova	2024	152
Bosnia and Herzegovina	2023	351	Mongolia	2025	601
Botswana	2023	622	Montenegro	2023	151
Brunei Darussalam	2025	150	Morocco	2023	598
Bulgaria	2023	710	Mozambique	2025	621
Burkina Faso	2024, 2025	333; 47	Namibia	2024, 2025	195; 112
Burundi	2025	161	Nepal	2023	582
Cabo Verde	2024	163	New Zealand	2023	358
Cambodia	2023	519	Niger	2025	156
Cameroon	2024, 2025	596; 19	Nigeria	2025	1,043
Canada	2024, 2025	273; 742	North Macedonia	2023	354
Central African Republic	2023	151	Pakistan	2022	1,300
Chad	2023	164	Papua New Guinea	2024, 2025	60; 85
China	2024, 2025	1,504; 685	Paraguay	2023, 2024	289; 89
Colombia	2023, 2024	804; 115	Peru	2022, 2023	396; 591
Comoros	2025	153	Philippines	2023, 2024	913; 89
Congo	2024	365	Poland	2024, 2025	242; 1,483
Costa Rica	2023	357	Portugal	2023, 2024	921; 86
Croatia	2023, 2024	411; 63	Qatar	2025	480
Cyprus	2024	246	Romania	2023, 2024	933; 14
Czechia	2024, 2025	141; 127	Rwanda	2023, 2024	291; 67

Continued on next page

Table 18 continued

Country	Year(s)	Firms	Country	Year(s)	Firms
CÃ´te d'Ivoire	2023	649	Samoa	2023	157
Dem. Rep. Congo	2024	1,025	Sao Tome and Principe	2025	164
Denmark	2025	822	Saudi Arabia	2022, 2023, 2025	793; 780; 1,002
Djibouti	2025	153	Senegal	2024, 2025	479; 124
Ecuador	2024, 2025	315; 30	Serbia	2024	502
Egypt	2025	1,005	Seychelles	2023	103
El Salvador	2023	729	Sierra Leone	2022, 2023	91; 118
Equatorial Guinea	2024	173	Singapore	2023	623
Estonia	2023	351	Slovak Republic	2023, 2024	227; 237
Eswatini	2024	156	Slovenia	2024	402
Ethiopia	2025	1,011	Solomon Islands	2025	150
Fiji	2025	151	Somalia	2025	241
Gabon	2025	151	South Sudan	2024	156
Gambia	2023	162	Spain	2022, 2024, 2025	33; 1,143; 288
Georgia	2023	592	Sri Lanka	2025	607
Germany	2022, 2025	268; 1,943	St. Lucia	2025	150
Ghana	2023, 2024	698; 15	Suriname	2025	152
Greece	2023	598	Sweden	2024	598
Guinea	2025	160	Switzerland	2025	492
Guinea-Bissau	2025	153	Türkiye	2024, 2025	1,080; 334
Hong Kong SAR China	2023	598	Tajikistan	2024	364
Hungary	2023, 2024	758; 73	Tanzania	2023	600
Iceland	2024	359	Thailand	2025	597
India	2022	9,111	Timor-Leste	2022	213
Indonesia	2022, 2023	54; 2,901	Togo	2023	148
Iraq	2022	1,009	Tonga	2024	150
Ireland	2024	609	Trinidad and Tobago	2024, 2025	27; 87
Israel	2024	388	Tunisia	2024, 2025	237; 408
Italy	2024	1,216	Turkmenistan	2024, 2025	53; 258
Jamaica	2023, 2024, 2025	54; 107; 54	Uganda	2025	605
Jordan	2024	592	United Kingdom	2024, 2025	326; 677
Kazakhstan	2024	1,013	United States	2024, 2025	727; 1,970

Continued on next page

Table 18 continued

Country	Year(s)	Firms	Country	Year(s)	Firms
Kenya	2025	1,024	Uruguay	2024, 2025	299; 61
Kiribati	2025	150	Uzbekistan	2024	1,008
Korea, Rep. of	2024, 2025	1,450; 68	Vanuatu	2023, 2024	34; 77
Kosovo	2025	269	Viet Nam	2023	1,028
Kuwait	2025	150	West Bank and Gaza	2023	360
Kyrgyz Republic	2023	354	Zimbabwe	2025	359
Total (144 countries)					89,380

For multi-wave countries, firm counts correspond to the listed years in order.

Table 19: Top 30 Countries by AI/Automation Adoption Rate

Rank	Country	Firms	Adopters	Rate (%)	OAD	TFR
1	Denmark	822	127	15.5	32.9	1.47
2	Switzerland	492	52	10.6	30.8	1.29
3	Belgium	320	30	9.4	32.4	1.44
4	Czechia	268	25	9.3	32.7	1.36
5	Moldova‡	152	13	8.6	25.3	1.73
6	Slovenia	402	33	8.2	34.3	1.52
7	China	2,189	172	7.9	21.2	1.01
8	Sweden	598	44	7.4	33.3	1.43
9	Mongolia	601	41	6.8	8.2	2.50
10	New Zealand	358	23	6.4	25.9	1.57
11	Spain	1,464	93	6.4	32.0	1.10
12	Peru	987	60	6.1	13.4	1.99
13	Ecuador	345	20	5.8	12.4	1.81
14	Kyrgyz Rep.	354	19	5.4	8.8	2.70
15	Ireland	609	31	5.1	24.2	1.47
16	Somalia‡	241	12	5.0	5.1	6.01
17	Germany	2,211	110	5.0	36.7	1.37
18	Kosovo	269	12	4.5	14.6	1.54
19	Zimbabwe	359	16	4.5	6.5	3.67
20	Iceland	359	16	4.5	23.5	1.56
21	Bahrain‡	150	6	4.0	5.0	1.80
22	Albania	365	14	3.8	25.5	1.34
23	Malta‡	241	9	3.7	30.3	1.01
24	Slovak Rep.	464	17	3.7	27.8	1.47
25	Costa Rica	357	13	3.6	17.0	1.33
26	Austria	601	21	3.5	31.6	1.31
27	Sierra Leone‡	209	7	3.3	5.5	3.83
28	Solomon Islands‡	150	5	3.3	6.1	3.51
29	Uruguay	360	12	3.3	24.4	1.40
30	Fiji‡	151	5	3.3	9.8	2.27

Notes: OAD = old-age dependency ratio (elderly per 100 working-age). TFR = total fertility rate. Sample restricted to countries with at least 30 surveyed firms. Adoption rates use the LLM broad measure (`ai_process_llm_broad`). ‡Fewer than 250 surveyed firms; adoption rate may reflect sampling noise.

PUBLICATIONS

Aging Economies and AI Adoption: Firm-Level Evidence from the World Bank Enterprise Surveys

Working Paper No. WP/2026/176