



EVIDENCE REVIEW / PART II

The International Comparison

How U.S. patterns of AI use compare with Europe, Asia and the emerging economies, and which findings survive the crossing.

Companion to AI Use Across Generations, Professions and Disciplines

September 2026

Version 1.0

The short answer

The United States is not the frontier. On population-level adoption it ranks twenty-fourth. Its workplace intensity is at or slightly above the global average and well below the emerging-economy leaders. Its enterprise adoption rate is indistinguishable from the European Union average. Where the United States is genuinely distinctive is in two places, and neither is the one usually assumed.

First, the American age gradient is the shallowest of any comparably measured market. Adults 65 and over in the United States use chatbots at roughly three times the rate of the equivalent European band. Second, American attitudes are among the most negative in the world, with trust in domestic AI regulation at a level Ipsos flags as an international outlier.

The finding that should most concern anyone building an argument on cohort exclusion appears in the European data. When Eurostat asked the sixty-seven percent of Europeans who had not used generative AI why not, 64 percent said they did not feel the need. Fourteen percent said they were unsure how to use the tools, 9 percent were unaware of them, and 7 percent cited privacy. That is the largest official dataset in existence bearing on whether non-use is a capability and access problem or a relevance judgment, and it points toward relevance. Section 6 treats this seriously rather than routing around it.

1. Why cross-country comparison is harder than it looks

The comparability problem flagged in Part I becomes acute internationally. Four specific hazards apply to everything that follows.

- Reference period. Eurostat asks about use in the previous three months. Pew asks whether a respondent has ever used a chatbot. This alone will widen any transatlantic gap, and it widens it most in low-frequency cohorts, which is precisely where the age comparison sits.
- Age ceiling. Eurostat caps its population at 74. Pew does not cap at all. The European 65-74 figure therefore excludes the oldest and lowest-using segment, which should narrow rather than widen the gap. That it does not is what makes the comparison interesting.
- Sampling base. Workplace surveys in emerging economies sample employed, connected, frequently urban respondents. In countries where formal salaried employment covers a minority of working adults, an 80 percent workplace figure and a 20 percent population figure are both correct and describe different populations.
- Unit of measurement. Population adoption, employee use, enterprise deployment and platform telemetry are four different things. Ranking countries by combining them, as several commercial indices do, produces nonsense. Every table below states its unit.

One outright conflict should be noted. Stanford’s own AI Index materials report the United Arab Emirates at 54 percent in some summaries and 64 percent in others. Neither figure is used as load-bearing here.

2. Population adoption: the United States ranks 24th

Stanford's 2026 AI Index puts global generative AI adoption at 53 percent of the population within three years of mass-market release, faster than the personal computer, which took roughly fifteen years to reach comparable penetration, and the internet, which took about ten. Adoption correlates strongly with GDP per capita, but the correlation has notable exceptions in both directions, and the United States is the largest negative one.

Market	Population adoption	Note
Singapore	61%	Highest measured; far above GDP prediction
United Arab Emirates	54%	Reported as 64% in some Stanford summaries
Global average	53%	Stanford AI Index 2026
Norway	56%	Eurostat, prior 3 months, ages 16-74
Denmark	48%	Highest in the EU (Eurostat)
Estonia / Malta	47%	Eurostat
Finland	46%	Eurostat
Switzerland	32.4% to 34.8%	Ranked 15th (Stanford); 47% on Eurostat basis
EU average	33%	Eurostat, prior 3 months
United States	28.3%	Ranked 24th (Stanford)
Bulgaria	23%	Eurostat
Italy	20%	Eurostat
Romania	18%	Lowest in the EU (Eurostat)

Units differ between the two sources. Stanford measures population adoption on its own basis; Eurostat measures use in the previous three months among ages 16 to 74. Switzerland appears on both and the two methods disagree by roughly thirteen points, which is a fair estimate of the noise in any cross-source comparison in this table.

The eight-fold spread inside the European Union, from Romania at 18 percent to Denmark at 48 percent, is larger than the spread between any two major blocs. Whatever explains national variation is operating more powerfully within Europe than between Europe and North America. Eurostat attributes it principally to general digital foundation, with Nordic and Baltic states leading and Southern, Central-Eastern and Balkan states trailing.

The American ranking deserves care rather than alarm. The United States leads the world in AI investment, model development and research output while ranking twenty-fourth in population use, and the estimated consumer value of generative AI tools to Americans reached \$172 billion annually by early 2026, with median value per user tripling in a year. A low population adoption rate alongside high measured consumer surplus suggests intensity concentrated in a narrower segment, which is consistent with the American education and income gradients documented in Part I.

3. Europe has a far steeper age gradient than the United States

This is the most directly relevant comparison for any instrument measuring cohort distance, and it runs opposite to the intuition that the United States is the outlier on generational division.

Cohort	EU (Eurostat, prior 3 months)	U.S. (Pew, ever use)	U.S. (Pew, daily use)
Youngest measured band	64% (ages 16-24)	66% (ages 18-29)	31% (ages 18-29)
Oldest measured band	7% (ages 65-74)	23% (ages 65+)	7% (ages 65+)
Ratio, youngest to oldest	9.1 to 1	2.9 to 1	4.4 to 1
All adults	33%	49%	24%

The youngest cohorts are effectively identical, 64 percent against 66 percent. The oldest cohorts are not close. The honest caveat is that Eurostat’s three-month window is stricter than Pew’s ever-use question, and that difference will inflate the European ratio. But the comparison survives a conservative correction: even measured against American daily use, which is a far stricter standard than Eurostat applies, the European ratio is more than twice as steep. And Eurostat caps its sample at 74, excluding the segment that would drag the European figure lower still.

Two further Eurostat cuts sharpen the picture and both echo the American findings in Part I. Education is a powerful predictor: 49 percent of Europeans with high formal education used generative AI against 26 percent with medium and 22 percent with low. Gender is a weak one: 35 percent of men against 30 percent of women. And across every working-age band from 25 to 64, more than half of European users report using generative AI for work purposes, which means the work-use pattern is stable across the working-age population even as overall use declines with age. The decline is in whether people use the tools at all, not in what they use them for once they do.

One structural asymmetry is worth naming. Eurostat records 25.1 percent of Europeans using generative AI for personal purposes against 15.1 percent for work. The European gap between personal and professional use is wider than the American one, and the EU AI Act’s compliance regime is the most commonly offered explanation. If that explanation is right, part of the European age gradient is a regulatory artifact operating through employers rather than a cohort characteristic.

4. Workplace intensity inverts the wealth gradient

Population adoption tracks GDP per capita. Workplace intensity does not. It runs the other way.

Stanford reports 58 percent of employees globally using AI at work on a semi-regular or regular basis, with the share exceeding 80 percent in India, China, Nigeria, the United Arab Emirates, Egypt and Saudi Arabia, against 40 to 48 percent across most of North America and Europe. ADP Research’s People at Work 2026, surveying more than 39,000 working adults across 36 markets, corroborates this independently: about 50 percent of employees globally use AI multiple times a week and 20 percent almost daily, while India records 80 percent weekly and 41 percent daily, the highest levels in the study. Nigeria follows at 39 percent daily and Vietnam at 36 percent.

Set against that, Anthropic’s platform telemetry shows the opposite ordering. Measured as usage per capita relative to working-age population, Singapore runs at 4.6 times expected and Canada at 2.9 times, while Indonesia sits at 0.36, India at 0.27 and Nigeria at 0.2.

These are not contradictory findings, and the reconciliation is the useful part. Workplace surveys in emerging economies sample formal-sector, connected, largely urban employees. In India and Nigeria that population is a small minority of all working adults. Platform telemetry measures the entire population. The two instruments therefore describe a formal-sector professional elite and a nation, and both descriptions are accurate. The practical implication is that a global benchmark built from workplace surveys will systematically overstate emerging-economy adoption relative to one built from population data, and any normative comparison must hold the sampling frame constant or it is measuring national formal-employment rates dressed up as AI readiness.

Anthropic’s January 2026 report adds a finding that bears directly on the quality of use rather than its volume. Higher-income, higher-usage countries show more augmentation, meaning collaborative back-and-forth use, while lower-income, lower-usage countries show more automation, meaning directive task completion. The report frames this explicitly as an adoption-curve effect rather than a cultural one, on the grounds that the same shift toward directive delegation is visible over time within markets as they mature: directive task delegation rose from 27 percent to 39 percent globally over eight months. Whether the income pattern is a stage or a settled difference is not resolved by the data.

Firm size and sector are universal

Three official instruments measuring enterprise adoption by three different statistical agencies land within half a percentage point of one another.

Instrument	Enterprise AI use	Reference
Eurostat, EU27 enterprises with 10+ employees	20.0%	2025 reference year

Instrument	Enterprise AI use	Reference
U.S. Census Bureau, Business Trends and Outlook Survey	19.8%	May 2026
OECD firm average, ICT Access and Usage	20.2%	2025

The convergence of these three figures is the strongest evidence in this review that enterprise AI adoption in advanced economies is a single phenomenon rather than a set of national ones. Beneath the average, national variation is wide: Denmark at 42.0 percent, Finland at 37.8, Sweden at 35.0, Belgium at 34.5, Luxembourg at 33.6 and the Netherlands at 33.2, against Romania at 5.2, Poland at 8.4, Bulgaria at 8.6 and Greece at 8.9. Twenty-six of twenty-seven EU member states recorded increases over 2024, with Denmark rising 14.5 points in a single year.

Sector concentration is identical everywhere it has been measured. In the EU, Information and Communication leads at 48.7 percent, followed by professional, scientific and technical services at 29.5 percent and financial and insurance activities at 26.9 percent, with manufacturing at 17.5 and wholesale and retail at 15.1. The American ordering, Information at 39.7 percent and Finance and Insurance at 33.9 against a national 19.8, is the same ordering at slightly different levels. Firm size dominates in both. ADP found the same pattern inside a single country: in India, daily AI use runs at 54 percent in firms of 1,000 or more employees against 32 percent in small firms.

5. Attitudes run opposite to capability

The Ipsos AI Monitor, now in its fifth year and covering 32 countries as of June 2026, documents a divide that has held stable across waves. People in Asia and Latin America are more likely to think AI has more benefits than drawbacks and to be excited about products using it. People in Europe and North America are more likely to be nervous. Enthusiasm is highest where frontier models are not built and lowest where they are.

The United States is the sharpest case. Ipsos flags it as an outlier in its low-trust and high-nervousness combination, with only 31 percent trusting the government to regulate AI responsibly against a 30-country average of 54 percent, in markets where Singapore, Indonesia, Malaysia and Thailand sit far higher. Trust that companies using AI will protect personal data ranges from 72 percent in Thailand to 32 percent in France, Japan and the United States. Stanford records the same shape from a different direction: a fifty-point gap between AI experts, 73 percent of whom expect a positive impact on jobs, and the American public, at 23 percent.

The generational attitude finding from Part I is now confirmed as international rather than American. Ipsos 2026 finds people under 35 the most likely of any age group to be nervous about AI, at 52 percent across 32 countries, and simultaneously the most likely to be excited, at 56 percent. This is the same pattern Pew found in the United States, where the heaviest-using cohort is also the most pessimistic about societal impact. It is not an artifact of American politics or of a single instrument. Younger cohorts everywhere hold both positions at once, and older cohorts are more likely to hold neither: Pew found adults 65 and over the most likely to say they were unsure.

Two attitude findings cut across the enthusiasm map rather than along it. Across 32 countries, 46 percent of respondents say they would trust a generative AI tool less if its answers were influenced by advertisers, with a majority saying so in fifteen markets. And only 43 percent worldwide believe business leaders tell the truth, which is a plausible upstream cause of AI-specific distrust rather than a consequence of it.

6. The strongest challenge to a cohort-exclusion thesis

Eurostat asked non-users of generative AI why they had not used it. The answer distribution is the most consequential single result in the international literature for anyone arguing that non-use reflects denied access rather than exercised judgment.

Main reason for not using generative AI (EU, 2025)	Share of non-users
Did not feel the need to	64%
Not sure how to use the tools	14%
Not aware the tools existed	9%

Main reason for not using generative AI (EU, 2025)	Share of non-users
Privacy and security concerns	7%

Read straight, this says that skill and awareness together account for less than a quarter of European non-use, and that the dominant reason is a judgment about relevance. That is difficult to reconcile with a pure exclusion account, and an evidence review that suppressed it would not be worth commissioning.

Three qualifications are legitimate, and none of them is a dismissal.

- The question asks for a main reason and offers a socially costless option. "Did not feel the need" is the answer that requires no admission of unfamiliarity, and it will attract respondents whose actual position is closer to not being sure what the tools would do for them. This is a known limitation of single-reason non-use items, and it is why the Age Intelligence Index item bank should separate perceived relevance from perceived permission rather than collapsing them.
- Relevance is itself partly organizational. A worker whose employer has not connected AI to any task in their role has been given no occasion to need it. "I did not feel the need" and "no one showed me what it was for" are not distinguishable in this instrument, and the distinction is exactly what an organizational diagnostic exists to draw.
- The finding is about the general population, not about workers inside an organization that has deployed AI. The access and training data in Part I, drawn from Randstad, AARP, Generation, ManpowerGroup and the New York Fed, all concern people who are inside such organizations. The two bodies of evidence describe different populations and are not in direct conflict.

The defensible position is narrower than the one usually asserted and stronger for being narrower. Most population-level non-use is elective. Most within-organization variation in access, training and permission is not. An instrument that measures the second while acknowledging the first is credible. One that treats all non-use as exclusion is not, and the Eurostat table is the specific evidence a skeptical reader will bring.

7. What travels and what does not

Findings that hold in every market measured

- The education gradient. High-education adoption exceeds low-education adoption by roughly twenty-five points in the EU and by thirty-six points in the United States. No exception has been reported.
- The firm-size gradient. Larger firms adopt more, everywhere, including within high-adoption emerging markets.
- Sector ordering. Information and communication first, professional and technical services second, finance third, manufacturing and retail well behind. Identical in the EU, the United States and the OECD aggregate.
- The direction of the age gradient. Younger cohorts use more in every market with published age data. Only the steepness varies.
- The training gap. ManpowerGroup found 56 percent of the global workforce across nineteen countries reporting no recent training and 57 percent no mentorship access, with confidence falling 18 percent even as use rose 13 percent, and the sharpest confidence declines among boomers at 35 percent and Gen X at 25 percent. Nothing in the international data suggests any market has solved this.
- The gap between adoption and integration. Microsoft's 2026 Work Trend Index surveyed 20,000 knowledge workers who already use AI at work, across ten markets, fielded by Edelman Data X Intelligence between February and April 2026. Only 16 percent qualified as what Microsoft terms Frontier Professionals, users who have redesigned workflows around the technology. Because the sample is drawn entirely from existing AI users, the figure describes depth of use among adopters and understates how rare that depth is across the workforce as a whole. Stanford found fewer than 10 percent of organizations fully scaled in any single function.
- The seniority effect in research. The Nature and Fudan survey of 10,480 researchers found frequent use declining with experience in every region, with different levels but the same direction: among researchers with fewer than three years of experience, heavy content-discovery use reached 21.2 percent in Japan and 15.2 percent in China against 5.7 percent in the United States and 8.9 percent in Europe, and among those with more than ten years, 8.6, 6.0, 5.2 and 4.3 percent respectively.

Findings that do not travel

- Adoption levels. An eight-fold spread inside the EU alone.
- The steepness of the age gradient. Roughly 9 to 1 in Europe against 3 to 1 in the United States on ever-use.

- Attitudes. Excitement and nervousness are near mirror images between Asia-Pacific and Latin America on one side and Europe and North America on the other.
- Tool preference at the margin. General-purpose models dominate everywhere and ChatGPT ranks first in every country measured, but DeepSeek accounts for 48.8 percent of Chinese researchers' tool mentions against under 10 percent everywhere else.
- The relationship between workplace intensity and national wealth, which inverts entirely depending on whether the sampling frame is employees or population.

And a caution that travels further than the enthusiasm

Yotzov, Barrero, Bloom, Bunn, Davis and colleagues, in *Firm Data on AI* (NBER Working Paper 34836, February 2026), surveyed almost 6,000 CFOs, CEOs and executives from stratified firm samples in the United States, the United Kingdom, Germany and Australia, fielded in January 2026. They report 69 percent of firms actively using AI while 89 percent report no impact on their labor productivity, measured as volume of sales per employee, over the previous three years; the abstract states the broader result as more than 80 percent reporting no impact on either employment or productivity. This sits alongside task-level experimental results, 15 percent in customer support from the published Brynjolfsson, Li and Raymond study and a reported 26 percent in software development, which measure individual throughput rather than firm outcomes. The reconciliation is presumably the integration gap: measured task gains are real and firm-level effects have not yet appeared because workflows have not been redesigned. But that is an interpretation, not a finding, and four advanced economies now report the same null result.

8. What this means for the Index

Norms must be national, and the roadmap should say so. A 45+ Delta calibrated on U.S. data will understate the gap in Europe by a wide margin and may misread it entirely in the highest-intensity emerging markets, where ADP found 43 percent of Indian workers aged 18 to 39 in near-daily use with the rest of the workforce not far behind. The Phase Two normative stage described in the Index roadmap should either scope itself explicitly to the United States or plan for national norm sets from the outset. A single global percentile band would be actively misleading.

The Permission domain should separate three things the Eurostat instrument collapses. Perceived relevance, perceived permission and perceived capability produce identical non-use and require completely different interventions. Eurostat's 64 percent "did not feel the need" is the largest official measurement in this space and it cannot distinguish them. An item bank that can is a genuine methodological contribution, and it is also the Index's best answer to the sharpest external objection it will face.

The Access domain reads differently outside the United States. Only 11.1 percent of researchers worldwide report institutional funding for their primary AI tool, with 41.3 percent

paying personally. In markets with high individual use and low institutional provision, an Access score will diverge sharply from a Fluency score, and the diagnostic quadrant will place organizations differently than a U.S.-calibrated reading would. This is a feature if the instrument is designed for it and a defect if it is not.

The Accountability thesis is the most portable finding in the review. The seniority effect held direction across Japan, China, the United States and Europe. Deloitte’s judgment-versus-technical-skill split was measured in Australia, not the United States, and Part I applied it to an American argument. The proposition that experience-linked judgment holds up better than tool recency is one of very few findings in this literature with cross-national support, and it remains under-measured everywhere. That combination, portable and unmeasured, is what an instrument should be built on.

One position to abandon. The claim that generational AI gaps reflect organizational failure rather than individual choice is well supported inside organizations and poorly supported at population level. The Eurostat non-use data is public, official, and will be found by anyone who looks. Conceding it explicitly, and scoping the thesis to the within-organization case where the evidence is strong, costs very little and forecloses the most damaging line of attack.

Additional sources for this part

Source	Population and n	Field or reference date
Eurostat, Use of AI by individuals (isoc_ai_iaiu)	EU and EFTA individuals aged 16-74	2025 ref. year, extracted Mar 2026
Eurostat, Use of AI in enterprises (isoc_eb_ai)	EU enterprises with 10+ employees	2025 reference year
Eurostat, Digitalisation in Europe 2026 edition	EU households and businesses	2026
Stanford HAI, AI Index 2026, Public Opinion and Economy chapters	Global, multi-source compilation	2026
Ipsos AI Monitor 2026 (Global Advisor)	32 countries	Jun 2026
Ipsos AI Monitor 2025 (Global Advisor)	30 countries, n = 23,216 online adults under 75	Mar to Apr 2025
ADP Research, People at Work 2026	Working adults, 36 markets, n = 39,000+	2026
Google and Ipsos Multi-Country AI Survey 2026	21 countries	Sep to Oct 2025
Microsoft, 2026 Work Trend Index	Knowledge workers who use AI, 10 markets, n = 20,000	Feb to Apr 2026
Yotzov et al., Firm Data on AI, NBER WP 34836	CFOs, CEOs and executives, 4 countries, n = ~6,000	Fielded Jan 2026
OECD, ICT Access and Usage by Businesses	OECD firms	2025
Anthropic Economic Index, geographic and enterprise report	Claude.ai and API, 150+ countries	Sep 2025
Anthropic Economic Index, January and March 2026 reports	Claude.ai and API, global	through Feb 2026

Verification note. Eurostat and Census figures were taken from the statistical agencies directly. Stanford AI Index figures were taken from HAI’s own summary pages where available and from secondary reporting otherwise; the UAE discrepancy noted in Section 1 arises between two HAI-published summaries and should be resolved against the full report before any external citation. NBER Working Paper 34836 and the Microsoft Work Trend Index have since been verified against the working paper and Microsoft’s own published materials respectively, and the figures here are corrected accordingly. ADP figures remain drawn from press materials rather than the primary report and still warrant confirmation. The Ipsos 2026 figures come from the published report PDF; the 2025 comparison figures come from the prior wave and should not be presented as a single trend series without checking that country coverage is consistent between waves.