

Beyond Touching ChatGPT: Measuring the Real Depth of AI Adoption Across Firms, Industries, and Countries

The Economy Research Editorial^{1,2}

¹The Economy Research, 71 Lower Baggot Street, Dublin 2, Co. Dublin, D02 P593, Ireland

²Swiss Institute of Artificial Intelligence, Chaltenbodenstrasse 26, 8834 Schindellegi, Schwyz, Switzerland

Abstract

AI adoption has traditionally been measured as equivalent to access, experimentation, and organizational change. This article argues that this definition is conceptually mistaken and politically undesirable. Generative AI tools, through the visibility provided to workers, consumers, and managers, may have made it possible to touch the technology, but do not necessarily lead to labor-saving or a lasting competitive advantage. What ultimately constitutes "adoption" relies upon integration of the technology into workplace processes, data systems, managerial routines, staffing configurations, and decision-making architectures. The developing distinction is, therefore, not one between adopters and non-adopters but between a superficial and a system-level adoption. This observation also serves to illustrate the distinct rate of adoption between and among sectors and countries. Digitized, text-based and modular sectors fare better and proceed with a greater pace, given that their tasks can be rendered more easily into machine-readable procedures, while physical, regulated, and tacit knowledge-based industries exhibit more of a lag. At the country level, while the number of individual workers using AI may rise quickly, productive utilization will continue to be concentrated in regions where the institutions, infrastructure, computing power, managerial talent, and model-development capability remain highest. While policymakers can increase a minimum adoption floor through the introduction of training programs, infrastructure investment, public data and standards, and guidance for appropriate use, the capacity to recreate private routines and expertise of leading-edge firms is still lacking. It follows that the core policy challenge is not about increasing the sheer number of people who have had contact with the technology, but about measurement and governance of the depth to which it alters work, who profits from this change, and whether this adoption exacerbates rather than ameliorates existing inequalities.

1 Introduction - From Chatbot Access to Real AI Adoption

The most pervasive and consequential mistake in current discussions about artificial intelligence is confusing access with adoption. While it is abundantly clear that by 2025, generative AI had clearly reached mass consumer status: more than a third of all individuals in the OECD used generative AI tools and 33% in the EU (per Eurostat data) used them within the past three months. That number falls dramatically when the subject shifts from workers to the firm-level, where the firm decides to raise productivity and where labor demand is based and a competitive advantage of the firm exists. The OECD showed that by 2025, only 20.2% of firms were using AI, with Eurostat finding a 19.95% number for EU firms using at least one form of AI technology.^[1] This is not statistical noise. It highlights the fundamental differences between an employee touching a chatbot and that employee's company implementing AI into production. Just because one worker has touched a chatbot does not necessarily imply a firm has already reorganized workflows, cut certain job roles, or formed a dependence on the output of these programs.

This distinction is vital since many current numbers either measure different things or treat these different metrics as equivalent. One survey showed 78% of organizations use AI in some function, but explained that what was meant by "adopted" varied from occasional, isolated experimentation to a firm-wide integration with revamped work processes. The Federal Reserve used data from several U.S. surveys, which highlighted the size of this measurement gap. A firm-weighted estimate from the Census Bureau said 18% of firms use AI, while a labor-weighted survey of senior management suggested that 78% of all employees work for an AI-adopting firm. Making the gap wider, the Census Bureau's wording for AI use changed from its definition in production to simply "any use" in business functions, requiring new time series measurements starting in November 2025. The implication is clear, that measuring AI adoption is the argument itself, not just a part of it.^[2]

A more accurate assessment of AI adoption must rely on a multi-tiered index rather than a single percentage. Access permission to use, paid subscriptions, or availability. Usage intensity: hours spent per work session and the number of functions that rely on AI. Depth of usage: how high-stakes the work performed by AI is, ranging from simple writing assistance to managing delegated workflows. Integration within the organization: how the use of AI impacts job design, decision-making workflows, organizational structures and data systems. Finally, the impact of AI changes in cost structure, workforce size and composition and per-worker revenue. AI's capabilities are further illustrated by a B2B Signals report that showed firms using the most AI at the 95th percentile utilized 3.5x the amount of intelligence per employee as a typical firm. While prompt volume accounts for only 36% of this gap, it is the greater depth and complexity of work delegated to AI that explains the bulk of the difference, with leading firms sending 16x as many Codex messages per worker as the average firm. Further, a recent survey found that while 70% of surveyed firms were actively using AI, top managers averaged only 1.5 hours per week using it themselves and 80% said AI had had zero effect on either jobs or output in the preceding 3 years. In short, usage is not synonymous with impact.^[3]

This shift is crucial given the acceleration between 2023 - 2026, where capability development, consumer access and capital inputs all advanced at a rapid pace. The Brookings Institution analysis on this phenomenon is useful for differentiating frontier capability, diffusion and distributive impacts-three concepts that are routinely

collapsed into one narrative of AI progress. Brookings pointed out that macro trends might lag despite the sharp increases in capability due to complementary investments needed to enable diffusion of AI; such investments include skills training, data accessibility, managerial expertise and structural reforms. Stanford's AI Index report mirrored these findings by detailing significant growth in business usage and private investment, in addition to compute-training growth of doubling every 5 months. The question is not whether AI exists. It is who is creating the necessary complementary investments to leverage broad access to transform that access into sustainable output growth, increased bargaining power and labor market advantages. This is the real adoption gap and where the issue of redistribution becomes evident.^[4]

2 Degrees of AI Adoption - Firms: Access Adopters and Systemic Adopters

As it stands, the evidence seems to suggest a dichotomy in AI adoption at the firm level. Access adopters—those firms that have secured licenses, given their employees permission to explore, or used AI in a singular, standalone capacity but have not transformed their operational structures. Systemic adopters—those firms that are integrating AI through new organizational processes, codifying internal guidelines, reallocating workforces and relying on it not simply for assistance but for a substantive shift in the border between human discretion and machine output. These two are not simply semantic differences; they represent distinct phenomena. Evidence from the 2026 NBER survey indicated that AI usage was particularly high for younger and highly productive firms. The Federal Reserve's data synthesis also pointed to smaller firms using it more frequently than size alone would suggest, though larger firms continue to have the highest overall adoption levels. Such a finding implies both that some smaller firms may be empowered by AI to enter the market and grow quickly and that the advantages of size remain in terms of resources necessary to deeply integrate AI across the enterprise.^[5]




Figure 1: AI adoption appears much higher when measured by workers or employment-weighted firms than when measured by firm count.

The key variable, therefore, is not whether firms say they are using AI, but rather the maturity of this adoption into business capabilities. Here, the findings are far less optimistic than many public discussions claim. McKinsey’s research revealed that only 1% of survey respondents considered their organization’s deployment of generative AI to be mature. Fewer than one-third reported using most of the twelve different AI adoption/scaling practices that predicted strong financial impact and fewer than one in five thought they were well on their way to hitting clearly defined key performance indicators for their generative AI initiatives. Despite 71% of respondents in the same McKinsey survey claiming to use generative AI regularly in at least one function, the reality was that of shallow usage and lack of firm-wide consolidation rather than robust adoption.^[6] Thus, simply measuring firms that are using AI somewhere is not an effective metric. A firm that uses ChatGPT for marketing copy drafting, for example, is in a different empirical category from a firm that is reorganizing its purchasing, software development, customer service and management oversight structures around AI.




Figure 2: Firm size remains a major adoption divider, showing that access to AI tools does not remove organizational capacity gaps.

The top-performing firms highlighted in OpenAI’s B2B Signals data serve as an excellent example of deeper adoption.^[7] Their success derives not from simple prompt volume but from providing more contextual input for AI models, delegating more complex tasks and shifting from simply interacting with AI to enabling agentic, workflow-oriented systems. This necessitates complementary investments, which would be impossible to gauge through a basic adoption metric, such as enhanced internal documentation, clearer separation of tasks, greater oversight over AI usage, clearer risk management protocols and managerial willingness to adjust performance standards. Such firms are truly embedding AI into their production process rather than just using it as an optional convenience layer over an existing framework. When an organization reaches this point, the know-how required is implicitly organizational rather than explicit, in-house and commercially sensitive. The ‘best practices’ for using AI are not easily replicable by outside organizations. Instead, they reside within firm-specific data, internal motivations and organizational routines.




Figure 3: Most organizations remain in experimentation rather than scaled agentic deployment, reinforcing the gap between surface use and operational transformation.

This is likely why a few digital-native companies have successfully condensed the scope of operations that had until now necessitated larger, specialist teams. Marketing and sales, product development and service offerings, customer service and support, software development and IT all witnessed the widest deployment of generative AI. Text is the dominant output, though over a third of firms are creating images and over a quarter are producing code.^[8] This indicates a clear shift towards more condensed functions: Pipelines for content, task automation for translation, client-facing material production, rapid prototyping and part of the debugging process may be managed by considerably smaller workforces than were previously required. However, condensing a process is not synonymous with full autonomy: the burden simply moves. Automated drafting, coding or graphic production now increases value for the steps of specification, validation, evaluation of correctness, security inspection and responsibility for risks. This decreases reliance on some specialists while simultaneously raising dependence on others to manage those areas.^[9]

This provides narrow, but critical, policy guidance: Policy intervention is most impactful at the beginning of adoption. Policies that assist in managing informational obstacles and providing complementary resources are critical. Roughly 75% of firms using AI cited public support for information or advisement, with 58% using such programs for training and 42% using programs designed to support financial access.^[10] Those with significant obstacles were more likely to utilize publicly available information and advisement services, whereas larger firms used such programs less often than small or medium-sized firms. Simultaneously, broad majorities agreed on the need for public support for retraining and AI training at university and vocational levels, for a robust marketplace of software vendors, for reliable internet infrastructure and for access to high quality public data-but this all supports moving organizations from an early, immature stage to moderate operational use; this is unlikely to bridge the gap separating basic adopters from cutting-edge adopters already incorporating

AI as a distinct competitive advantage.

The result is a divergence at the firm level. This gap does not just fall between adopters and non-adopters, but between organizations that view AI as merely a tool for their workers and those that consider AI a means of transforming organizational structure. This impacts output and profit margins but also the competitive landscape. The Federal Reserve already hinted that this differential adoption pattern among firms may lead to increased competition and concentration. While smaller, newer firms may leverage AI to reach scale faster, large organizations are still best positioned to deploy AI broadly across their workforce and leverage the scale of internal data and organizational systems. The end result likely will not be mass democratization, but a recalibration of advantage, wherein some smaller firms accelerate past peers and many conventional firms lag at superficial levels while a core minority breaks away through deep, systemic adoption.^[11]

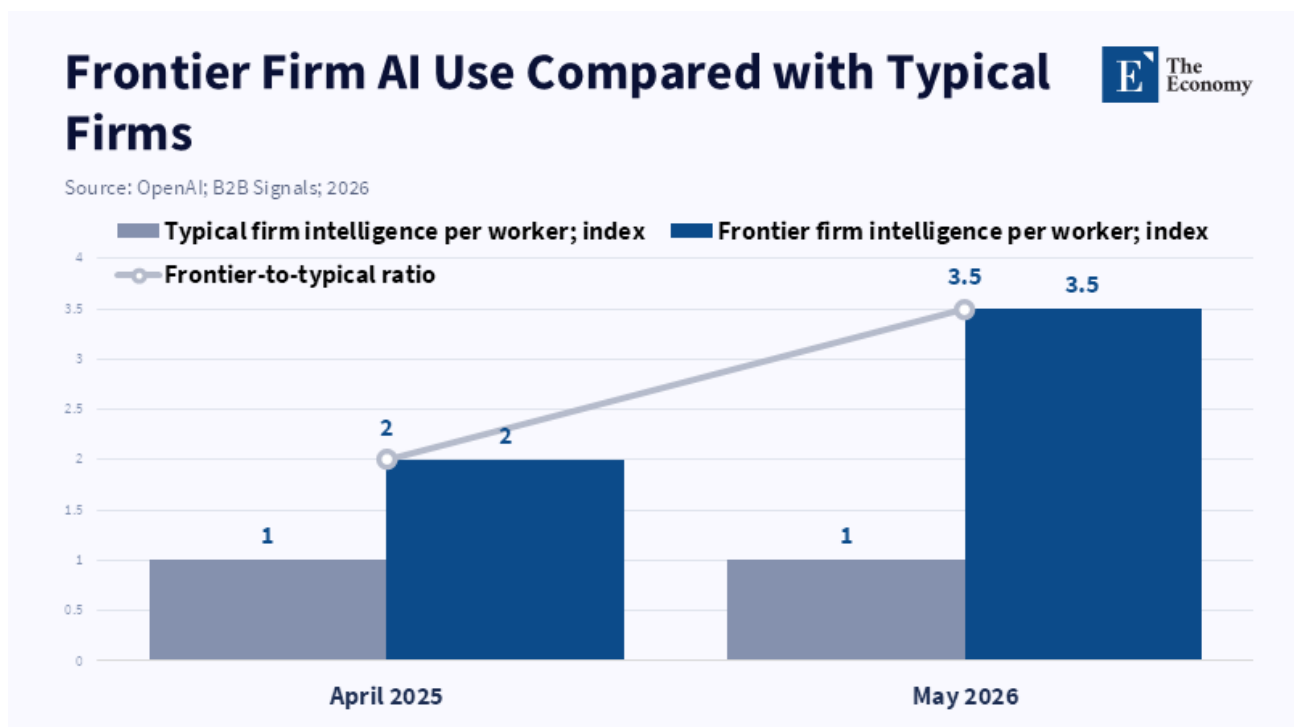


Figure 4: Frontier firms are pulling away not simply through more frequent AI use, but through deeper delegation, richer context and more complex workflow integration.

3 Degrees of AI Adoption - By Industry: Digitized Work, Selective Compression and Physical Bottlenecks

While some researchers have emphasized that at the industry level AI adoption is more a function of the nature of the work rather than of comparisons between high-tech and low-tech sectors, such patterns appear to be largely consistent. The research suggests that LLM adoption has proliferated within sectors of business that already contain high proportions of text-intensive, data-rich, modular and digitized workflows. There is already some data to support this: in 2025, 57.3% of ICT firms reported having used AI, compared to 36.8% of professional and scientific service firms (OECD data) with EU firm-level data from Eurostat showing similar figures of 62.52% of firms in Information Communication and 40.43% in Professional, Scientific and Technical Services with no other sectors showing figures over 25%.^[12] In the US professional services and finance sectors

are also reported to have some of the highest use cases (around 33% and 30% respectively) by the Federal Reserve (using the BTOS data). Thus it is evident that the patterns of initial adoption are driven by the inherent ability of AI to facilitate structures of work that resemble digitized and analyzed work with a large component of linguistic interpretation.




Figure 5: AI adoption is highest in sectors where work is already digitized, modular and information-intensive, while more physical sectors remain slower to transform.

This evidence is compounded when considering not just if firms are using AI, but how much. The NBER report suggests that in the early adopting Computer, Mathematical and Management occupations, users spend anywhere between 9-12% of their work day interacting with generative AI and save between 2.1-2.5% of their time relative to baseline, compared to personal service workers who spend 1.3% of their time on the task and saves 0.4% of their time. Similarly, information services use generative AI for 14.0% of work hours, the highest usage, to save 2.6% of hours, whereas Leisure and Hospitality use it for only 2.3% of hours and save a relatively small 0.6% of hours. The implication is that currently, the use of LLM technologies is not diffuse across the board, but is concentrated in fields with task structures that lend themselves readily to automation and data analysis.^[13]

The IT sector clearly shows very early adoption patterns; however, this does not indicate that all of its work can be considered routinized search jobs ripe for automation. The field experiment evidence also points to high levels of variation, with MIT data showing that when provided with tools, 4,867 software developers could produce 26.08% more tasks (this benefit was even higher for less experienced workers), whereas the study by METR showed that providing these tools to expert developers with significant experience of existing systems and problem areas led to a 19% increase in work time required. This seeming contradiction reveals the critical fact about AI tools that task structures should readily translate well to machine completion with human interaction:^[14] this would involve work tasks that can be effectively insulated from the broader system,

that do not require detailed knowledge of other parts of the project or code base to function properly and that do not entail extensive decision making.




Figure 6: AI productivity effects vary sharply by task context, rising in structured work but weakening where tacit knowledge, system familiarity and contextual judgment dominate.

As such, instead of job destruction in the IT sector, the best description appears to be "selective compression." Routinized tasks like basic code writing, first-level tech support, automatic document creation and generation and preliminary drafting are ripe for replacement or significant reduction. However, things like software design and integration architecture, security concerns and the establishment and validation of standards are clearly resistant to automation. This does not necessarily mean mass job losses in IT, but it could have dramatic effects on employment and working conditions within certain tiers. IT technicians will face a reduction in some core tasks, while those overseeing processes and making high-level design and security decisions may find themselves with even more responsibility and work. This will likely mean decreased traditional IT career progression and could impact those starting their careers in the field the most^[15] (congruent with the IMF analysis showing more usage of AI skills among less-exposed workers leads to lower demand).

In many other knowledge-intensive fields, the core effect has been augmentation, with field experiments supporting this across multiple sectors. One of the most promising such studies from the NBER found that providing a customer support AI tool to workers leads to 14% increased productivity.^[16] This effect is greater for less experienced workers; the quality of the task performed for the customer and employee is increased and employee satisfaction also goes up. This suggests that augmentation has potential and new technologies could even diffuse best practices to those less experienced workers, rather than replacing them altogether. The World Bank report suggests a growing trend of using AI as a partner and assistant and not as a direct replacement of labor (especially in developing nations)^[17] and this behavior is similar among the research/professional/business analytics sectors. Speed improvements in things like drafting, summarization and analysis are certainly happening,

yet human oversight, judgment and final review remain central to these tasks.

Outside of the information industries, there has been real, though much slower, adoption with more critical limitations related to the physicality of embodiment, regulatory environment and capital limitations. OECD data show that certain formerly slower moving industries have seen the biggest increase in AI usage; however, 62.5% for accommodation food services and 59.1% for construction do not represent the rate of adoption but the year-on-year growth.^[18] ICT is still out in front with 57.3% of all OECD firms in ICT utilizing AI in 2025, compared with professional, followed by professional and scientific services at 36.8%. Autonomous driving is not yet purely experimental and has over 150,000 weekly autonomous rides with a single US carrier. However, the eventual complete human displacement in the transport industry is distant due to issues with fleet economics, liability issues, mapping technology and weather conditions, among others; physical AI is certainly an important part of future discussion, but progress here will be significantly slower and more hardware dependent.

Taken together, the industry-level evidence should be interpreted not as a defense or condemnation of technology, but as a nuanced picture that highlights the current limitations of AI beyond mere access. AI is no longer simply a search tool, with systems now able to write code, create images, generate text and suggest actions, with the frontier systems even executing parts of work processes for top firms. It is not, however, a wholesale destroyer of jobs. Its impacts have been greatest where tasks already existed in digital forms and could be readily translated to machine processing and it is less effective when work requires embodiment, tacit knowledge, trust and faces significant regulatory hurdles. Industries early on are therefore not seeing universal replacement but instead a greater potential to optimize job compositions to increase labor efficiency and value distribution within the firm.

4 Degrees of AI Adoption - By Country: Usage Gaps, Productive Control and the New Diffusion Divide

While data on a global scale may be confusing when interpreting industry-level data, it remains difficult to avoid such data when policy is examined at the country level. The data below, however, aggregates three potentially distinct levels-individual usage, firm-level adoption and frontier production-which can lead to misinterpretation.^[19] Although a country's population may have high individual access to consumer AI tools, its firm penetration may be weak and its diffusion throughout all levels of the firm can be limited, even with advanced local AI production. This suggests that the issue isn't simply the level of individual use, but its efficiency, in which countries this use case results in value capture through its translation into wider organizational adoption, in which countries value is captured through adaptation and localization, customization and downstream participation. Indeed, as highlighted earlier, the Brookings Institution points to this discrepancy and complexity, though its findings on how country comparative analysis is made difficult remain generally relevant to policy implications.

At the individual level use has already shown the great diversity across countries. In 2025 nearly 1/3rd of individuals from OECD countries report having used generative AI. Within Europe, this picture is even starker

showing a 48% use rate among Danes versus 18% in Romania and 47% in Malta and Estonia, versus 46% in Finland.^[20] Age disparity was also extremely significant showing a 53.6 percentage points gap across OECD countries, with 64% usage among 16-24 year-olds versus 7% among 65-74 year-olds in the EU. While these numbers help explain some trends in labor adoption through the channel of consumer use, they explain why a lack of demographic parity does not necessarily mean a lack of diffusion or adaptation to a certain task.

Business-level AI adoption still significantly lags, with variation according to nation and other factors. OECD data suggests only 20.2% of all firms report having used AI, where this percentage reaches over half (52.0%) for large firms and only 17.4% for small firms.^[21] Sectorally, IT has the highest percentage of users at 57.3%, but elsewhere this number has dropped, with Eurostat data for the EU showing only 19.95% of firms adopting AI in 2025, yet for Denmark this figure is 42.96%, whereas for Romania it is only 5.21%. Similarly, the World Bank analysis shows some interesting insights, with average general firm adoption of AI standing at just 8% in 2023 (across OECD), but with South Korea at 28% and Romania at just 2% (and many other variations in between). Clearly, not only individual use, but also the penetration rate among businesses remains extremely varied due to the country’s current technological position, business size and level of infrastructure sophistication.




Figure 7: European enterprise adoption remains highly uneven, with digitally mature economies far ahead of the lowest-adoption member states.

The core inequality in the system, however, stems from the upstream level of producing AI. AI production is concentrated in far fewer locations than even the usage. US firms invested over \$109 billion in AI in 2024, compared to only \$9.3 billion in China and \$4.5 billion in the UK and developed 40 models, compared to China’s 15. The World Bank continues to observe this imbalance with a growing number of the world’s largest users originating from middle-income countries such as India, Brazil, Indonesia and Vietnam, making up over 40% of ChatGPT traffic compared to less than 1% for low-income countries. More importantly, however, is that the production of large AI models remains concentrated still more than the consumption of AI tools,^[22] with the

US the frontier producer and only a handful of middle income countries appearing amongst those considered key producers of models with the World Bank listing Argentina, China and India.

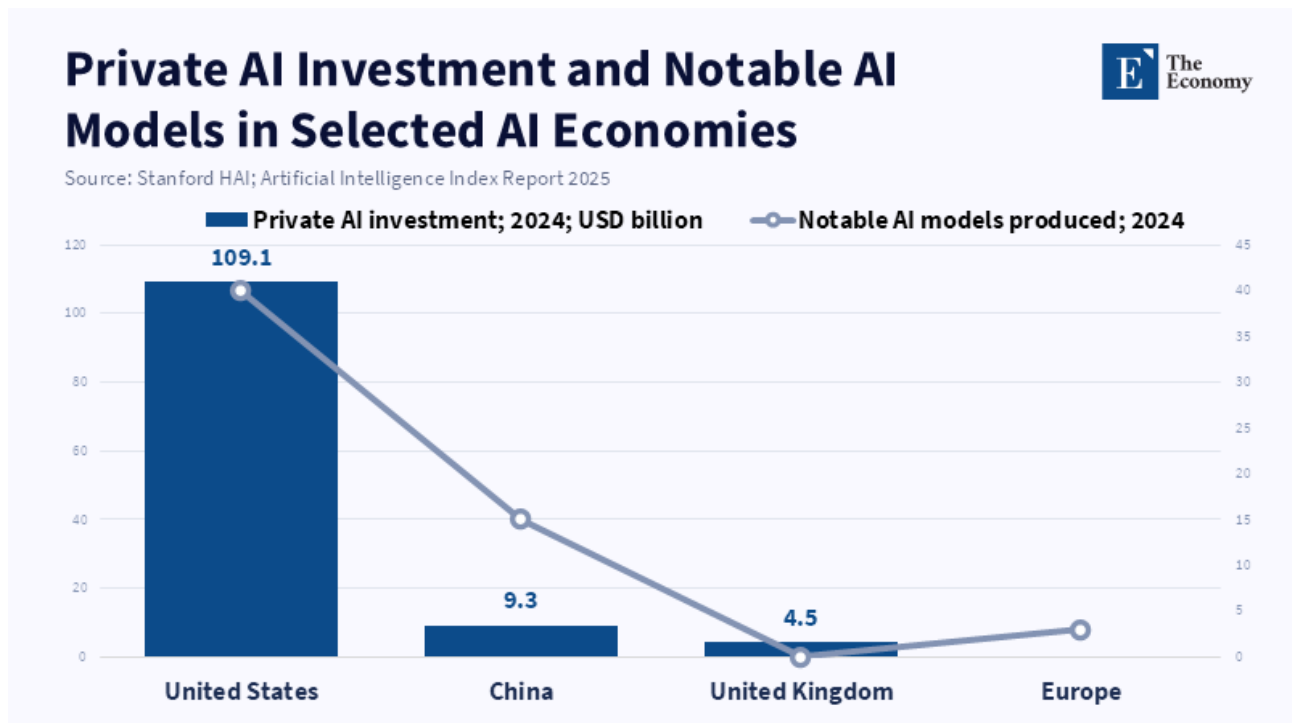


Figure 8: AI use is spreading more widely than frontier production, which remains concentrated in a small number of economies with large investment capacity.

As a result, it may be overly optimistic to assume that the disparity between countries and different strata of adoption will automatically be narrowed. While disparity in individual and demographic use, as well as in experience with new technologies, can eventually bring usage differences somewhat closer, disparity in work-related use and organization is largely driven by complementary resources rather than anything else. Complementary resources, as the World Bank defines them, are: compute, context, connectivity and competency.^[23] Applying this concept further, the OECD suggests a strong dependence of broad adoption on investment and on national leaders in the technology. A similar finding comes from the Federal Reserve Bank, where management quality and incentives explain a significant portion of the work-related AI adoption gap between Europe and the US, suggesting that a change in management quality of one standard deviation causes an increase of 9.6 percentage points in worker AI usage, clearly not something easily remediated.^[24]

It follows that for policy purposes, it is not simply enough to focus on individual use, or even firm-level adoption, but it is more effective to think of how these things translate into system-wide changes that have systemic advantages and implications for that nation as a whole. Developing countries can only truly reap the benefits of AI if they invest in capabilities not just for usage but for adaptation, localization and for the entire pipeline, including upstream production, to avoid premature deprofessionalization, which the World Bank suggests is a real concern. Ultimately, the country-level debate will always be focused on which nations are capitalizing on these trends for genuine economic growth.

5 Can Policy Boost AI Adoption? From Raising the Floor to Closing the Frontier Gap

One of the key drivers of whether to use policy intervention depends on the task ahead and whether policy has the potential to close important adoption gaps, thereby mitigating key hindrances to adoption. In fact, firms using AI are already stating that they are taking advantage of it in ways that the OECD report suggests policy needs to intervene more directly in; the study shows 83% of firms found that retraining and continuing education in use of AI was valuable, 82% found higher education and vocational training helpful and 78% indicated that they were helped by their improved IT infrastructure and the existence of competition within the market.^[25] Public provision of guidance, from reliable specialized sources, also helps, especially those with greater barriers, with the OECD recommending more readily available data to firms on matters of regulation, compliance and practical business applications. However, whether policy should be geared towards enhancing a specific group's technological access or artificial competitive advantage remains somewhat of an issue.

Public policy cannot help organizations create in-house capacity through delegation like leading firms already have, via their already-developed unique datasets, specialized custom workflows, experienced workforce and management capable of utilizing those unique proprietary datasets. OpenAI indicates that the productivity gap widens because of this in-house capability as top firms develop deeper forms of usage, beyond the increase in user counts.^[26] McKinsey's report confirms this by noting that the profits derived from these new technologies are for organizations capable of implementing business strategies for AI, dedicated teams focused on adoption and use, clear metric targets and job-specific training and trust. In other words, the key benefit is still derived internally by the firm itself. Firm-level (and even public institution) policy has also tended to focus less on how the system changes within firms and more on tracking user numbers and defining acceptable-use policies. These indicators need to become more concrete and measure work outcomes—the fraction of business processes currently performed with AI, total hours of AI use per role, the types of tasks that will be automated, amount of rework stemming from errors and creation of new jobs versus outsourcing;^[27] as shown by the variability found in Federal Reserve Bank research on the issue, performance metrics are extremely sensitive and depend entirely on context.

More critical views must be taken into account. AI can indeed create massive gains in productivity, as suggested by the NBER study showing a 14% improvement for those in customer support tasks, especially for less experienced workers^[28] and improved conditions for employees, customer support and the employer's satisfaction. It is also undeniable that AI adoption is growing rapidly, faster than the internet and personal computers did relative to their time of launch, as noted by the NBER. Lastly, it is reasonable to assume that country-level and inter-firm adoption gaps will close somewhat with decreasing prices, improved ubiquity and integration into everyday software.^[29] None of these arguments, however, makes up for the central weakness of these points; even in the most successful stories, AI adoption has been slow and steady, with two-year workplace usage comparable to PC usage in the first three years of introduction. Furthermore, evidence from the software development sector suggests that productivity improvements are not a given outcome of a simple increase in usage and depend heavily on task structure, worker experience and organization. However, even

though AI follows some common pathways as other technologies before it, there is something inherent about the capabilities that may separate it from them: the output. AI generates language, code and analytical answers that directly relate to cognitive tasks—this leads to the potential to transform the concept of delegation and work based on human judgment at all levels of the organization, with the risk that entry-level tasks, in particular, will be greatly diminished without comparable gain from newer jobs.

Given this background, policy interventions need to be more refined, sophisticated and realistic than simple calls for increased adoption or arguments about the inherent danger of this new technology. Policy can indeed play a vital role in establishing trustworthy frameworks and best practices, developing a skilled workforce that understands the value of AI, increasing overall investment, ensuring adequate competition between suppliers and building a better system of data sharing, alongside improved performance measures that better reflect reality rather than simplistic usage numbers. At the same time, policy cannot ignore that leading firms and countries have access to vast resources, have already built customized solutions and datasets for AI integration and have management teams with specific expertise to drive this forward. AI is beginning to become a feature primarily of the front-line, where those firms that are investing most are doing so to augment their own existing knowledge, not necessarily to replace their current workforces, or by integrating specific, customized applications. Governments cannot and should not, try to recreate this advantage purely from public policy; instead, the primary goal of policy should be to ensure that this new technological revolution does not widen, rather than narrow, existing global and national disparities.

6 Conclusion - From Contact to Transformation

The central mistake of the debate around AI adoption is conflating availability with transformation. Generative AI is already globally mainstream, but evidence suggests that organizational adoption remains uneven, patchy and heavily dependent on complementary assets. The key divide is not between firms, sectors, or countries that have and have not heard of AI; rather, between those that use it trivially and those that re-engineer their work processes, data infrastructure, managerial processes, and labor distribution in this regard. This matters, since the benefits of enhanced productivity, job gains and losses and greater competitiveness are derived from institutional reform, not from the isolated use of a chatbot.

This logic plays out across sectors and states alike: digitized, text-heavy, modular sectors are quick to adapt, while physical, regulated sectors do so more slowly. High-income, digitally mature economies are likely to convert use into profit more effectively; countries with less developed infrastructure, computation, management capacity and skills may be consigned to being downstream consumers of foreign AI applications. This policy can increase the adoption baseline through training, data provision, infrastructure investment, standards and guidance. It cannot replicate the idiosyncratic routines of leading firms. The policy focus should not be on measuring how many have engaged with AI, but on quantifying the extent to which work has been transformed, the distribution of profits, and the impact on inequality.^[30]

References

- [1] OECD (2026) *AI use by individuals surges across the OECD as adoption by firms continues to expand*. Paris: OECD.
- [2] Eurostat (2025) *20% of EU enterprises use AI technologies*. Luxembourg: Eurostat.
- [3] Eurostat (2025) *32.7% of EU people used generative AI tools in 2025*. Luxembourg: Eurostat.
- [4] Allen, J.S. (2026) 'Monitoring AI adoption in the U.S. economy', *FEDS Notes, Board of Governors of the Federal Reserve System*, 3 April.
- [5] U.S. Census Bureau (2025) *BTOS AI Core Question Updates*. Washington, DC: U.S. Census Bureau.
- [6] Brookings Institution (2026) *AI growth acceleration versus distributional fairness*. Washington, DC: Brookings Institution.
- [7] Stanford Institute for Human-Centered Artificial Intelligence (2025) *Artificial Intelligence Index Report 2025*. Stanford: Stanford University.
- [8] OpenAI (2026) *How frontier firms are pulling ahead: B2B Signals*. San Francisco: OpenAI.
- [9] McKinsey & Company (2025) *The State of AI in 2025: Agents, Innovation, and Transformation*. New York: McKinsey & Company.
- [10] Yotzov, I., Barrero, J.M., Bloom, N., Bunn, P., Davis, S.J., Foster, K.M., Jalca, A., Meyer, B.H., Mizen, P., Navarrete, M.A., Smietanka, P., Thwaites, G. and Wang, B.Z. (2026) *Firm Data on AI*. NBER Working Paper No. 34836. doi:10.3386/w34836.
- [11] OECD/BCG/INSEAD (2025) *The Adoption of Artificial Intelligence in Firms: New Evidence for Policy-making*. Paris: OECD Publishing. doi:10.1787/f9ef33c3-en.
- [12] Bick, A., Blandin, A. and Deming, D.J. (2025) *The Rapid Adoption of Generative AI*. NBER Working Paper No. 32966. doi:10.3386/w32966.
- [13] Federal Reserve Bank of St. Louis (2025) *The Impact of Generative AI on Work Productivity*. St. Louis: Federal Reserve Bank of St. Louis.
- [14] Brynjolfsson, E., Li, D. and Raymond, L.R. (2023) *Generative AI at Work*. NBER Working Paper No. 31161. doi:10.3386/w31161.
- [15] Cui, Z., Demirer, M., Jaffe, S., Musolff, L., Peng, S. and Salz, T. (2025) *The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers*. Microsoft Research.
- [16] Becker, J., Rush, N., Barnes, B. and Rein, D. (2025) *Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity*. METR.

- [17] World Bank (2025) *Artificial Intelligence and Development*. Washington, DC: World Bank.
- [18] International Monetary Fund (2026) *Bridging Skill Gaps for the Future: New Jobs Creation in the AI Age*. Washington, DC: IMF.
- [19] OECD (2025) *Generative AI and the SME Workforce: New Survey Evidence*. Paris: OECD Publishing. doi:10.1787/2d08b99d-en.
- [20] PwC (2025) *AI adoption could boost global GDP by an additional 15 percentage points by 2035 as global economy is reshaped*. London: PwC.
- [21] Economist Impact (2025) *Tracking AI's Economic Impact*. London: Economist Impact.
- [22] Eurostat (2025) *Use of artificial intelligence in enterprises. Statistics Explained*. Luxembourg: Eurostat.
- [23] Federal Reserve Bank of St. Louis (2026) 'Why does AI adoption differ so much across countries?', *On the Economy*. St. Louis: Federal Reserve Bank of St. Louis.
- [24] Bick, A., Blandin, A., Deming, D.J. and related authors (2026) *Cross-country evidence on AI adoption and management quality*. Federal Reserve Bank of St. Louis.
- [25] OECD (2025) *AI and the global productivity divide*. Paris: OECD.
- [26] GPTZero (2025) *AI adoption by industry*. GPTZero Research.
- [27] INSEAD (2025) *Mapping AI adoption by industry: What is your company's next move?* Fontainebleau: INSEAD Knowledge.
- [28] OpenWebUI (2026) *AI adoption leaderboard by country*. OpenWebUI.
- [29] Visual Capitalist (2025) *AI adoption rates by country*. Vancouver: Visual Capitalist.
- [30] Alice Labs (2026) *AI adoption by country 2026*. Alice Labs.