




BIS Working Papers
No 1343

The geography of AI firms

by Kumar Rishabh and Vatsala Shreeti

Monetary and Economic Department

April 2026



JEL classification: O33, C81, L86, F23, L16

Keywords: artificial intelligence, AI supply chain, firm geography, AI measurement

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS.

This publication is available on the BIS website (www.bis.org).

© *Bank for International Settlements 2026. All rights reserved. Brief excerpts may be reproduced or translated provided the source is stated.*

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

The geography of AI firms*

Kumar Rishabh[†] Vatsala Shreeti[‡]

April 15, 2026

Abstract

In this paper, we trace the geography and economic characteristics of firms that produce artificial intelligence (AI) products and services. Many economies around the world are evaluating their strategic priorities in AI, yet relatively little is documented about the global distribution of AI production. We construct a new database that identifies 1,246 AI-producing firms across 32 economies. We map these firms in each economy into the five layers of the AI supply chain: compute, cloud and related infrastructure, data tools, AI models and AI applications. The biggest markets for AI production are China and the US. Most economies specialise only in a few supply chain layers and many focus largely on compute. AI firms in all economies exhibit strong home bias in investment activity, with a focus on downstream applications. Finally, we find that venture capital inflows are strongly correlated with the presence and density of AI firms in a given economy.

Keywords: artificial intelligence, AI supply chain, firm geography, AI measurement

JEL codes: O33, C81, L86, F23, L16

*The authors thank Iñaki Aldasoro, Sebastian Doerr, Jon Frost, Gaston Gelos, Yi Huang and participants at the OECD competition division seminar for helpful comments. We also thank Adam Cap and Rudraksh Kansal for excellent research assistance. The views expressed here are those of the authors and not necessarily of the Bank for International Settlements.

[†]HEC Lausanne, University of Basel, Bank for International Settlements. Email: kumar.rishabh@unibas.ch.

[‡]Bank for International Settlements. Email: vatsala.shreeti@bis.org.

1 Introduction

Artificial intelligence (AI) is advancing rapidly and is being adopted across a wide range of sectors, from finance and healthcare to manufacturing and public administration. A growing body of research examines the implications of AI adoption for labour markets, productivity and macroeconomic outcomes.¹ However, not much is known about the firms that *produce* AI. Policymakers need to understand not only how AI is being used, but also where the capacity to develop and supply AI technologies resides and what the economic characteristics of AI-producing firms are. Many jurisdictions are increasingly evaluating their comparative advantage within the AI supply chain and the trade-offs associated with regulating AI firms and some are pursuing sovereign AI strategies. Yet a firm-level analysis of global AI production and its economic characteristics that can inform these discussions remains largely absent from the literature. As AI is becoming a strategic priority for governments around the world, this is an important research gap.

In light of this gap, in this paper, we build and provide a new dataset of 1,246 AI producer firms across 32 economies. We construct this dataset using a large language model (LLM)-assisted classification and extensive manual verification of data sourced from PitchBook. We use these data to identify the location of AI-producing firms and trace their economic weight in their respective economies over time. We then analyse the economic and institutional factors that are correlated with a greater density of AI-producing firms. We also look at the variation in each economy’s AI supply chain and document the sub-markets in which different economies specialise: i) compute, ii) cloud and related infrastructure, iii) data tools, iv) AI models and v) AI applications. Finally, we analyse the investment linkages between AI firms across different economies. The resulting map of AI firms and their economic and financial characteristics is not only useful for understanding AI production within each economy, but can also inform research seeking to quantify the cross-economy macroeconomic impact of AI.

Consistent with previous work, we find that the United States (US) and China are

¹See for example, [Bick, Blandin, and Deming \(2026\)](#); [Babina, Fedyk, He, and Hodson \(2024\)](#); [Brynjolfsson, Li, and Raymond \(2025\)](#); [Noy and Zhang \(2023\)](#); [Acemoglu \(2025\)](#); [Aldasoro, Doerr, Gambacorta, and Rees \(2024\)](#); [Cerutti, Garcia Pascual, Kido, Li, Melina, Mendes Tavares, and Wingender \(2025\)](#)

the biggest AI markets in terms of both the number and valuation of AI firms (Frost, Shreeti, and Rishabh (2026)). Beyond this, the focus of our analysis is on documenting new stylised facts. We provide four sets of results in line with our research questions.

First, we find that AI firm density (ie the presence and number of AI firms) is positively correlated with the size of the economy, share of research and development (R&D) expenditure in the economy and the size of venture capital (VC) inflows. VC inflows, in particular, are positively correlated both with the probability of a jurisdiction having an AI firm and the number of AI firms in the jurisdiction.

Second, the economic footprint of AI firms is especially large in the US, Korea and Chinese Taipei. Compared with other economies, AI firms in the US and Korea account for significantly higher shares of market capitalisation at 40% in the US and 39% in Korea in 2025. The same pattern holds for AI firms' share of total revenues and capital expenditure (capex). In 2024, AI firms accounted for around 13% of total revenue in both the US and Korea and 23% and 26% of capital expenditure respectively. This share was significantly higher than other economies. In Chinese Taipei, the combined market capitalisation of AI firms is more than twice its GDP.

Third, there is substantial heterogeneity across economies in terms of AI supply chain specialisation. Some economies like Chinese Taipei, Korea, Japan, the United Kingdom (UK), Switzerland and Hong Kong SAR focus largely on compute and others like the Euro area, Canada, Australia, Israel, Singapore and Sweden are focusing more on AI applications. Supply chain diversification in an economy is positively correlated with the extent of university-industry collaboration, size of VC inflows and its AI preparedness (as measured by the IMF AI preparedness index).

Fourth, AI firms in most economies exhibit home bias in investment activity. Notably, AI firms in the US and China make 64% and 74% of their investment deals with target companies based within their own jurisdiction. Beyond their home jurisdictions, AI firms in most economies (except China) also invest extensively in companies that are based in the US. Moreover, AI firms in all economies invest heavily in target firms operating in the AI applications sub-market, with at least 44% of all deals targeting such firms. We also

find that investment deals reinforce specialisation in many economies: in all economies except the US, India and Japan, the majority of investment deals made by AI firms target companies in the same layer of the supply chain that the investor operates in.

The remainder of the paper is organised as follows. Section 2 describes our firm classification methodology and the construction of the dataset. Section 3 explores the geography of AI firms and the economic and institutional factors associated with AI firms’ presence in an economy. Section 4 describes the economic weight of AI firms in different economies. Section 5 dives into differences in AI supply chain specialisation across economies and Section 6 focuses on the nature of AI firms’ domestic and cross-border investment activity. Section 7 concludes.

We also make our dataset and the data behind the figures publicly available, as a resource for other researchers and policymakers.

2 Firm classification and data description

2.1 Identifying AI firms

Most firm-level work on AI producers focuses on publicly listed companies in the US (e.g., [Babina et al., 2024](#); [Rishabh, Mihet, and Jang-Jaccard, 2025](#); [Mihet, Rishabh, and Gomes, 2026](#)). But a large part of the AI story involves private companies and AI firms beyond the US. To analyse both public and private AI firms globally, we begin with PitchBook, a commercial database with broad coverage across ownership types and geographies and treat its output as a seed set that we refine through systematic screening.

We identify AI firms globally in three steps.

Step 1: PitchBook vertical tags. The starting filter uses three PitchBook vertical tags: AI-machine learning (AI-ML), Big Data and Cloud and Development Operations (Cloud-DevOps). These tags are assigned by PitchBook’s research team based on firm descriptions, product offerings and industry activity. The AI-ML tag alone can miss firms operating in the upstream layers of the AI supply chain. Chip designers, cloud providers and data infrastructure companies often fall under Big Data or Cloud-DevOps rather

than AI-ML. Using all three tags leads to a more comprehensive set of firms.

Step 2: LLM-based screening. At the same time, the three vertical tags inevitably include firms that merely use AI products rather than produce them. To separate AI producers from users, an LLM classifier (GPT-5 Pro) reads each firm’s PitchBook description and classifies it as an AI supply chain producer, an AI user or unrelated to AI. We drop the firms classified as user and unrelated from the analysis.

Step 3: Manual verification. We conduct a manual review for each economy, over several rounds. This aims to resolve ambiguous cases where the LLM classifier lacked sufficient context and to add firms that the initial PitchBook filter missed.

We restrict the sample to firms with a valuation above \$500 million, using market capitalisation for public firms and the latest available private valuation from PitchBook.² The final dataset contains 1,246 AI producer firms across 32 economies, or 25 economies when Euro area member states are treated as a single entity.

2.2 Classifying firms along the AI supply chain

The selected AI firms play varied roles in the production of AI. A chip designer and a healthcare AI start-up both qualify as AI producers, but they sit at opposite ends of the supply chain. To capture this heterogeneity, we classify each firm into one of five supply chain layers following the method in [Rishabh and Shreeti \(2026\)](#). [Appendix B](#) details the classification procedure and [Gambacorta and Shreeti \(2025\)](#) provides further discussion of the layer definitions.

Compute includes firms that design, manufacture, or sell the core hardware for AI workloads: chips, circuit boards, servers and networking equipment. *Infrastructure* includes companies that provide data centre facilities and cloud services for hosting and scaling AI workloads. *Data tools* capture firms that provide data, labelling and reinforcement learning from human feedback (RLHF) services and the software tools used in training and deployment of AI models. *AI models* capture firms developing general-purpose machine-learning models. *AI applications* cover firms that integrate AI models

²Data quality for firms with lower valuations is relatively poor and as such we exclude them from our sample.

into end-user products in sectors such as healthcare, finance, legal, education and autonomous vehicles.

We assign a primary classification to each firm based on its core business activity. When a firm operates in more than one layer of the supply chain, we assign a secondary role to it as well. The classification follows a similar methodology as above: an LLM classifier processes the PitchBook firm description and manual review resolves disagreements and edge cases.³

2.3 Other data

Public firm financials: We obtain annual firm-level financial data from S&P Global Market Intelligence for publicly listed companies in jurisdictions that host AI firms, covering 2010 through 2024. Market capitalisation data are extended through 2025 using end-of-year pricing. The panel is restricted to ultimate parent companies to avoid double-counting across subsidiaries. We match our AI firm sample to this panel using a manual crosswalk. For each jurisdiction and year, we compute the AI share of key financial aggregates (revenue, capital expenditure and market capitalisation) as the sum across AI firms divided by the sum across all firms.

Economy-level AI activity variables: We construct several economy-level variables on AI activity using the firm-level data. To measure the extensive margin of an economy's AI firm activity, we build a binary indicator variable which equals one when a given jurisdiction hosts at least one AI firm. We measure the intensive margin of AI firm activity by the total count of AI firms and the aggregate valuation of AI firms in a given economy normalised by nominal GDP.

Supply chain diversification: For the supply chain diversification analysis, we compute a diversification index: one minus the Herfindahl-Hirschman index (HHI) of valu-

³Some subjectivity with the classification is unavoidable. For instance, Indian IT services firms such as Wipro and HCL Technologies are classified into the cloud and related infrastructure layer. Their business is closer to systems integration and cloud consulting than to building or operating cloud platforms. Within the five-layer taxonomy, infrastructure is the closest fit given these firms' cloud services revenue, though interpretation can differ.

ation shares across the five primary-role layers. This index ranges from zero, when all value is concentrated in a single layer, to 0.8 for a perfectly even distribution across the five categories.

Macroeconomic indicators: We match the firm-level data to a cross-section of jurisdiction characteristics drawn from five external sources. Macroeconomic indicators come from the World Bank’s World Development Indicators: GDP in current US dollars, R&D expenditure as a share of GDP, stock market capitalisation as a share of GDP and the manufacturing and services sectors’ shares of value added. For each indicator, we use data from the most recent year available. The World Bank data do not have information on Chinese Taipei; we extract its GDP from the IMF World Economic Outlook. We also use data on private credit extended by banks (as a share of GDP), which comes from the World Bank’s Global Financial Development Database.

AI preparedness: The IMF’s AI Preparedness Index (AIPI) provides a composite readiness score for 174 economies on a zero-to-one scale, constructed from sub-indicators spanning digital infrastructure, human capital, innovation capacity and the regulatory environment. A value closer to 1 indicates a higher degree of AI preparedness.⁴

Venture capital intensity and research collaboration: We use the World Intellectual Property Organisation (WIPO) Global Innovation Index 2025 for two indicator-level scores, each normalised to a 0-to-100 scale. The first captures VC deals received per billion US dollars of GDP by purchasing power parity (PPP), which proxies for the depth of the risk capital ecosystem. The second captures university–industry R&D collaboration and is a survey-based measure drawn from the World Economic Forum’s Executive Opinion Survey, where respondents rate the extent to which businesses and universities collaborate on R&D in their jurisdiction.

⁴See [Cazzaniga, Jaumotte, Li, Melina, Panton, Pizzinelli, Rockall, and Tavares \(2024\)](#) for the methodology underlying the AIPI and its application to AI readiness measurement across economies.

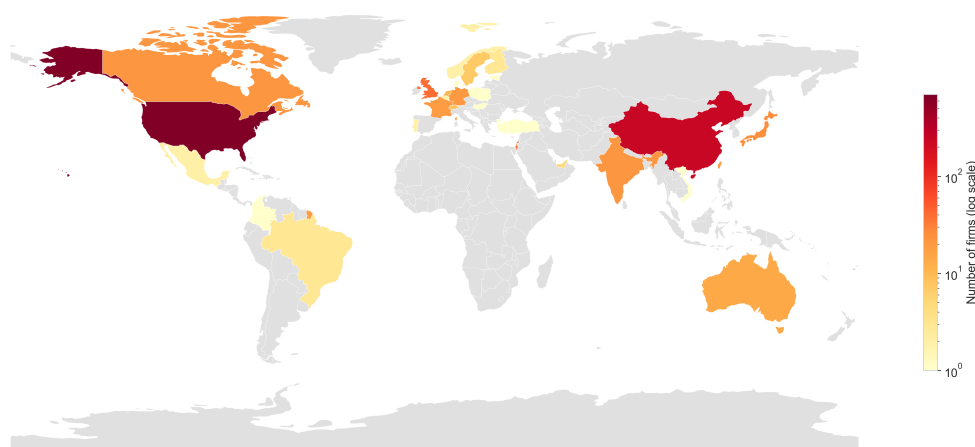
Unit of observation: The unit of observation for the macroeconomic data, the AI preparedness index and the WIPO data is a jurisdiction, with one exception: the 21 euro area member states enter as a single entity. Where an external source reports a pre-built Euro area aggregate (as in the World Development Indicators), we use it directly. Where it does not (the AIPI and the Global Innovation Index), we construct the aggregate as a GDP-weighted average of member state scores. The sample of macroeconomic data retains all economies with GDP above \$50 billion. The final sample contains 76 economies, with 32 having at least one AI firm.

3 Where are AI firms located?

The US and China have the largest number of AI firms.

Of the 1,246 AI firms across 32 economies, nearly 700 are based in the US and around 250 in China. The next biggest hubs by firm count include the Euro area (59 firms), Israel (42), the United Kingdom (40), Japan (26), India and Canada (22 each), Chinese Taipei (20) and Korea (15). Figure 1 provides a map of AI firms across the world by the number of AI firms in each jurisdiction in 2025. Figure A1 in the appendix shows the top 15 jurisdictions by AI firm count.

Figure 1: Global AI firms: count by jurisdiction in 2025

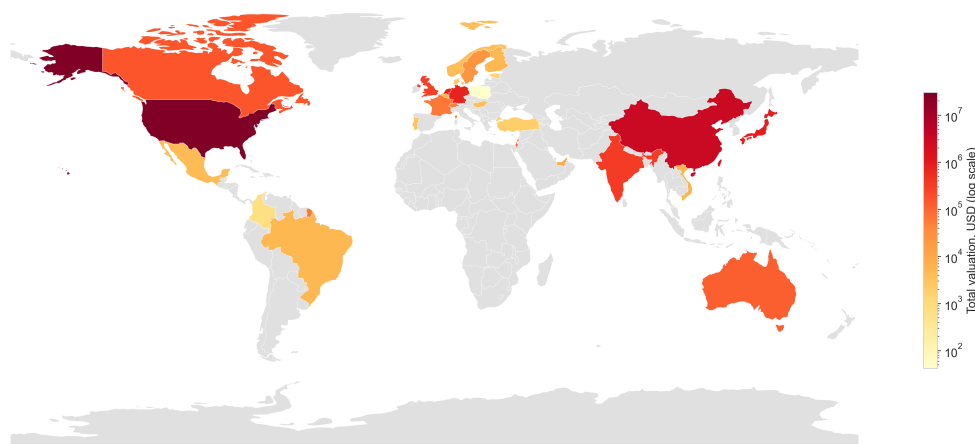


Notes: Colour scale is logarithmic. Jurisdictions with no AI firms in the dataset appear in grey. Euro area members appear individually. The use of this map does not constitute, and should not be construed as constituting, an expression of a position by the BIS regarding the legal status of, or sovereignty of any territory or its authorities, to the delimitation of international frontiers and boundaries and/or to the name and designation of any territory, city or area.

The most highly valued AI firms are located in the US and China.

The top of the geographic distribution of AI firms in terms of firm valuation looks similar to the distribution by firm count, with the US, China and the Euro area having the most highly valued AI firms (Figure 2). This is followed by Chinese Taipei and Korea, even though they have a smaller number of AI firms compared to other jurisdictions. The next biggest markets in terms of valuation are Japan, India, United Kingdom, Canada and Australia (Figure A2).

Figure 2: Global AI firms: aggregate valuation by jurisdiction in 2025



Notes: Colour scale is logarithmic. Valuation is latest market capitalisation for public firms and latest estimated valuation for private firms. Jurisdictions with no AI firms in the dataset appear in grey. Euro area members appear individually. The use of this map does not constitute, and should not be construed as constituting, an expression of a position by the BIS regarding the legal status of, or sovereignty of any territory or its authorities, to the delimitation of international frontiers and boundaries and/or to the name and designation of any territory, city or area.

3.1 Institutional and economic factors correlated with greater regional AI firm activity

To shed light on institutional and economic factors associated with greater AI firm activity, we estimate two simple regressions on AI firm density.

Bigger economies with larger VC inflows are more likely to have AI firms.

First, we focus on the extensive margin of AI firm activity (columns 1-4 of Table 1). Our sample consists of a cross-section of jurisdictions and the unit of observation is a jurisdiction. The dependent variable in this case is an indicator which equals 1 if there

is at least one AI firm in that jurisdiction. The main independent variables include i) the size of the economy, measured by the logarithm of its GDP, ii) the jurisdiction’s AI preparedness index, which is a composite of its quality of digital infrastructure, human capital, innovation and regulatory readiness, iii) the R&D expenditure as a share of the GDP and iv) the VC received per billion US dollars of GDP (measured in PPP).

Our estimates show that the probability of an economy having at least 1 AI firm is positively correlated with the size of the economy and the amount of VC funding the economy receives (column 4 of Table 1).

Table 1: Factors associated with AI firm location

| | AI firm presence (Probit) | | | | AI firm count (PPML) | | | |
|-----------------------|---------------------------|---------------------|---------------------|---------------------|----------------------|---------------------|---------------------|---------------------|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| Log(GDP) | 0.917*** (0.265) | 0.902*** (0.252) | 1.150*** (0.268) | 1.079*** (0.283) | 1.213*** (0.137) | 1.269*** (0.095) | 1.352*** (0.078) | 1.282*** (0.063) |
| AI preparedness index | 7.366*** (1.909) | | | | 7.289*** (2.782) | | | |
| R&D / GDP | | 1.226*** (0.403) | | 0.423 (0.599) | | 0.630*** (0.149) | | 0.343** (0.155) |
| VC received | | | 0.070*** (0.024) | 0.057* (0.030) | | | 0.030*** (0.009) | 0.018** (0.009) |
| Observations | 74 | 72 | 71 | 68 | 74 | 72 | 71 | 68 |
| Pseudo R^2 | 0.610 | 0.635 | 0.670 | 0.668 | 0.933 | 0.946 | 0.948 | 0.955 |

Notes: Robust standard errors in parentheses. Constants not reported. Columns (1)–(4) estimate a probit where the dependent variable equals one if the economy hosts at least one AI producer firm. Columns (5)–(8) model the count of AI firms using Poisson pseudo-maximum likelihood. All specifications control for log GDP in current US dollars. AIPI overall is the IMF AI Preparedness Index (0–1), a composite of digital infrastructure, human capital, innovation, and regulatory readiness. R&D / GDP is gross R&D expenditure as a share of GDP (World Bank). VC received is venture capital deals received per billion PPP GDP (WIPO GII 2025, scored 0–100). The Euro area enters as a single observation. Columns (4) and (8) report the preferred specification. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Turning to the intensive margin, in the second regression, the dependent variable is the AI firm count in a given jurisdiction (columns 5-8 in Table 1). The independent variables remain the same as before. Since many economies have no AI firms, we estimate the regression using Poisson pseudo-maximum likelihood (PPML), which accounts for the dependent variable having many zeroes.

Our estimates show that an economy with a greater number of AI firms also has higher GDP, higher share of R&D expenditure and higher inflows of venture capital.

These findings are intuitive, as a flourishing AI ecosystem rests both on active R&D and the availability of financing for early-stage, high-risk AI firms.

VC inflows are correlated with the presence of AI firms but bank credit is not.

Table 2 looks deeper into the financing of AI firms across economies with the same dependent variables but a different set of independent variables. As above, columns 1-4 focus on the extensive margin and columns 5-8 focus on the intensive margin of AI activity in a given jurisdiction. The independent variables now include the size of the economy, the ratio of bank credit to GDP, the VC inflows and the ratio of total stock market capitalisation to GDP.

Our results show that at the extensive margin, the probability that an economy has an AI firm is positively correlated with its size and its VC inflows. Bank credit and stock market capitalisation are not significantly correlated with this probability in our sample. On the intensive margin, in addition to the size of the economy, the proportion of bank credit and venture capital inflows are positively correlated with the number of AI firms in a given economy. Notably, the magnitude of the correlation is much higher for VC inflows than bank credit.

Given the rapid pace of technological development in AI and the presence of young firms, it is intuitive that larger pools of venture capital are able to support the AI ecosystem to a greater extent than bank credit. Bank lending is collateral-based and structured around predictable cash flows. Many firms active in AI are young (Figure A3) and may lack the collateral or predictability of cash flows needed to obtain bank credit. VC has higher tolerance for the risk profile of young AI firms and hence can fill this gap. As opposed to banks, venture capital may also be willing to finance firms that have a higher share of intangible capital, as might be the case for many AI firms (outside of compute and infrastructure).

Table 2: Financial structure and AI firm location

| | AI firm presence (Probit) | | | | AI firm count (PPML) | | | |
|---------------------|---------------------------|---------------------|---------------------|---------------------|----------------------|---------------------|---------------------|---------------------|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| Log(GDP) | 1.040*** (0.269) | 1.150*** (0.268) | 1.018*** (0.234) | 1.114*** (0.281) | 1.468*** (0.140) | 1.352*** (0.078) | 1.422*** (0.120) | 1.410*** (0.099) |
| Bank credit / GDP | 0.015*** (0.005) | | | 0.004 (0.008) | -0.001 (0.002) | | | 0.005* (0.003) |
| VC received | | 0.070*** (0.024) | | 0.060** (0.026) | | 0.030*** (0.009) | | 0.033*** (0.009) |
| Stock mkt cap / GDP | | | 0.002 (0.002) | 0.000 (0.002) | | | 0.003*** (0.001) | 0.001 (0.001) |
| Observations | 75 | 71 | 64 | 64 | 75 | 71 | 64 | 64 |
| Pseudo R^2 | 0.589 | 0.670 | 0.454 | 0.654 | 0.897 | 0.948 | 0.908 | 0.953 |


Notes: Robust standard errors in parentheses. Constants not reported. Columns (1)–(4) estimate a probit where the dependent variable equals one if the economy hosts at least one AI producer firm. Columns (5)–(8) model the count of AI firms using Poisson pseudo-maximum likelihood. All specifications include log GDP in current US dollars. Bank credit / GDP is private credit extended by deposit money banks as a share of GDP (World Bank GFDD). VC received is venture capital deals received per billion PPP GDP (WIPO GII 2025, scored 0–100). Stock mkt cap / GDP is stock market capitalisation as a share of GDP (World Bank WDI). The Euro area enters as a single observation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

4 How has the economic weight of AI firms evolved in their jurisdictions?

AI firms’ valuation as a share of GDP is the highest in Chinese Taipei.

The combined market valuation of AI firms in Chinese Taipei was more than twice its GDP in 2025. The valuation share of AI firms was also high in the US and Korea, with the combined market valuation of these firms roughly equal to GDP and about half of GDP respectively (Figure 3). Japan, Israel and China follow at a distance on this measure.

Figure 3: AI firm valuation as a share of 2025 GDP




Notes: Total AI firm valuation expressed as a percentage of nominal GDP for each jurisdiction. Top 15 economies selected by aggregate AI firm valuation.

AI firms in the US and Korea have seen the biggest market cap gains.

Compared to other economies, AI firms in the US and Korea stand out in terms of the proportion of total market capitalisation they account for (at 40% and 39% in 2025 respectively) as well as its growth since 2010 (Figure 4). AI firms in other economies account for less than 15% of total market cap, although this share has been rising in China, Israel and the Euro area. We cannot show Chinese Taipei in Figure 4 since the denominator capturing the total market capitalisation is not available.

Figure 4: AI firms’ share of total market capitalisation by jurisdiction, 2010–25




Notes: Share is computed as the sum of end-of-year market capitalisation across AI firms divided by the sum across all public firms in each jurisdiction-year. Euro area is treated as a single economic bloc. “Other” reports the unweighted average of jurisdiction-level shares. Market capitalisation data from S&P Global Market Intelligence end-of-year file.

AI firms’ share of capex and total revenues is the highest in the US and Korea.

As with market cap, AI firms in the US and Korea also account for higher proportions of total revenue and capital expenditure (capex) relative to other economies (Figure 5 and Figure 6). In 2024, AI firms accounted for around 13% of total revenue in both the US and Korea, but 23% and 26% of capex respectively. AI firms in the US stand out, even compared to Korea, in terms of the rapid increase in the proportion of total revenues and capex they account for between 2010 and 2024.


Figure 5: AI firms' share of total revenue by jurisdiction, 2010–24



Notes: Share is computed as the sum of revenue across AI firms divided by the sum across all public firms in each jurisdiction-year. Euro area is treated as a single economic bloc. “Other” reports the unweighted average of jurisdiction-level shares. Revenue data from S&P Global Market Intelligence.

AI firms' share in total revenues was less than 6% in China, the Euro area, India, Japan and Israel in 2024. AI firms' share of capex in economies other than Korea and the US was less than 5% in 2024 and has stayed relatively stable over time.

Figure 6: AI firms' share of total capital expenditure by jurisdiction, 2010–24




Notes: Share is computed as the sum of capital expenditure across AI firms divided by the sum across all public firms in each jurisdiction-year. Euro area is treated as a single economic bloc. “Other” reports the unweighted average of jurisdiction-level shares. Capex data from S&P Global Market Intelligence.

Price-to-earnings ratios of AI firms have been increasing disproportionately more than non-AI firms.

Figure 7 depicts the median price-to-earnings ratio of AI firms and non-AI firms across economies in two time periods: 2010 to 2016 and 2017 to 2024. The median price-to-earnings ratio of AI firms is higher than that of non-AI firms in all economies in our analysis, suggesting higher expected growth of AI firms compared with non-AI firms. Moreover, this ratio for AI firms has increased over time for the US, China, the Euro area, Japan, UK and India but not for Korea.

Figure 7: Median price-to-earnings ratio: AI vs non-AI firms by jurisdiction, 2010–16 and 2017–24



Notes: Each jurisdiction shows four points: AI and non-AI medians for two periods (2010–2016 and 2017–2024). Arrows connect the same group across periods, pointing from the earlier to the later period. The price-to-earnings ratio is computed at the firm level as market capitalisation divided by net income. Data from S&P Global Market Intelligence.

5 Regional differences in AI supply chain specialisation


The AI supply chain consists of five sub-markets that include i) compute, ii) cloud and related infrastructure, iii) data and related tools, iv) AI models and v) AI applications (Gambacorta and Shreeti (2025)). Economies differ in their specialisation in AI, with some having AI firms active in all sub-markets and others with more specialised firms

dedicated to a few sub-markets.

Some economies specialise only in compute, others are more diverse in terms of AI firm count

In terms of the number of AI firms, economies like Chinese Taipei, Australia, Switzerland, Hong Kong SAR and Sweden specialise in at most two AI sub-markets (Figure 8). Other economies like the US, China, the Euro area, Japan and the UK have active AI firms in every sub-market, though the number of firms varies substantially across the supply chain within each economy. Chinese Taipei and Korea stand out in terms of the high share of AI firms active in the market for compute.

Figure 8: AI supply chain layer composition by jurisdiction, by firm count in 2025




Notes: Vertical 100% stacked bar chart showing the share of each supply chain layer within each jurisdiction, by firm count. Top 15 jurisdictions selected by aggregate AI firm valuation and ordered left to right by decreasing total valuation. Euro area members are pooled into a single observation.

In terms of AI firm valuations, most economies are not very diversified

Market valuation of AI firms is not evenly distributed across different sub-markets, even in economies which are diverse in terms of AI firm counts across sub-markets (Figure 9). In Chinese Taipei, Korea, Japan, United Kingdom, Switzerland and Hong Kong SAR, most of the AI valuation resides in the market for compute. In India, AI firms in the cloud and other infrastructure markets contribute the most to the total valuation of AI

firms. On the other hand, in China, the Euro area, the UK, Canada, Australia, Israel, Singapore and Sweden, most of the AI valuation lies in the front-end AI applications market. In the US and to a lesser extent in China, AI valuation is more diversified across sub-markets.

Figure 9: AI supply chain layer composition by jurisdiction, by valuation in 2025



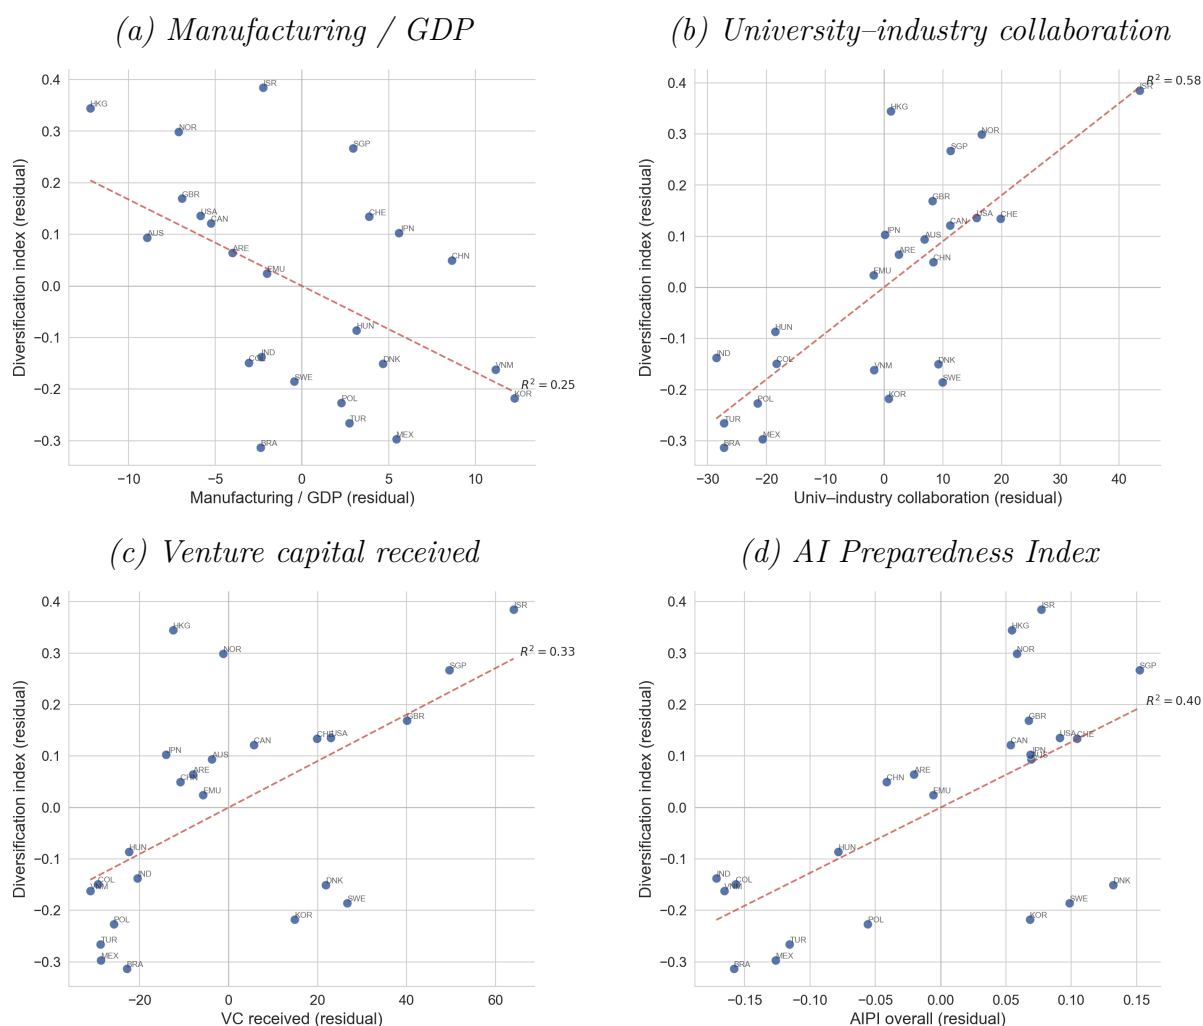
Notes: Vertical 100% stacked bar chart showing the share of each supply chain layer within each jurisdiction, by valuation as reported by PitchBook. Top 15 jurisdictions selected by aggregate AI firm valuation and ordered left to right by decreasing total valuation. Euro area members are pooled into a single observation.

Diversification is positively correlated with university-industry collaboration, venture capital inflows and AI preparedness.

Figure 10 focuses on economies' supply chain diversification. We compute an AI diversification index based on the valuation of AI firms across the five sub-markets. The index is defined as one minus the sum of squared valuation shares of AI firms across the five sub-markets. The index equals 0 if the totality of AI firms' valuation accrues to one sub-market. The index equals 0.8 if the valuation of AI firms is evenly distributed across the five sub-markets.

We compute partial correlations between the value of the diversification index in a given economy and i) the manufacturing share of GDP in the economy, ii) the extent of university-industry collaboration in the economy, iii) the amount of venture capital received in the economy and iv) the economy's AI preparedness index. In each case, the correlations control for the size of the economy.

Figure 10: Supply chain diversification and economy characteristics: partial correlations



Notes: Each panel shows the partial correlation between the diversification index and the labelled variable after residualising both on log GDP, effectively controlling for economy size. The diversification index is one minus the Herfindahl-Hirschman index of valuation shares across five primary-role supply chain layers, ranging from 0 (fully concentrated) to 0.8 (even distribution). The dashed line is a linear fit through the residuals. Sample restricted to economies with at least 1 AI firm. Manufacturing share from World Bank WDI (latest available year). University–industry collaboration and VC received from WIPO GII 2025 (scored 0–100). AII from IMF (2023). The Euro area enters as a single observation.

Economies which have extensive university–industry collaboration, higher venture capital inflows and rank higher on the AI preparedness index also have more diversified AI supply chains. Strong knowledge spillovers between universities and industries may help break self-reinforcing specialisations, for example, an economy with a strong service sector gravitates towards AI applications absent these linkages. Moreover, different sub-markets in the AI supply chain have different capital requirements and risk profiles and vibrant venture capital ecosystems may be suited to support them. Together, these findings suggest two potentially complementary mechanisms. University–industry collaboration may diversify the knowledge base in AI and VC availability may diversify the financing base.

We also find a negative correlation (conditional on market size) between an economy's AI supply chain diversification and the share of the manufacturing sector in the economy. A stronger manufacturing sector may pull an economy's AI sector towards compute and infrastructure, which are hardware-intensive markets and away from diversifying into more downstream AI markets. This pattern plays out in Chinese Taipei and South Korea, which are both manufacturing-heavy economies and have an outsized share of their AI firm valuation in the compute layer.

6 What is the nature of AI firms' investment activity?

Most economies have a home bias in AI investment activity.

Figure 11 shows a heatmap of AI firms' cross-border investment activity. The Y axis depicts the jurisdiction that the investing AI firm belongs to and the X axis depicts the jurisdiction of the target company the AI firm invests in. The diagonal elements of the heatmap reflect the proportion of investment deals where the investor and target are in the same jurisdiction. AI firms with more than 30% of their investment in their home economy are said to have significant home bias.

Figure 11: Geographic deal flows: investor jurisdiction to target jurisdiction, 2010–2025

| Investor region | United States | China | Euro area | Japan | Korea | United Kingdom | Canada | India | Sweden | Other |
|-----------------|---------------|-------|-----------|-------|-------|----------------|--------|-------|--------|-------|
| United States | 63.9 | 1.3 | 8.0 | 0.9 | 0.5 | 6.2 | 3.5 | 3.8 | 0.6 | 11.1 |
| China | 7.8 | 74.4 | 2.3 | 0.6 | 0.8 | 2.0 | 0.4 | 1.9 | 0.5 | 9.3 |
| Euro area | 29.6 | 2.1 | 39.0 | 1.4 | 0.2 | 5.9 | 2.3 | 2.2 | 1.2 | 16.0 |
| Japan | 23.1 | 1.0 | 9.6 | 46.5 | 0.6 | 2.1 | 1.4 | 2.9 | 0.3 | 12.5 |
| Korea | 34.8 | 3.0 | 6.3 | 1.6 | 37.6 | 2.0 | 0.8 | 2.2 | 0.6 | 11.0 |
| United Kingdom | 30.1 | 2.0 | 9.8 | 0.0 | 0.6 | 43.8 | 1.4 | 0.6 | 0.6 | 11.2 |
| Canada | 47.8 | 0.5 | 10.0 | 1.3 | 0.5 | 10.0 | 15.2 | 1.8 | 1.0 | 11.8 |
| India | 41.7 | 0.0 | 8.0 | 1.7 | 0.0 | 5.2 | 0.9 | 27.9 | 0.6 | 14.1 |
| Sweden | 22.2 | 0.0 | 35.6 | 0.0 | 0.0 | 8.9 | 0.0 | 0.0 | 22.2 | 11.1 |
| Other | 26.4 | 18.2 | 8.1 | 2.7 | 0.6 | 2.6 | 1.5 | 1.3 | 0.6 | 37.9 |

Notes: Cell values are normalised: each investor jurisdiction’s row sums to 100%. Nine named jurisdictions shown individually; all others grouped as Other. The diagonal captures domestic investment; off-diagonal cells capture cross-border flows. A cross-border deal is defined as one where the investor’s jurisdiction differs from the target’s jurisdiction. Euro area member countries are pooled into a single observation on both the investor and target side.

Home bias in investment activity is the strongest in the US and China, but also sizeable in Japan, the UK, the Euro area and Korea. The only economies that do not show significant home bias are Canada, India and Sweden. AI firms in all economies except China invest extensively in US-based target companies.

AI firms in most economies invest heavily in the applications sub-market.

Most economies are investing heavily in the AI applications sub-market. Figure 12 shows the distribution of AI firms’ investment deals across different AI sub-markets in different jurisdictions. At least 44% of the deals in all economies targeted AI applications firms. AI firms in Korea, Japan, China, the UK and the Euro area are also active in deal-making in the compute sub-market, with compute target firms accounting for at least 11% of the deals.

Figure 12: Target layer composition by investor jurisdiction, 2010–2025




Notes: Heatmap showing, for each investor jurisdiction, the share of deals directed at each target supply chain layer. Cell values are row-normalised: each investor jurisdiction’s row sums to 100%. Both domestic and cross-border deals are included. Targets without a classified supply chain layer are grouped as Other/Unclassified.

Investment deals reinforce specialisation in many economies.

In Figure 13, we calculate the share of deals where the investor’s supply chain layer matches its targets’ supply chain layer. In all economies except the US, India and Japan, the majority of investment deals made by AI firms are in the same layer of the supply chain that they operate in. For example, in Sweden, over 90% of AI firms’ deals target the same layer as the one they operate in. For Canada and China, this number is between 70% to 80%. On the other hand, in India, Japan and the US, AI firms invest in a relatively broader range of AI sub-markets.

Figure 13: Share of deals in same layer as investor, by jurisdiction, 2010–2025



Notes: For each investor jurisdiction, the bar shows the share of deals where the investor’s primary supply chain layer matches the target’s classified supply chain layer. Higher values indicate that a jurisdiction’s AI firms predominantly invest in targets operating in the same part of the supply chain. The dashed vertical line marks 50%.

In principle, firms can use investment deals to acquire capabilities that they lack and diversify their operations in adjacent markets. However, we observe that in most jurisdictions, investment activity is concentrated within the same AI sub-market, pointing towards the role of informational advantages over the incentive to diversify.

7 Conclusion

In this paper, we construct a novel firm-level dataset of 1,246 AI producer firms across 32 economies and map them to the five sub-markets within the AI supply chain including compute, cloud and related infrastructure, data tools, AI models and AI applications. By combining LLM-assisted classification with extensive manual verification and drawing on PitchBook, S&P Global Market Intelligence and several cross-economy indicator databases, we provide a comprehensive overview of where AI is produced, how the economic weight of AI firms has evolved and what institutional and financial conditions are

associated with greater AI activity. Our findings shed light on topical policy debates, particularly those surrounding sovereign AI and allow jurisdictions to assess where their economies sit in the global AI supply chain.

The dataset and analysis also open several avenues for further research. A natural extension would be to link firm-level AI production data to measures of AI adoption, enabling researchers to study how the geography of AI supply interacts with the geography of AI demand and, ultimately, with productivity and labour market outcomes. The economy-level data on AI firm density can also be useful for other cross-country studies where measuring regional AI exposure is important.


References

- Acemoglu, Daron.** 2025. “The simple macroeconomics of AI.” *Economic Policy* 40 (121): 13–58.
- Aldasoro, Iñaki, Sebastian Doerr, Leonardo Gambacorta, and Daniel Rees.** 2024. “The impact of artificial intelligence on output and inflation.” *BIS Working Papers* (1179): .
- Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson.** 2024. “Artificial intelligence, firm growth, and product innovation.” *Journal of Financial Economics* 151 103745.
- Bick, Alexander, Adam Blandin, and David J Deming.** 2026. “The rapid adoption of generative AI.” *Management Science*.
- Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond.** 2025. “Generative AI at work.” *The Quarterly Journal of Economics* 140 (2): 889–942.
- Cazzaniga, Mauro, Florence Jaumotte, Longji Li, Giovanni Melina, Augustus J. Panton, Carlo Pizzinelli, Emma 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.
- Cerutti, Eugenio M., Antonio I. Garcia Pascual, Yosuke Kido, Longji Li, Giovanni Melina, Marina Mendes Tavares, and Philippe Wingender.** 2025. “The Global Impact of AI: Mind the Gap.” IMF Working Paper 2025/076, International Monetary Fund.
- Frost, Jon, Vatsala Shreeti, and Kumar Rishabh.** 2026. “Global giants in the AI supply chain.” BIS Bulletin 122, Bank for International Settlements.

- Gambacorta, Leonardo, and Vatsala Shreeti.** 2025. “The AI supply chain.” *BIS Papers*.
- Mihet, Roxana, Kumar Rishabh, and Orlando Gomes.** 2026. “Is it AI or data that drives firm market power?” *Journal of Monetary Economics* 157 103878. <https://doi.org/10.1016/j.jmoneco.2025.103878>.
- Noy, Shakked, and Whitney Zhang.** 2023. “Experimental evidence on the productivity effects of generative artificial intelligence.” *Science* 381 (6654): 187–192.
- Rishabh, Kumar, Roxana Mihet, and Julian Jang-Jaccard.** 2025. “Cyberrisk and AI Firms.” *Swiss Finance Institute Research Paper* (25-39): .
- Rishabh, Kumar, and Vatsala Shreeti.** 2026. “Market Structure and Firm Outcomes in Artificial Intelligence.” Mimeo.


A Appendix: Additional Figures and Tables

Figure A1: AI firm count in 2025 by jurisdiction, public and private




Notes: Top 15 jurisdictions selected by aggregate AI firm valuation. Euro area members are pooled into a single observation.

Figure A2: Aggregate AI firm valuation in 2025 by jurisdiction, public and private firm



Notes: Stacked horizontal bar chart of aggregate AI firm valuation (USD), split by public and private status. Horizontal axis is log-scaled. Top 15 jurisdictions selected by aggregate AI firm valuation. Euro area members are pooled into a single observation. Valuation is market capitalisation for public firms and latest estimated valuation for private firms.

Figure A3: Share of young AI firms by jurisdiction, valuation and firm count in 2025



Notes: Grouped horizontal bar chart showing the share of young firms (founded after 2012) within each jurisdiction, measured by valuation and by firm count. Top 15 jurisdictions selected by aggregate AI firm valuation. Jurisdictions are sorted by the valuation-based young share. Euro area members are pooled into a single observation.

B Appendix: LLM prompt for supply chain classification

Each firm in the PitchBook sample was classified into an AI supply chain layer using OpenAI’s GPT-5 Pro model (released August 2025) with temperature set to zero for deterministic output. The model received one firm at a time, with inputs limited to the short business description and keyword tags available in PitchBook. No external information was supplied. The prompt below was designed to separate the classification decision (which depends on what a firm does) from the AI flag decision (which depends on whether the description explicitly mentions AI). All outputs were returned as structured JSON and parsed programmatically. A random subset of classifications was reviewed manually. The prompt used for the purpose is given below.

You will receive ONE company description as `text_input` containing:

- "DESC:" (short description)
- "KW:" (keywords)

CRITICAL PRINCIPLE (repeat + follow strictly):

- 1) Role classification is NOT dependent on AI language.
 You MUST classify the company into an AI supply-chain role even when AI is not mentioned. This includes firms like semiconductors, OSAT packaging/test, chip manufacturing services, networking, data centers, cloud, hosting, etc. Even if the text is "old-fashioned" and never says "AI", still map the business to the supply-chain role based on what it does.
- 2) AI flags ARE dependent on explicit AI language.
 Be conservative: set AI flags to 1 only if AI/ML is explicitly stated in the text_input.

TASK

From text_input ONLY (no outside knowledge), output ONE JSON object with:

- role_primary: exactly one label from the allowed set
- role_secondary: either a different allowed label, or null
- ai_producer_flag: 1 only if text explicitly indicates the firm sells/provides AI/ML products/services/models to customers; else 0
- ai_user_flag: 1 only if text explicitly indicates the firm uses AI/ML internally (operations/R&D/etc.); else 0
- notes: short string, "" if none

ALLOWED ROLE LABELS (use exactly these strings):

"compute", "infra", "data_tools", "models", "apps", "other", "unclassified"

ROLE DEFINITIONS (decide roles even without AI mentions)

- compute:
 Hardware + semiconductor ecosystem used for compute/networking.
 Includes: chips/semiconductors, wafers/foundry/fab, OSAT, packaging, assembly, electrical test, semiconductor manufacturing/testing, boards/servers, networking equipment sold as products.
- infra:
 Infrastructure/services to run workloads.
 Includes: cloud hosting/IaaS/PaaS, data centers/colocation, managed compute clusters, edge/CDN, telecom/ISP, managed networking services.
- data_tools:
 Data and tooling for data/ML systems.
 Includes: datasets/data licensing, labeling/annotation/RLHF, synthetic data, ETL/warehousing, databases (incl vector DB), MLOps, monitoring/observability/evaluation tools.
- models:
 Develops/sells AI/ML models or model APIs.
 Includes: LLMs/foundation models, embeddings,

training/fine-tuning, inference, "model API".

- apps:
End-user software/products (usually SaaS or packaged software).
If AI is mentioned as a feature, keep role as apps; AI involvement goes in the AI flags.
- other:
Clearly unrelated to AI supply chain (e.g., mattresses, food, apparel) and not fitting the roles above.
- unclassified:
Not enough information to confidently assign a role.

PRIMARY VS SECONDARY RULES

- Always choose exactly ONE role_primary.
- role_secondary is OPTIONAL: set only if a second role is clearly supported; otherwise null.
- role_secondary must be different from role_primary.
- If role_primary is "other" or "unclassified", role_secondary MUST be null.

AI FLAG RULES (explicit only)

- ai_producer_flag = 1 ONLY if text_input explicitly includes AI/ML terms (AI, artificial intelligence, machine learning, LLM, NLP, computer vision, model, embeddings, inference, etc.) AND clearly indicates customer-facing offering (provide/offer/sell/API/platform/service) of AI/ML capability.
- ai_user_flag = 1 ONLY if text_input explicitly states internal use of AI/ML (use/leverages AI/ML) for its own operations/products, without clear selling of AI itself.
- If AI/ML is not explicitly stated, BOTH flags must be 0.

OUTPUT REQUIREMENTS

- Output ONLY valid JSON.
- Use exactly these keys and no others:
role_primary, role_secondary, ai_producer_flag, ai_user_flag, notes
- ai_producer_flag and ai_user_flag must be integers 0 or 1.
- notes must be a string (" " if none).