How Is AI Adopted Across Countries and Occupations?

Early Evidence from One Million Usage Conversations in Over 100 Countries

Rachel Yuting Fan

May 2026

Abstract

How is AI adopted across countries and occupations? This paper provides the first large-scale early evidence, using a sample of one million Claude AI conversations from more than 100 countries during one week in February 2026. Three findings emerge. First, per-capita AI usage varies more than 200-fold across countries; income level, economic structure, scientific output, and AI-policy readiness jointly determine the level of AI penetration. Second, occupational breadth follows volume: countries with more total AI conversations use AI across more occupations. Third, countries naturally follow a global diffusion path, entering occupations that are structurally proximate to their existing AI-occupation portfolio and for which AI tools are widely available globally. Compute capacity, policy vision, and targeted sector-specific interventions determine which occupations adopt AI next.

Keywords: generative AI; AI adoption; technology diffusion; product space; AI policy; cross-country

1 Introduction

Generative AI has been adopted faster than any previous digital technology, yet direct measures of how it is adopted across countries and occupations remain scarce. The dominant approach in the literature relies on *exposure indices*: expert assessments of which occupational tasks AI *could* perform, constructed from US labor-market data and applied uniformly across economies (Felten et al., 2021; Eloundou et al., 2024). These indices measure potential exposure, not actual use; they cannot distinguish a country where AI is widely adopted from one where the same occupations exist but no one uses AI; and they cannot reveal what users actually ask AI to do. Firm-level surveys cover only a single country and miss consumer usage (Calvino et al., 2022; Calvino and Fontanelli, 2023). Recent studies have begun to close this gap: Chugunova et al. (2026) survey 6,000 researchers at two German institutions to ask who uses AI in research and for what; Calvino and Fontanelli (2026) characterize the difference between French firms that buy AI and those that develop it. But no study has yet observed what people *actually ask AI to do* at a global scale, across all sectors and countries, using revealed behavioral rather than self-reported data.

This paper provides the first such evidence. Using a random sample of one million Claude AI conversations from more than 100 countries during one week in February 2026, observed usage data rather than self-reports or expert projections, I document how AI adoption varies across countries and how it diffuses across occupations. The data come from the Anthropic Economic Index (AEI), a publicly available dataset in which each conversation is mapped to the closest O*NET task statement, which in turn maps to a six-digit O*NET-SOC occupation and a two-digit SOC major group (Anthropic, 2026; Appel et al., 2026). The resulting cross-country variation among this one million sample is large: per-capita AI usage ranges 244-fold, from 3.7 conversations per million people in Niger to 911 in Singapore.[1] The mix of occupations varies with economic structure. Computer & Mathematical occupations (SOC 15) account for 66% of classified conversations in Egypt, 48% in India, 39% in Germany, and 27% in the United States. In one country, AI usage spans law, health, education, finance, and creative production; in another, it concentrates almost entirely in software debugging and data analysis.

To visualize the different paths of AI diffusion through occupations within each country, I adapt the product-space methodology of Hidalgo et al. (2007) to construct an *AI-diffusion space*: a network of 22 SOC two-digit occupations linked by revealed co-adoption patterns. The AI-diffusion space shows that lower-income countries concentrate in a software-and-STEM core, while higher-income countries activate a periphery of legal, health, creative, and financial occupations.

The question is what accounts for these differences, and how AI adoption spreads from early concentration in a few occupations toward broader occupational coverage. This question matters because AI usage today starts almost everywhere in Computer & Mathematical occupations, and a country that never moves beyond this software core captures productivity gains in only a narrow slice of its economy. The general-purpose-technology literature argues that the economic returns to a new technology depend less on how much it is used than on how many domains it permeates (Bresnahan and Trajtenberg, 1995): the steam

[1] Conversations per million population is computed as the country's conversation count in the AEI sample divided by total population (World Bank, 2024) and multiplied by 10^6. The 244-fold ratio is computed from unrounded values (910.9 / 3.7 = 243.8); the per-million figures shown are rounded. Section 2.2 details the intensity measure.

engine transformed manufacturing only after it spread from mining to textiles, transport, and agriculture; electricity required decades of diffusion across industries before aggregate productivity gains materialized (David, 1990). The same logic applies to AI. A country that uses AI only for software debugging puts a small share of its workforce in contact with the technology; a country that uses AI across law, health, education, and creative production multiplies the occupations where AI-driven productivity gains can emerge. Understanding what drives this spreading is therefore central to whether AI will deliver broad-based economic benefits or remain confined to a technology-sector enclave. I ask three questions: what determines how much a country uses AI (*intensity*), how broadly that usage spreads across occupations (*breadth*), and which specific occupations adopt AI as usage grows (*direction*).

On intensity, countries with higher income, stronger scientific output, a more services-oriented economy, and more developed AI-policy readiness tend to use AI more per capita. Singapore, Luxembourg, Switzerland, Israel, and the United States lead in per-capita usage, while the United States (27.5%), India (7.6%), and the United Kingdom (4.4%) lead by total volume. Sub-Saharan Africa and Central Asia show the lightest usage. Scale leads to breadth: the total volume of AI conversations in a country is the dominant predictor of how broadly usage spreads across occupations; countries that use AI more heavily in total will see AI activity across a wider set of occupations. A natural question that follows is which occupations adopt AI as usage grows. Using three AEI waves spanning six months, I show that countries naturally follow a global diffusion path, entering occupations that are structurally proximate to their existing AI-occupation portfolio, following the same principle of relatedness that governs export diversification and patent entry. Widely adopted occupations are easier to enter, and the two forces reinforce each other. Beyond structural proximity and ubiquity, compute capacity and policy vision determine which occupations adopt AI next, while targeted sector-specific interventions can further boost AI adoption into occupations that the structural model alone would not predict.

This paper contributes primarily to the measurement of AI adoption. Chugunova et al. (2026) measure who uses AI within German research organizations (survey of 6,000 researchers, single country); Calvino and Fontanelli (2026) characterize AI-adopting firms in France (firm-level survey, single country). I complement these studies by measuring AI usage across all sectors and over 100 countries, using over 800 thousand revealed-preference conversations, two orders of magnitude larger than existing survey samples, and covering the full global income distribution rather than a single advanced economy. I also connect to the cross-country technology-diffusion and catch-up literature (Comin and Hobijn, 2010; Fagerberg, 1994; Lee, 2024) by providing the first direct measurement of how a general-purpose technology diffuses across countries, and to the product-space literature (Hidalgo et al., 2007; Juhász et al., 2026) by showing that AI adoption across occupations follows the same relatedness structure documented for trade, patents, and software. I proceed as follows. Section 2 describes the data, sample construction, and measurement. Section 3 presents the AI-diffusion space and country-level snapshots. Section 4 reports the intensity regression and the breadth–volume relationship. Section 4.3 introduces relatedness density and tests the principle of relatedness predictively across three AEI waves. Section 5 discusses implications, and Section 6 concludes.

2 Data

2.1 The Anthropic Economic Index and occupation classification

The primary data source is the Anthropic Economic Index (AEI), a publicly available dataset released by Anthropic on HuggingFace (Anthropic, 2026). The AEI periodically samples approximately one million conversations from Claude, Anthropic's AI assistant, and classifies each conversation by the type of work it most closely resembles. The wave used in the cross-sectional analysis covers the week of February 5–12, 2026. Anthropic summarizes each conversation into a short abstract stripped of personally identifying content, then tags it with the closest task from the O*NET-SOC taxonomy, a standardized database of occupational tasks maintained by the U.S. Department of Labor. Each O*NET task maps to a six-digit occupation code, which in turn rolls up to a two-digit SOC major group. The February 2026 sample covers 1,576 distinct O*NET tasks across 537 six-digit occupations and 22 of the 23 two-digit SOC major groups.[2] I use the 22 SOC two-digit groups as the primary occupation unit in this paper.

Geographic location is derived from each conversation's IP address using ISO-3166 country codes, with VPN, anycast, and hosting-service traffic excluded. Anthropic applies a privacy suppression floor: any country-occupation cell with fewer than 15 conversations or five unique users is suppressed. After this suppression, over 800 thousand of the one million sampled conversations are attributable to a specific country and enter the cross-country analysis (Appendix Figure A1 details the full classification pipeline). Table 1 illustrates the classification with worked examples: two SOC-2 occupation groups and two of the O*NET tasks that fall inside each.

Table 1: Occupation classification: SOC two-digit groups and representative O*NET tasks.

SOC-2 category	Representative O*NET task description
Computer & Mathematical (SOC 15)	Write new programs or modify existing programs to meet customer requirements, using current programming languages and technologies.
	Design, build, or maintain web sites, using authoring or scripting languages, content creation tools, management tools, and digital media.
Educational Instruction & Library (SOC 25)	Assist students who need extra help with their coursework outside of class.
	Review class material with students by discussing text, working solutions to problems, or reviewing worksheets or other assignments.

Notes: Two of the 22 SOC two-digit occupation groups in the sample and two representative O*NET-SOC tasks that map into each group. Task descriptions are reproduced verbatim from the O*NET-SOC taxonomy.

2.2 Measuring intensity and breadth

Intensity is measured by Claude conversations per million people within this one million conversations sample, the natural per-capita measure that makes usage comparable across economies of different sizes.

Breadth, how widely AI usage is distributed across occupations, is measured by the Herfindahl–Hirschman

[2] SOC-55 (Military Specific) is excluded by Anthropic for operational-security reasons.

Index (HHI) over the 22 SOC two-digit occupation groups.[3] An HHI close to 10,000 indicates that a country's AI usage is concentrated in one or two occupations; lower values indicate broader usage.

Table 2 reports summary statistics for the country-level covariates used in the regressions.

Table 2: Summary Statistics

	N	Mean	SD	Min	Median	Max
Panel A: AI usage measures (country level)						
Conversations per million	189	180	224	0	87	1268
log(HHI over SOC-2 occupations)	117	8.42	0.55	7.16	8.43	9.21
Panel B: Intensity regression (country level)						
log(GDP per capita, PPP)	171	9.86	1.13	7.09	9.97	11.96
log(Population)	189	15.21	2.42	9.17	15.61	21.10
Old-age dependency ratio	189	16.5	11.6	2.0	12.3	71.8
Services VA (% GDP)	159	56.9	11.5	14.5	57.5	87.2
Internet users per 100	168	73.6	23.9	8.6	82.3	100.0
log(Sci. articles per million pop.)	177	4.78	1.98	0.17	4.93	7.90
AIPI: Regulation	156	0.128	0.047	0.022	0.127	0.230
Panel C: Entry regression (country-occupation pair level, $N = 1{,}653$)						
Entry (0/1)	1653	0.080	0.272	0.00	0.00	1.00
Relatedness density (ω)	1653	0.147	0.168	0.006	0.083	0.953
Ubiquity (# countries in R5)	1653	25.87	20.62	2	22	94
Compute capacity (0–1)	1653	0.103	0.094	0.00	0.074	0.57
Policy vision (0–1)	1653	0.530	0.262	0.00	0.50	1.00

Sources: AI usage measures from the Anthropic Economic Index (AEI), R5 release (Anthropic, 2026). GDP per capita (PPP), population, services value added, and internet users from the World Development Indicators (World Bank, 2024). Old-age dependency ratio from UN Population Division. Scientific articles from SCImago/Scopus. AIPI Regulation sub-index from Cazzaniga et al. (2024). Relatedness density (ω) and ubiquity computed from the R3 and R5 AEI waves following Hidalgo et al. (2007). Compute capacity and policy vision from the Oxford Government AI Readiness Index 2025 (Oxford Insights, 2025). Panel C reports statistics at the country-occupation pair level (occupations not adopted in R3); country-level variables are repeated across pairs.

3 Motivating facts

3.1 Cross-country variation in AI usage intensity and concentration

Per-capita AI usage varies enormously across countries. Among the one million conversations globally during the week of February 5–12, 2026, the five leading countries by per-capita usage are Singapore (911 conversations per million people), Luxembourg (768), Switzerland (698), Israel (689), and the United States (654). The lowest-usage country with observed conversations, Niger, records 3.7 per million, a 244-fold gap. The GDP-per-capita gap between Singapore and Niger is approximately 74-fold (PPP); the AI-usage gap is more than three times steeper. The interquartile range runs from 45 (Indonesia) to 283 (Italy) conversations per million, with a median of 112 (Malaysia). Higher per-capita usage concentrates in North America,

[3] HHI is computed after excluding the `not_classified` residual and renormalizing the remaining shares to 100%, so concentration is measured over substantively classified occupations only.

Western and Northern Europe, four East Asian economies (Japan, South Korea, Singapore, Taiwan), and a small set of high-income outliers (Israel, Ireland, the Nordics, and the UAE). Sub-Saharan Africa, Central Asia, and large parts of the Middle East show the lightest usage. The five largest users by total volume are the United States (27.5% of all country-level conversations), India (7.6%), the United Kingdom (4.4%), France (4.0%), and Germany (3.7%).

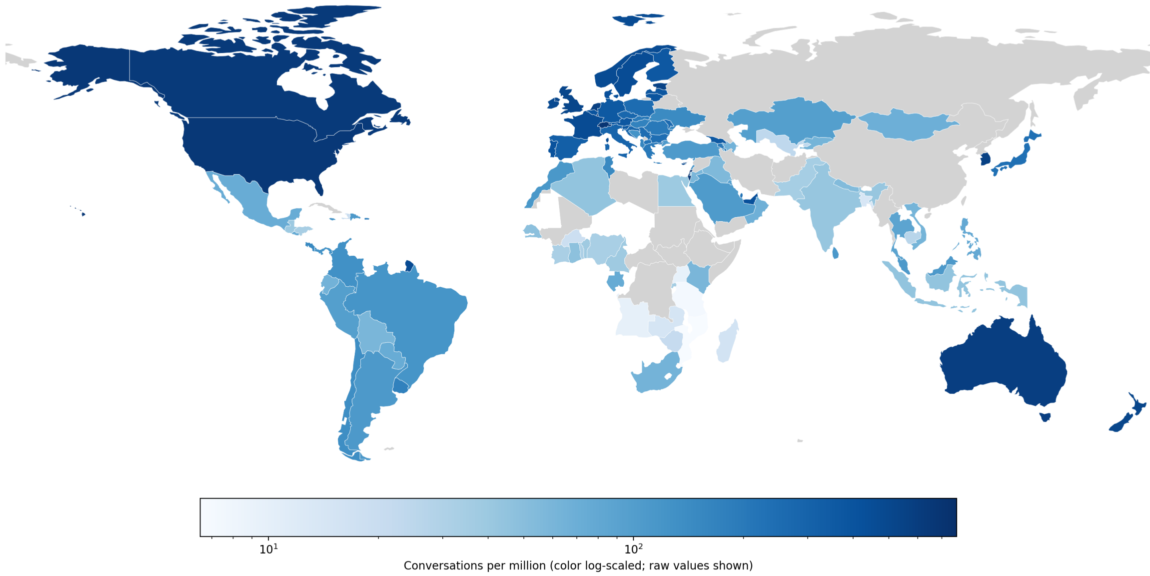

Figure 1: Where is AI adopted? Among one million global sample, conversations per million people during the week of February 5–12, 2026. Darker shading indicates higher per-capita usage. Gray indicates no data.

The composition of AI usage also varies. Countries that use AI more intensively per capita also tend to use it more broadly across occupations.

3.2 Country-by-country AI-diffusion-space snapshots

To visualize how AI adoption clusters across occupations, I adapt the *product space* methodology of Hidalgo et al. (2007) to construct an *AI-diffusion space*: a network of the 22 SOC two-digit occupation groups where two occupations sit closer together if countries that use AI intensively in one also tend to use AI intensively in the other. The proximity between two occupations is defined as the minimum of the two conditional probabilities that a country with revealed comparative advantage (RCA ≥ 1) in one also has RCA ≥ 1 in the other, following the canonical Hidalgo–Hausmann construction.[4]

Figure 2 presents six country snapshots. The United States activates 18 of 22 occupations, spreading well beyond software into healthcare, construction, and food service. Germany activates 8 occupations concentrated in engineering, protective service, and administrative support. Brazil activates only 4, dominated by its exceptionally strong legal node, reflecting AI-assisted document review in a judicial system

[4]RCA in an occupation is the ratio of a country's share of AI usage in that occupation to the global share. The graph backbone retains the maximum spanning tree plus all edges with proximity $\phi \geq 0.55$. Layout is computed once using the Kamada–Kawai algorithm (Kamada and Kawai, 1989) and reused across all country panels. Construction details appear in the note to Figure 2.

with 80 million pending cases (CNJ, 2025). India activates just 2 occupations, both in the software core. Vietnam also activates 2 but in a different pair: Computer & Math and Arts/Media. Tanzania activates 2 in Computer & Math and Education.

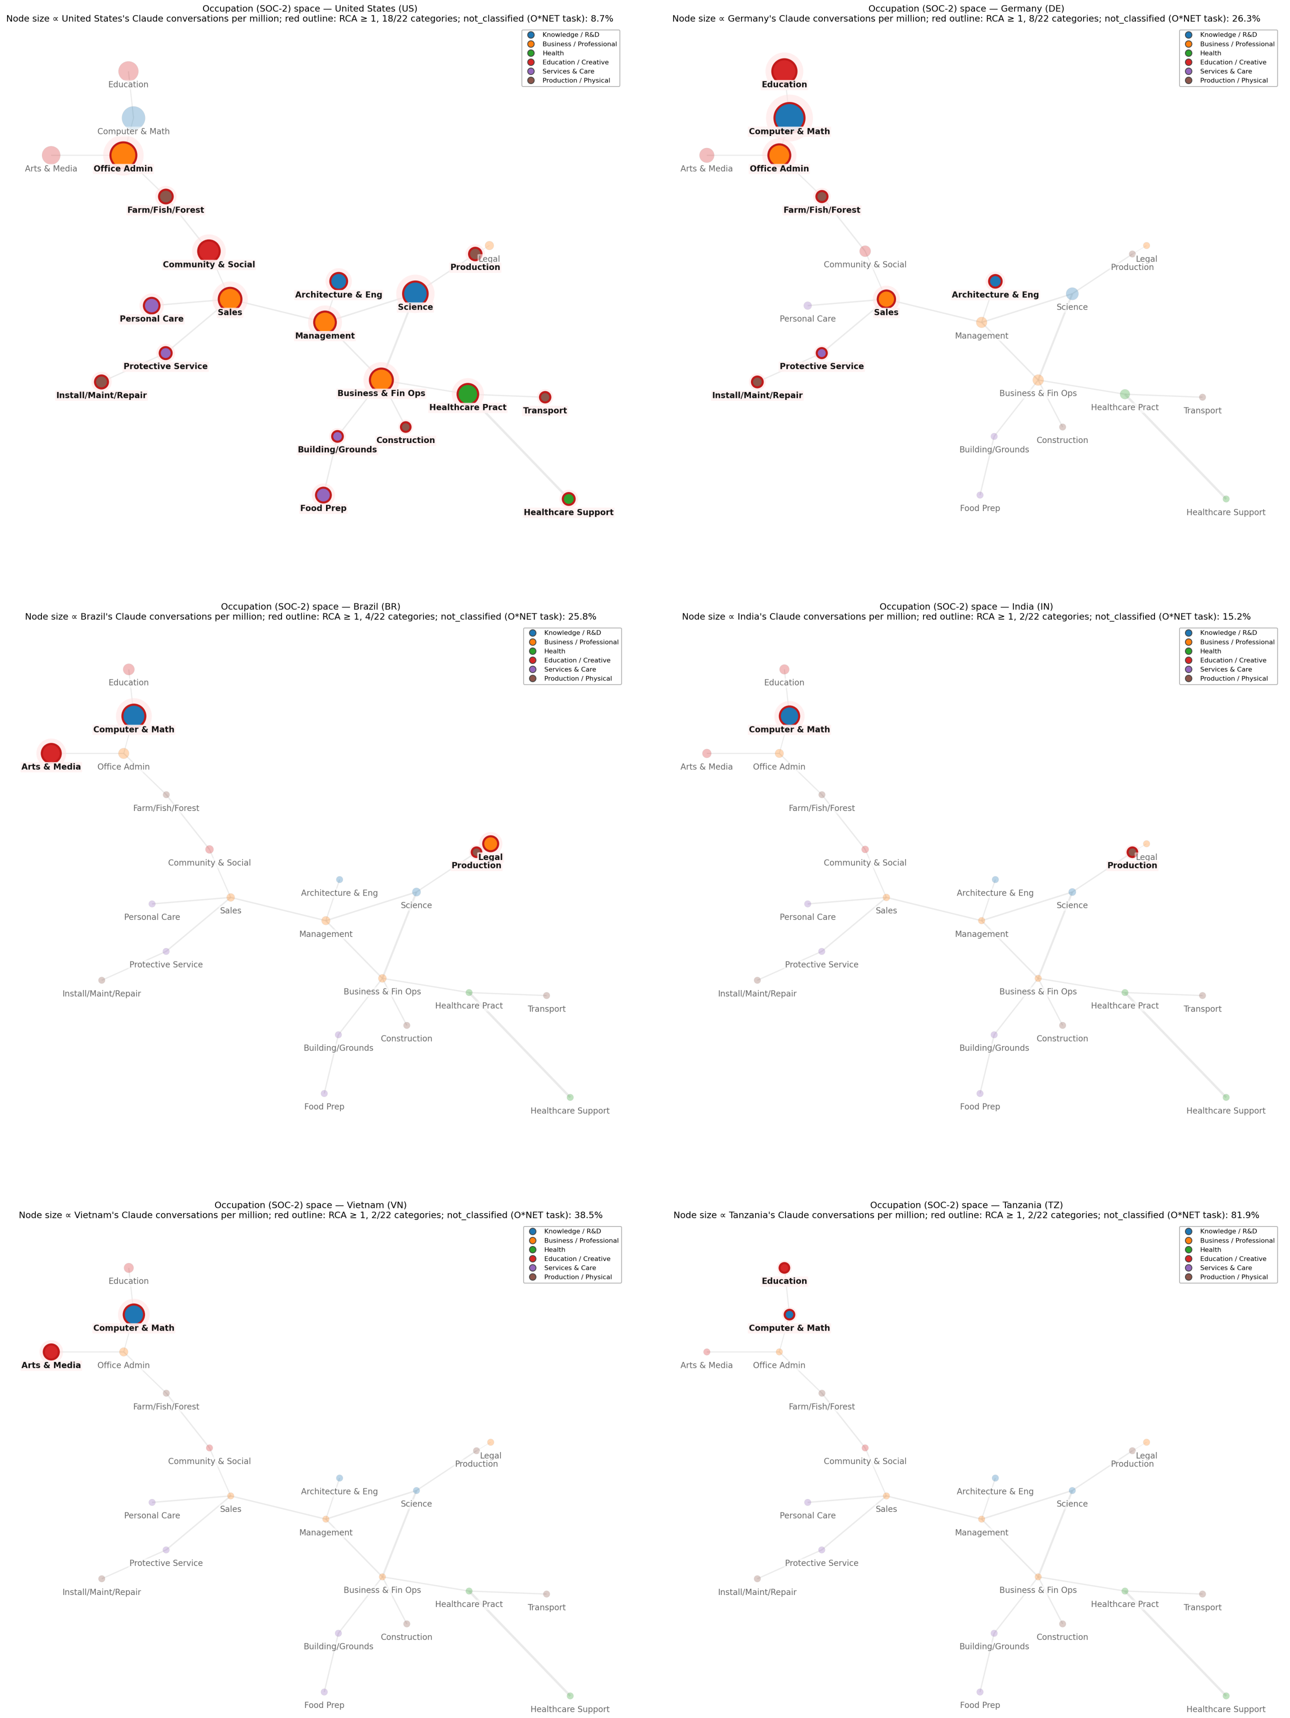

Figure 2: Country snapshots of the SOC-2 AI-diffusion space. Same 22-node layout in each panel. Node size $\propto$ country's conversations per million in each occupation; red outline = RCA $\geq$ 1. Edges connect occupations with proximity $\phi \geq 0.55$. Layout: Kamada–Kawai (Kamada and Kawai, 1989), computed once and reused. Row 1: US, Germany. Row 2: Brazil, India. Row 3: Vietnam, Tanzania.

The six panels reveal two regularities. First, the number of occupations that adopt AI increases with income: the United States activates 18 of 22 occupations, Germany 8, Brazil 4, and India, Vietnam, and Tanzania each 2. Second, *which* occupations adopt AI reflects the country's existing economic structure, not a universal sequence. Germany's AI adoption concentrates in engineering and administrative support, mirroring its manufacturing economy; Brazil's concentrates in Legal, reflecting a judicial system with 80 million pending cases; India's concentrates in the software core, reflecting its IT-services export sector. AI does not diffuse uniformly from a single origin; it follows the economy's existing occupational contours.

4 What determines AI adoption?

4.1 Intensity: per-million population AI usage

Which country characteristics are associated with higher per-capita AI usage? The dependent variable is the log of Claude conversations per million people. I estimate:

$$\log(\text{ConvPerMCap}_c) = \alpha + X_c'\boldsymbol{\beta} + \varepsilon_c, \tag{1}$$

where X_c is the vector of country-level regressors, with Huber–White (HC1) robust standard errors. Table 3 reports eight progressive models.

Table 3: Progressive log-OLS for per-capita AI usage

	(M1)	(M2)	(M3)	(M4)	(M5)	(M6)	(M7)	(M8)
log(GDP per capita, PPP)	1.064***	1.033***	0.829***	0.676***	0.441***	0.294***	0.230**	0.232**
	(0.047)	(0.049)	(0.063)	(0.057)	(0.085)	(0.096)	(0.098)	(0.098)
log(Population)		-0.078***	-0.083***	-0.062***	-0.065***	-0.089***	-0.099***	-0.099***
		(0.030)	(0.028)	(0.022)	(0.021)	(0.022)	(0.021)	(0.021)
Old-age dependency ratio			0.027***	0.018***	0.019***	0.012**	0.011**	0.012**
			(0.005)	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)
Services VA (% GDP)				0.035***	0.036***	0.033***	0.029***	0.029***
				(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Internet users (%)					0.012***	0.009***	0.011***	0.012***
					(0.003)	(0.003)	(0.003)	(0.003)
log(Sci. articles/million pop.)						0.154***	0.106***	0.099**
						(0.038)	(0.039)	(0.039)
AIPI: Regulation							4.407***	4.303***
							(1.301)	(1.331)
English-official (0/1)								0.062
								(0.094)
Observations	129	129	129	123	122	122	122	122
Adjusted R^2	0.817	0.823	0.856	0.899	0.909	0.918	0.924	0.924

Notes: Dep. var.: log(Claude conversations per million people). Robust HC1 standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Income is the dominant predictor. Model (M1) shows that a 10% higher GDP per capita is associated with approximately 11% more AI conversations per million people; income alone explains over four-fifths of the cross-country variation. Richer countries have more knowledge workers, better digital infrastructure, and greater exposure to the professional tasks that current AI tools assist.

Models (M2)–(M3) add population and the old-age dependency ratio. Larger countries show somewhat lower per-capita usage. Aging economies use AI more intensively, consistent with Acemoglu and Restrepo (2022)'s finding that countries with older workforces adopt labor-substituting technologies faster. For generative AI, the mechanism may run through both supply and demand: aging societies have fewer young workers to perform routine cognitive tasks, and AI is especially capable of assisting the entry-level and junior work that these societies most need to fill. Japan and Italy, the OECD's two most aged economies, are among the heaviest per-capita AI users in their income brackets.

Model (M4) introduces services value added. Service-oriented economies use AI substantially more: a 10-percentage-point higher services share is associated with roughly 35% more per-capita AI usage. Countries whose GDP comes from finance, consulting, media, and professional services have the occupations for which AI is most directly a labor-augmenting tool; countries with large extractive or manufacturing sectors have a narrower set of AI-complementary roles. This helps explain why oil-exporting Gulf economies, despite high incomes, show lower per-capita AI usage than comparably wealthy service-oriented economies such as Singapore or Luxembourg.

Models (M5)–(M6) add internet penetration and research intensity (scientific articles per million population). Both are significant: connectivity lowers the access cost of cloud-based AI tools, and countries with higher research intensity both produce and absorb AI tools more readily.

Model (M7) adds the IMF AI Preparedness sub-index for Regulation (Cazzaniga et al., 2024), which captures data-protection laws, AI-governance institutions, and public-sector AI strategies. Countries with more developed regulatory frameworks use AI substantially more per capita. Singapore and the United Kingdom sit near the top of both the AIPI-Regulation and per-capita-usage rankings; countries with underdeveloped AI-regulatory institutions in Sub-Saharan Africa and Central Asia sit near the bottom of both. I take (M7) as the preferred model: all seven regressors are individually significant and jointly explain 92% of the cross-country variation.

Model (M8) adds English as a primary language, which is indistinguishable from zero. Despite Claude being an English-first product, English-speaking countries do not use AI more intensively once income, economic structure, and institutional readiness are controlled.

As a robustness check, I re-estimate the full progressive intensity model using PPML, which retains countries with zero usage as structural zeros (Appendix Table A2); all signs and significance patterns on the key regressors are confirmed.

4.2 Breadth: breadth across tasks

Occupation concentration and total conversation volume

How broadly does AI usage spread across occupations, and what drives that breadth? A natural measure is the Herfindahl–Hirschman Index (HHI) over the 22 SOC two-digit occupation groups in which each country's conversations are classified. A lower HHI indicates broader usage. Figure 3 plots log HHI against log total Claude conversations for the 115 countries with occupation-level data.

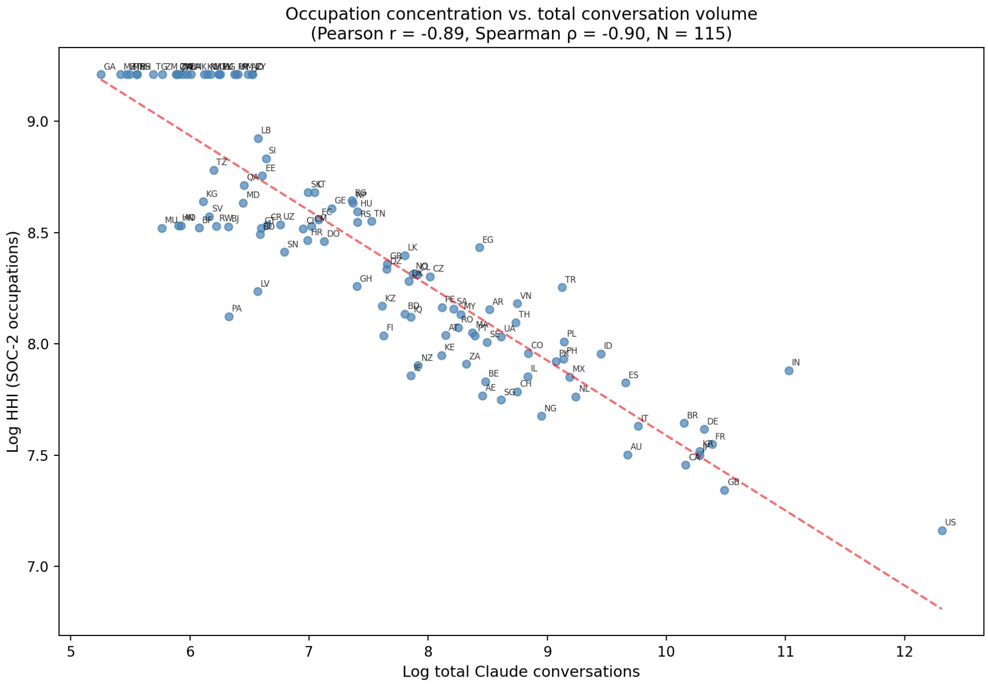

Figure 3: Occupation concentration vs. total conversation volume. Each point is one country ($N = 115$). Horizontal axis: log total Claude conversations in the AEI sampling week (Feb 5–12, 2026). Vertical axis: log HHI over 22 SOC two-digit occupations (higher = more concentrated). Dashed line: OLS fit. Spearman rank correlation $\rho = -0.90$ ($p < 0.001$), indicating that the ranking of countries by conversation volume almost perfectly mirrors the ranking by occupational breadth.

Countries with more total conversations have lower occupation concentration. The cluster of roughly 30 countries in the upper-left corner all have HHI of 10,000, meaning that all classified conversations fall into a single occupation (in most cases Computer & Mathematical); these are countries with fewer than approximately 1,000 total conversations.[5] The middle of the distribution traces a nearly linear gradient as conversation volume rises and AI usage spreads across more occupations. At the lower-right tail, the largest AI user bases, the United States, United Kingdom, Germany, France, and Japan, have usage spread across nearly all 22 occupations, with RCA ≥ 1 in many (as shown in the US and Germany panels of Figure 2). India is a notable outlier: despite its large conversation volume, its usage remains concentrated, with 48% in Computer & Mathematical occupations, reflecting an economy where AI adoption is channeled through a globally oriented IT-services export sector rather than spreading across the broader domestic occupation structure.

The near-perfect rank correlation ($\rho = -0.90$) means that total conversation volume, the product of per-capita intensity and population, is effectively sufficient to predict occupation-level breadth. This contrasts with the intensity regression (Table 3), where structural and institutional variables explain substantial additional variation beyond income. Full regression results of HHI on these variables are available upon

[5]With low volume of usage, most of the 22 SOC-2 occupation cells could fall below Anthropic's 15-conversation privacy floor and are suppressed, leaving only the dominant occupation visible.

request.

The implication is that occupation-level breadth is primarily a *volume* phenomenon: countries with higher total AI usage show broader occupational coverage, while countries with low usage concentrate in the first occupation AI touched upon, the software core in most countries. A natural question follows: *which* occupations adopt AI as usage grows, and what structural forces govern that path. The entry model in Section 4.3 addresses this question directly.

4.3 The principle of relatedness in AI-task adoption

The AI-diffusion space constructed in Section 3 generates a testable prediction. The principle of relatedness, documented for trade (Hidalgo et al., 2007), patents (Boschma et al., 2014), scientific output (Guevara et al., 2016), and regional employment (Neffke et al., 2011), holds that countries (or regions) are more likely to develop comparative advantage in activities that are structurally proximate to their existing capability portfolio. If generative-AI tasks follow the same logic, countries should be more likely to adopt AI in occupations that are "nearby" in the AI-diffusion space.

I test this using relatedness density, the standard measure in the economic-complexity literature (Hidalgo et al., 2007; Balland et al., 2019). For each country–task pair (c, i), relatedness density is defined as

$$\omega_{c,i} = \frac{\sum_{j \neq i} \phi(i,j) \cdot M_{c,j}}{\sum_{j \neq i} \phi(i,j)}, \tag{2}$$

where $\phi(i,j)$ is the proximity between occupations i and j (Section 3.2) and $M_{c,j} = \mathbf{1}[\text{count}_{c,j} > 0]$ indicates whether country c has any AI conversations classified in occupation j. I depart from the standard product-space literature, which defines M using revealed comparative advantage (RCA ≥ 1), because AI adoption is still in its early stages and RCA measures *specialization*, not *adoption*. As a country's AI usage broadens across occupations, the share of any single occupation mechanically falls, potentially pushing its RCA below 1 even though the country did not reduce its usage in that occupation.[6] The binary count > 0 definition captures actual adoption and better fits a technology that is still diffusing. A high $\omega_{c,i}$ means that country c's existing AI-occupation portfolio is concentrated in the neighborhood of occupation i.

Does relatedness density computed from an earlier wave predict which occupations a country adopts in a later wave? We test this out of sample, using ω computed from August 2025 to predict adoption by February 2026.

Of the 1,653 country-occupation pairs that had not adopted AI by August 2025, 132 (8.0%) are adopted by February 2026.[7] Which pairs make the transition? Table 4 reports six progressive logit models.

[6] The United States, for example, uses AI heavily in Computer & Math, Education, Arts & Media, and Legal, yet has RCA < 1 in all four because its usage is so evenly spread that no single occupation exceeds the global share. Using RCA ≥ 1 would code these as "not adopted."

[7] As a complement, Appendix Table A5 reports an exit logit on the 701 pairs adopted in August 2025, confirming that ω also protects against loss of adoption.

Table 4: Entry model: new AI adoption by occupation (logit).

	(1)	(2)	(3)	(4)	(5)	(6)
Log GDP per capita	0.510^{***}	0.216	0.487^{**}	0.496^{**}	0.062	-0.121
	(0.153)	(0.173)	(0.237)	(0.243)	(0.242)	(0.281)
Relatedness density (ω)		5.529^{***}	8.253^{***}	5.581^{***}	2.524^{**}	2.146
		(0.664)	(0.775)	(1.012)	(1.227)	(1.346)
Ubiquity (# countries in R5)			0.064^{***}	0.039^{***}	0.037^{***}	0.041^{***}
			(0.007)	(0.011)	(0.012)	(0.011)
$\omega\times$ ubiquity				0.140^{***}	0.201^{***}	0.202^{***}
				(0.038)	(0.042)	(0.043)
Compute capacity					7.492^{***}	7.230^{***}
					(1.822)	(1.983)
Policy vision						1.734^{*}
						(1.002)
Constant	-7.616^{***}	-5.874^{***}	-11.674^{***}	-11.210^{***}	-7.466^{***}	-6.684^{**}
	(1.588)	(1.734)	(2.604)	(2.686)	(2.670)	(2.812)
Pseudo R^2	0.032	0.202	0.389	0.419	0.445	0.452
N	1653	1653	1653	1653	1653	1653

Notes: Logit with standard errors clustered by country. The dependent variable is one if the country-occupation pair has any classified conversations in R5 (February 2026) and zero otherwise. The sample consists of 1,653 country-occupation pairs with no conversations in R3 (August 2025), i.e. occupations not yet adopted. Relatedness density (ω) is computed from the R3 adoption matrix following Hidalgo et al. (2007). Ubiquity is the number of countries with positive AI usage in that occupation in R5, capturing the global breadth of AI tool availability; results are robust to using pre-determined R3 ubiquity (Appendix Table A3). Maximum VIF is 3.87 (Appendix Table A4). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Income alone is a weak predictor of entry (model 1). What matters is structural proximity: countries adopt AI in occupations that are close to what they already use, just as countries diversify exports into products related to their existing capabilities (Hidalgo et al., 2007). India, whose AI portfolio was concentrated in Computer & Mathematical work, entered structurally adjacent Education between August 2025 and February 2026, but not structurally distant Healthcare or Legal. Brazil, with an unusually strong Legal node reflecting AI-assisted document review in its 80-million-case judicial system, entered Community & Social Services, a neighboring occupation, but not Food Preparation or Protective Services.

Ubiquity, the number of countries that use AI in a given occupation, measures how broadly AI capability has been deployed globally in each occupation. It is used to capture supply-side shocks such as the emergence of new AI functions that make certain occupations suddenly accessible worldwide. In model (3), occupations where AI tools, training data, and demonstrated use cases are already widely available are easier to enter, regardless of a country's own starting position. Arts & Media is the clearest example: the explosion of image-generation AI models between 2023 and 2025 created a supply-side shock that pushed Arts/Media ubiquity from 59 countries in August 2025 to 70 by February 2026, making it one of the easiest occupations to enter even for countries with no prior creative-industry AI use. By contrast, Construction remains at just two countries, as the ecosystem of construction-oriented AI tools barely exists.[8] The interaction in model (4) reveals that structural proximity and ubiquity reinforce each other: a country is especially likely to adopt AI in an occupation that is both close to its existing portfolio *and* widely adopted globally. Sales, for instance, is highly ubiquitous (adopted by 53 countries) and structurally adjacent to many countries' existing portfolios; countries with even moderate proximity to Sales entered it readily. Protective Services has comparable proximity for many countries but much lower ubiquity (only 20 countries), and far fewer countries entered it. Neither proximity nor ubiquity alone is sufficient; the combination is what predicts entry.

Models (5) and (6) add two country-level measures: compute capacity, used here to capture a country's cumulative public and private investment in AI infrastructure, including data centers, cloud availability, and processing power; and policy vision, capturing whether the government has articulated a strategic direction for AI.[9] Compute capacity is the single largest improvement in the table and absorbs GDP per capita entirely. For example, both UAE and Bahrain are high-income GCC economies, but the UAE has invested heavily in data-center infrastructure, scoring four times higher than Bahrain on compute capacity, and entered more new AI occupations than Bahrain during this period. Policy vision provides further improvement: Türkiye scores high on policy vision despite low on compute capacity, and entered four new occupations, suggesting that a clear government AI direction can partially compensate for limited infrastructure, though the effect is modest.

Table 5: Countries with the most excess AI occupation entries (August 2025 to February 2026).

Country	Unpredicted occupations	Plausible driver
Nigeria	Food Prep, Prot. Svc, Biz/FinOps, Pers. Care	Fintech platform growth (47% of Africa's fintech deals); AI-driven credit scoring and digital lending
Switzerland	Prot. Svc, Healthcare, Biz/FinOps	Diversified, high-skill economy; professionals across many fields adopt new tools quickly
Pakistan	Prot. Svc, Biz/FinOps, Pers. Care	IT-services exports growing 24% (2024–25); AI startups in fintech, edtech, healthtech
Belgium	Healthcare, Biz/FinOps, Pers. Care	ARIAC program (EUR 32M) targeting AI in healthcare and finance
Türkiye	Prot. Svc	TÜBİTAK 1711 AI Ecosystem Call (41 projects, 215M TL)

Notes: "Unpredicted occupations" are entries with predicted probability below 20% under Table 4, model (6).

Even after accounting for structural proximity, ubiquity, compute capacity, and policy vision, there is still room for targeted sector-specific interventions to boost AI adoption further. Table 5 lists the five countries with the largest new entries that the structural drivers were not able to predict. Nigeria is the largest overachiever despite scoring near the bottom on compute capacity, driven instead by its fintech

[8] Ubiquity is measured contemporaneously to capture the current state of global AI tool availability. Because it is a property of the *occupation* aggregated across all other countries, mechanical feedback from any single country's entry is negligible (at most 1 added to denominators of 12–94). Results are robust to using the pre-determined R3 measure (Appendix Table A3).

[9] Both measures are sub-indicators from the Oxford Government AI Readiness Index 2025 (Oxford Insights, 2025), scaled 0–1.

platform ecosystem creating AI demand in consumer-facing services. Switzerland reflects a diversified, high-skill economy where professionals across many fields adopt new tools quickly. Pakistan and Türkiye are driven by IT-export growth and government research programs respectively. Belgium entered Healthcare and Business/Financial Operations, precisely the sectors targeted by its ARIAC program. These cases suggest that targeted interventions, whether through platform ecosystems, sector-specific research funding, or government procurement mandates, can push AI adoption into occupations beyond what structural proximity, ubiquity, and infrastructure alone would predict.

5 Policy implications

AI usage today is overwhelmingly concentrated in Computer & Mathematical occupations. For AI to deliver broad-based productivity gains, it must spread beyond this software core into healthcare, education, legal services, and other occupations where the bulk of employment and value added resides. The results of this paper suggest that such spreading depends on a combination of structural conditions and deliberate policy choices.

The most direct path to broader occupational coverage is scale. The breadth–volume relationship shows that total AI usage is the dominant predictor of how many occupations AI reaches in a country: countries that use AI more in aggregate see it spread across a wider set of occupations almost mechanically. The policy priority, therefore, is raising the overall volume of AI adoption.

To achieve that volume, income alone is a weak predictor of per-capita AI usage; what matters is the combination of a service-oriented economic structure, research capacity, internet connectivity, and AI-regulatory readiness. Among them, AI-regulatory readiness (data-protection laws, dedicated AI institutions, public-sector AI strategies) is the most directly actionable: governments can build regulatory clarity and institutional capacity without large fiscal outlays, and countries with more developed AI-governance frameworks use AI substantially more per capita, even after controlling for income. Internet connectivity and research capacity require more sustained public investment but are still within the reach of active industrial policy. Economic structure, by contrast, is a slow-moving variable; governments cannot quickly shift their economies toward services, but they can ensure that the service sectors they already have are equipped to adopt AI.

Even after raising overall volume, the occupational direction of AI adoption is not automatic. Which occupations adopt AI in a given country depends on the interaction of structural proximity, global tool availability, and country-level enablers. Compute infrastructure has a direct payoff for broadening AI into new occupations, independent of income, but requires substantial capital investment and fiscal space that many developing countries lack. Policy vision, however, is more immediately actionable: Vietnam, for example, has adopted a comprehensive national AI strategy and governance framework despite modest income levels, and its positive deviation from the structural diffusion path suggests that such deliberate policy commitment can meaningfully redirect AI adoption even without large-scale compute investments. Where structural and institutional drivers are not sufficient, targeted sector-specific interventions (platform ecosystems, sector-specific research funding, government procurement) can push AI adoption into occupations that the global diffusion structure would otherwise bypass.

6 Conclusion

This paper provides the first large-scale early evidence on how AI is adopted across countries and occupations. Like earlier general-purpose technologies, AI adoption reflects the broader national innovation system: income level, economic structure, research intensity, and AI policy readiness jointly determine how much a country uses AI. Because occupational breadth follows volume, countries with low total usage remain concentrated in the software core, while higher-usage countries spread AI across a widening set of occupations. This breadth, however, does not expand randomly. Countries follow a global diffusion path shaped by structural proximity and the worldwide availability of AI tools, entering occupations that are close to their existing portfolio and for which proven tools already exist, as the rapid spread of image-generation models into Arts & Media illustrates. Where countries deviate from this structural path, compute infrastructure, policy vision, and targeted sector-specific interventions explain much of the residual, suggesting that the direction of AI adoption is partly a policy choice, constrained by but not reducible to the global structure of AI diffusion.

Several caveats apply. The data cover Claude.ai web conversations only, excluding API-based enterprise usage (which remains concentrated in Computer & Mathematical occupations) and other platforms (ChatGPT, Gemini, Copilot); the paper's conclusions rest on an assumption that occupational penetration patterns are broadly similar across platforms. Claude is unavailable in China and has lower consumer penetration than ChatGPT in several markets, so the occupation-composition findings likely overrepresent professional knowledge work. Each wave covers a single week, and usage patterns may shift with seasonal variation, product launches, or pricing changes; the three-wave temporal analysis partially addresses this, but longer panel coverage will allow stronger tests. The Anthropic classifier that maps conversations to O*NET occupations is not independently validated on a hand-labeled test set.

As more data become available from future releases, I look forward to extending the findings of country-occupation entries at a larger scale, incorporating multi-platform and API-based usage to broaden coverage beyond consumer web traffic. These steps would transform the cross-sectional portrait offered here into a richer longitudinal account of how AI diffuses across the global economy.

Acknowledgments

I am indebted to Ha Nguyen (IMF) for inspirational discussions and invaluable comments. I am also grateful to Ben Zou (Purdue University), Caglar Ozden (World Bank), Nelly Elmallakh (World Bank), and Shu Yu (World Bank) for helpful discussions. The author used Claude (Anthropic) for inspiration, editorial assistance, and coding aid. All remaining errors are my own.

Disclosure of interest

The author reports there are no competing interests to declare.

Data availability statement

The data that support the findings of this study are publicly available. The Anthropic Economic Index (AEI) data used in this paper are available on HuggingFace at `https://huggingface.co/datasets/Anthropic/EconomicIndex` (release 2026-03-24). Country-level macroeconomic covariates are sourced from the World Development Indicators and other standard cross-country databases as described in Section 2.

References

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

Anthropic. 2026. "Anthropic Economic Index Data, Release 2026-03-24." HuggingFace dataset. `https://huggingface.co/datasets/Anthropic/EconomicIndex/tree/main/release_2026_03_24`.

Appel, Ruth, Maxim Massenkoff, Peter McCrory, Miles McCain, Ryan Heller, Tyler Neylon and Alex Tamkin. 2026. "The Anthropic Economic Index Report: Economic Primitives." Anthropic Research Report, January 2026.

Balland, Pierre-Alexandre, Ron Boschma, Joan Crespo and David L. Rigby. 2019. "Smart Specialization Policy in the European Union: Relatedness, Knowledge Complexity and Regional Diversification." *Regional Studies*, 53(9), 1252–1268.

Boschma, Ron, Gaston Heimeriks and Pierre-Alexandre Balland. 2014. "Scientific Knowledge Dynamics and Relatedness in Biotech Cities." *Research Policy*, 43(1), 107–114.

Bresnahan, Timothy F. and Manuel Trajtenberg. 1995. "General Purpose Technologies: 'Engines of Growth'?" *Journal of Econometrics*, 65(1), 83–108.

Calvino, Flavio, Lea Samek, Mariagrazia Squicciarini and Cody Morris. 2022. "Identifying and Characterising AI Adopters: A Novel Approach Based on Big Data." OECD Science, Technology and Industry Working Papers 2022/06.

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

Calvino, Flavio and Luca Fontanelli. 2026. "AI Users Are Not All Alike: The Characteristics of French Firms Buying and Developing AI." *Research Policy*, 55(5), 105473.

Conselho Nacional de Justiça. 2025. "Justiça em Números 2025." Brasília: CNJ.

Cazzaniga, Mauro, Florence Jaumotte, Longji Li, Giovanni Melina, Augustus J. Panton, Carlo Pizzinelli, Emma J. Rockall and Marina M. Tavares. 2024. "Gen-AI: Artificial Intelligence and the Future of Work."

IMF Staff Discussion Note SDN/2024/001, International Monetary Fund, Washington, DC. (Introduces the IMF AI Preparedness Index, AIPI, including its Regulation sub-index.)

Chugunova, Marina, Dietmar Harhoff, Katharina Hölzle, Verena Kaschub, Sonal Malagimani, Ulrike Morgalla and Robert Rose. 2026. "Who Uses AI in Research, and for What? Large-Scale Survey Evidence from Germany." *Research Policy*, 55(2), 105381.

Comin, Diego and Bart Hobijn. 2010. "An Exploration of Technology Diffusion." *American Economic Review*, 100(5), 2031–2059.

David, Paul A. 1990. "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." *American Economic Review*, 80(2), 355–361.

Eloundou, Tyna, Sam Manning, Pamela Mishkin and Daniel Rock. 2024. "GPTs Are GPTs: Labor Market Impact Potential of LLMs." *Science*, 384(6702), 1306–1308.

Fagerberg, Jan. 1994. "Technology and International Differences in Growth Rates." *Journal of Economic Literature*, 32(3), 1147–1175.

Fan, Rachel Yuting and Ha Nguyen. 2026. "The Aggregate and Distributional Gains from AI: Global Evidence from Usage Data." Working Paper.

Felten, Edward W., Manav Raj and Robert Seamans. 2021. "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses." *Strategic Management Journal*, 42(12), 2195–2217.

Guevara, Miguel R., Dominik Hartmann, Manuel Aristarán, Marcelo Mendoza and César A. Hidalgo. 2016. "The Research Space: Using Career Paths to Predict the Evolution of the Research Output of Individuals, Institutions, and Nations." *Scientometrics*, 109(3), 1695–1709.

Hidalgo, César A., Bailey Klinger, Albert-László Barabási and Ricardo Hausmann. 2007. "The Product Space Conditions the Development of Nations." *Science*, 317(5837), 482–487.

Juhász, Sándor, Johannes Wachs, Jermain Kaminski and César A. Hidalgo. 2026. "The Software Complexity of Nations." *Research Policy*, 55(3), 105422.

Kamada, Tomihisa and Satoru Kawai. 1989. "An Algorithm for Drawing General Undirected Graphs." *Information Processing Letters*, 31(1), 7–15.

Lee, Keun. 2024. "Innovation–Development Detours for Latecomers: Managing Global–Local Interfaces in the De-Globalization Era." Cambridge: Cambridge University Press.

Neffke, Frank, Martin Henning and Ron Boschma. 2011. "How Do Regions Diversify over Time? Industry Relatedness and the Development of New Growth Paths in Regions." *Economic Geography*, 87(3), 237–265.

Oxford Insights. 2025. "Government AI Readiness Index 2025." Oxford Insights, https://oxfordinsights.com/ai-readiness/government-ai-readiness-index-2025/.

Santos Silva, J. M. C. and Silvana Tenreyro. 2006. "The Log of Gravity." *Review of Economics and Statistics*, 88(4), 641–658.

Online Appendix

How Does AI Diffuse Across Countries and Occupations? Early Evidence from One Million Usage Conversations in Over 100 Countries

A AEI Sample Structure

Each AEI wave samples approximately one million Claude AI conversations (Fan and Nguyen, 2026). Not all receive a usable occupation assignment at every level of aggregation. Figure A1 shows the full breakdown for R5 (February 2026), the wave used in the cross-sectional analysis.

Panel A: Global Level

```
1,000,000 conversations

├─ ONET classified: 972,636                          Classifier successfully processed
│
│    ├─ Mapped to task: 929,714                      e.g. "Write a Python function to sort a list"
│    │
│    └─ "none": 42,922                               Classifier says not occupational, e.g. "What is the weather?"
│
└─ "not_classified": 27,364                          "Filtered for privacy protection or could not be classified"
                                                     (Anthropic, 2026)
```

Panel B: Country Level

```
1,000,000 conversations

├─ geo_id = country code: 805,387 (176 countries)
│
│    ├─ WITH ONET breakdown: 801,151 (117 countries)
│    │
│    │    ├─ ONET classified: 565,455
│    │    │
│    │    │    ├─ Mapped to task: 533,094            e.g. "Write a Python function to sort a list"
│    │    │    │
│    │    │    └─ "none": 32,361                     e.g. "What is the weather?"
│    │    │
│    │    └─ onet_task = "not_classified": 235,696   Small-cell suppression: "(country, task) intersections with
│    │                                               < 15 conversations and 5 unique user accounts" (Appel et al., 2025; 2026)
│    │
│    └─ WITHOUT ONET breakdown: 4,236 (59 countries) "Minimum Observations: 200 conversations per country,
│                                                    applied in enrichment step" (Anthropic, 2026)
│                                                    e.g. Malta, Iceland
│
├─ geo_id = "NONE": 2,340                            IP geolocation returned no country; possibly VPN/anycast/hosting
│
│    ├─ ONET classified: 963
│    │
│    │    ├─ Mapped to task: 895
│    │    │
│    │    └─ "none": 68
│    │
│    └─ onet_task = "not_classified": 1,377
│
├─ geo_id = "not_classified": 181,586 [no ONET breakdown]   Possibly countries below 200-conversation threshold (Anthropic, 2026)
│
└─ No country-level data: 10,687
```

Figure A1: AEI sample structure tree (R5, February 2026). Panel A shows the global classification pipeline: 972,636 of 1,000,000 conversations are classified by the O*NET classifier, of which 929,714 map to an occupational task. Panel B shows the country-level breakdown: 176 countries receive a geographic code (805,387 conversations), of which 117 receive O*NET task enrichment (801,151 conversations). The 59 countries without O*NET breakdown and the remaining categories ("NONE," "not_classified," and no country-level data) are detailed in the tree.

A non-trivial share of conversations cannot be placed confidently into any occupation. Anthropic tags these conversations `not_classified`. At the global aggregate the unclassified share is 2.7%, but this rate is dominated by the United States, which contributes 28% of conversations and where the Anthropic classifier performs well. At the country level, the occupation lens performs less well: the country-median `not_classified` share is 66%, and countries such as Togo, Jamaica, Bahrain, and Luxembourg have more than 93% of their conversations unclassified. This asymmetry arises because the O*NET-SOC taxonomy was trained on U.S. labor-market data and labels non-U.S. task mixes less reliably.

B Country Lists

Table A1: Countries with positive Claude usage in the 2026-03-24 AEI release, by World Bank region ($N = 165$).

World Bank region	Countries
East Asia & Pacific	Australia; Brunei Darussalam; Cambodia; Fiji; Guam; Indonesia; Japan; Korea, Rep.; Lao PDR; Malaysia; Mongolia; New Caledonia; New Zealand; Papua New Guinea; Philippines; French Polynesia; Singapore; Taiwan, China; Thailand; Timor-Leste; Viet Nam
Europe & Central Asia	Albania; Andorra; Armenia; Austria; Azerbaijan; Belgium; Bosnia and Herzegovina; Bulgaria; Croatia; Cyprus; Czechia; Denmark; Estonia; Finland; France; Georgia; Germany; Gibraltar; Greece; Hungary; Iceland; Ireland; Italy; Kazakhstan; Kosovo; Kyrgyz Republic; Latvia; Liechtenstein; Lithuania; Luxembourg; Malta; Moldova; Monaco; Montenegro; Netherlands; North Macedonia; Norway; Poland; Portugal; Romania; San Marino; Serbia; Slovak Republic; Slovenia; Spain; Sweden; Switzerland; Tajikistan; Turkiye; Ukraine; United Kingdom; Uzbekistan
Latin America & Caribbean	Antigua and Barbuda; Argentina; Aruba; Bahamas; Barbados; Belize; Bermuda; Bolivia; Brazil; Cayman Islands; Chile; Colombia; Costa Rica; Dominican Republic; Ecuador; El Salvador; Grenada; Guatemala; Guyana; Haiti; Honduras; Jamaica; Mexico; Panama; Paraguay; Peru; Puerto Rico; St. Lucia; St. Vincent and the Grenadines; Suriname; Trinidad and Tobago; U.S. Virgin Islands; Uruguay
Middle East & North Africa	Algeria; Bahrain; Djibouti; Egypt, Arab Rep.; Iraq; Israel; Jordan; Kuwait; Lebanon; Morocco; Oman; Qatar; Saudi Arabia; Tunisia; United Arab Emirates; West Bank and Gaza
North America	Canada; United States
South Asia	Bangladesh; Bhutan; India; Maldives; Nepal; Pakistan; Sri Lanka
Sub-Saharan Africa	Angola; Benin; Botswana; Burkina Faso; Burundi; Cabo Verde; Cameroon; Chad; Comoros; Congo, Rep.; Cote d'Ivoire; Eswatini; Gabon; Gambia, The; Ghana; Guinea; Kenya; Lesotho; Liberia; Madagascar; Malawi; Mauritania; Mauritius; Mozambique; Namibia; Niger; Nigeria; Reunion; Rwanda; Senegal; Sierra Leone; South Africa; Tanzania; Togo; Uganda; Zambia; Zimbabwe

C Robustness Checks

C.1 Intensity: PPML

Table A2 re-estimates the intensity regression using Poisson Pseudo-Maximum Likelihood (Santos Silva and Tenreyro, 2006), which retains countries with zero usage as structural zeros (171 countries including 17 zeros, versus 154 in log-OLS). All five key regressors are significant with the same sign. Old-age dependency and internet users lose significance under PPML because the zero-usage countries are predominantly small economies where these variables add noise. The income elasticity is lower under PPML (0.354 vs. 0.384) because PPML weights the lower tail more heavily.

Table A2: Robustness: Progressive PPML for per-million AI penetration

	(M1)	(M2)	(M3)	(M4)	(M5)	(M6)	(M7)	(M8)
log(GDP per capita, PPP)	1.034***	1.046***	0.912***	0.794***	0.656***	0.587***	0.354***	0.352***
	(0.052)	(0.050)	(0.064)	(0.057)	(0.084)	(0.107)	(0.087)	(0.087)
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
log(Population)		0.035	0.014	0.038	0.000	-0.066	-0.160***	-0.161***
		(0.031)	(0.032)	(0.032)	(0.027)	(0.077)	(0.048)	(0.049)
		[0.252]	[0.651]	[0.241]	[0.994]	[0.387]	[0.001]	[0.001]
Old-age dependency ratio			0.022***	0.016***	0.016***	0.013**	0.008	0.008
			(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.006)
			[0.000]	[0.001]	[0.000]	[0.027]	[0.146]	[0.153]
Services VA (% GDP)				0.032***	0.029***	0.030***	0.022***	0.022***
				(0.005)	(0.004)	(0.004)	(0.004)	(0.004)
				[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Internet users (%)					0.010**	0.009*	0.005	0.005
					(0.005)	(0.005)	(0.004)	(0.004)
					[0.034]	[0.059]	[0.208]	[0.197]
log(Scientific articles)						0.060	0.125***	0.125***
						(0.075)	(0.046)	(0.047)
						[0.422]	[0.007]	[0.007]
AIPI: Regulation							4.899***	4.841***
							(1.192)	(1.213)
							[0.000]	[0.000]
English-official (0/1)								0.022
								(0.086)
								[0.797]
Observations	171	171	171	158	151	151	143	143
Pseudo R2 (McFadden)	0.693	0.697	0.734	0.790	0.826	0.828	0.866	0.866

Notes: Dep. var.: Claude conversations per million people (levels). PPML retains countries with zero usage as structural zeros. Each cell reports coefficient, (robust HC1 SE), and [p-value]. * p<0.10, ** p<0.05, *** p<0.01.

C.2 Entry: pre-determined R3 ubiquity

Table A3 re-estimates the entry logit using ubiquity from the pre-determined R3 wave (August 2025) instead of contemporaneous R5. All signs, significance patterns, and relative magnitudes are preserved. The lower pseudo R^2 (0.31 vs. 0.39) reflects the staleness of the six-month-old measure rather than bias in the R5 model.

Table A3: Entry model robustness: pre-determined R3 ubiquity (logit).

	(1)	(2)	(3)	(4)
Log GDP per capita	0.510^{***}	0.216	0.389^{*}	0.378^{*}
	(0.153)	(0.173)	(0.209)	(0.212)
Relatedness density (ω)		5.529^{***}	7.228^{***}	5.923^{***}
		(0.664)	(0.692)	(0.876)
Ubiquity (# countries in R3)			0.054^{***}	0.038^{***}
			(0.007)	(0.012)
$\omega\times$ ubiquity				0.073^{**}
				(0.037)
Constant	-7.616^{***}	-5.874^{***}	-9.695^{***}	-9.223^{***}
	(1.588)	(1.734)	(2.237)	(2.297)
Pseudo R^2	0.032	0.202	0.310	0.318
N	1653	1653	1653	1653

Notes: Same specification as Table 4 but with ubiquity measured using the pre-determined R3 wave (Aug 2025) instead of contemporaneous R5 (Feb 2026). All signs and significance patterns are confirmed. The lower pseudo R^2 reflects the staleness of the R3 measure: some occupations saw large ubiquity changes between waves (e.g., Education: 73 → 94 countries; Personal Care: 16 → 29; Production: 47 → 14). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

D Variance Inflation Factors

Table A4: Variance inflation factors.

Intensity (Table 3, M7)		*Entry (Table 4, model 6)*	
Variable	VIF	Variable	VIF
Log GDP per capita	7.49	Log GDP per capita	1.73
Log population	5.45	Relatedness density (ω)	3.87
Old-age dependency	2.59	Ubiquity (R5)	2.14
Services VA (% GDP)	2.13	$\omega\times$ ubiquity	3.69
Internet users (%)	4.16	Compute capacity	2.63
Log research intensity	7.33	Policy vision	1.66
AIPI Regulation	3.66		

Notes: All VIF values are below 10 (intensity) and below 4 (entry), indicating no multicollinearity concern. The highest intensity VIF (7.49) reflects the correlation between GDP per capita and research intensity; the highest entry VIF (3.87) reflects the mechanical correlation between ω and its interaction with ubiquity.

E Exit Model

Table A5 reports the exit logit described in Section 4.3. The sample consists of 701 country-occupation pairs with count > 0 in R3; the dependent variable is $\text{kept}_{c,i} = 1$ if the pair retains count > 0 in R5. The intensity measure is pair-specific: log of classified conversations in the country-occupation pair per million national population in R3.

Table A5: Exit model: retention of AI adoption by occupation (logit).

	(1)	(2)	(3)	(4)
Log GDP per capita	0.436^{***}	0.228^{**}	-0.298^{**}	0.191
	(0.132)	(0.112)	(0.145)	(0.179)
Relatedness density (ω)		1.573^{***}	4.658^{***}	6.154^{***}
		(0.420)	(0.882)	(0.913)
Log pair AI intensity			1.984^{***}	0.954^{**}
			(0.407)	(0.431)
$\omega\times$ log pair AI intensity			-1.333^{**}	-1.100^{**}
			(0.572)	(0.485)
Ubiquity (# countries in R3)				0.045^{***}
				(0.011)
Constant	-2.431^{*}	-1.230	0.812	-6.128^{***}
	(1.295)	(1.110)	(1.339)	(2.027)
Pseudo R^2	0.022	0.049	0.190	0.229
N	701	701	701	701

Notes: Logit with standard errors clustered by country. Dependent variable: = 1 if the country-occupation pair retains any classified conversations in R5 (Feb 2026), = 0 if it had conversations in R3 but none in R5 (lost adoption). Sample: 701 pairs with positive conversation counts in R3 (Aug 2025), i.e. occupations already adopted. Log pair AI intensity = log(classified conversations in this specific country-occupation pair per million national population +1) in R3. Unlike the entry model, pair-specific intensity is used here because all exit-sample pairs have positive counts in R3. Ubiquity = number of countries with count > 0 in that occupation in R3. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.