V1.0 · AUGUST 2026

The adoption gap

The first three adoption figures I found were 88%, 19.8% and 5%. Twenty statistics graded on six criteria, to work out which one to plan against.

Nikolaos Broikos · broikos.gr August 2026 20 statistics graded on 6 criteria

88%

ADOPTION, SELF-SELECTED VENDOR PANEL

19.8%

US FIRMS, OFFICIAL PROBABILITY SURVEY

19.95%

EU ENTERPRISES, EUROSTAT, 2024 REF. YEAR

68pts

THE GAP · A FACTOR OF 4.4

01 The decision this was written for

I had a budget question. How much of my own operation should I be automating this year, and how far behind am I if I do nothing? It is the same question every business owner in Greece is being asked by somebody trying to sell them something, and I wanted to answer it from evidence rather than from the deck in front of me.

So I went looking for the adoption rate. **The first three numbers I found were 88%, 19.8% and 5%.**

Those are not estimates of the same thing arrived at with different precision. They are answers to different questions, published as if they were answers to mine. And the difference between them is not academic: **at 88% adoption you are dangerously late and should spend now; at 19.8% you are ordinary and can afford to be deliberate.** One of those readings costs real capital.

This paper is what happened when I tried to work out which number to plan against. It grades twenty widely-cited statistics on six criteria, who was asked, how they were selected, what exactly was counted, over what window, who paid, and whether the method is disclosed at all, and it reports what survives.

Why the answer matters more in 2026 than it did in 2023

In 2023 an inflated adoption figure cost you an uncomfortable board slide. In 2026 it is being used to justify capital allocation, headcount decisions and national policy, and the gap between the loudest number and the best-evidenced one has not narrowed: it has widened, because the loudest numbers are now published by companies whose revenue depends on the answer.

There is also a specifically Greek reason to ask now. The European figures have moved far enough that Greece's position can finally be stated with a real instrument rather than inferred from a vendor survey, and when you do that the story is not the one the headlines tell. That is section 05.

What this is, and what it is not

It is not an argument that AI adoption is overstated and everyone should relax. Two of the findings here run the other way, and the paper says so.

It is **a working method for deciding which number to trust**, demonstrated on the twenty statistics a Greek business owner is most likely to meet this year. The grading criteria are published, the sources are listed with their retrieval dates, and the statistics that failed verification are recorded alongside those that passed: including the ones this paper originally intended to use and had to drop.

The companion paper to this one censuses who publishes business research in Greece and whether anyone checks it. **This is that method turned on the loudest numbers in the market right now.**

02

Nobody agrees how many companies use AI

Ask how many companies use artificial intelligence and you will get an answer between 5% and 88% depending entirely on who you ask. The difference is not measurement error. It is the question.

THE SAME QUANTITY, THREE INSTRUMENTS

SOURCE	INSTRUMENT	FIGURE
McKinsey, State of AI	Self-selected online panel, n=1,993, no reference window	**88%**
US Census Bureau, BTOS	Official probability survey, “past two weeks”	19.8%
Eurostat	Official survey, 10+ employees, n=157,000, 2024 ref. year	19.95%

A gap of 68 percentage points. A factor of 4.4.

Two things make this more than a gotcha.

The two official instruments agree with each other. The US Census Bureau and Eurostat are separate agencies, on separate continents, using different questionnaires, in different languages, for different governments, and they land 0.2 points apart. Neither has anything to sell.

Eurostat's figure is **19.95%** of enterprises with ten or more employees, and it is a **2024 reference year** collected in early 2025: not a 2025 measurement, as it is almost universally reported. The survey is *ICT usage and e-commerce in enterprises*: 157,000 enterprises sampled from a population of 1.53 million. Dispersion across member states is enormous, from Romania at 5.21% to Denmark at 42.03%, with Finland at 37.82% and Sweden at 35.04%.

The definition matters more than the number. Eurostat counts an enterprise as using AI if it uses **at least one** of eight named technologies: text mining, speech recognition, natural language generation, image, video or audio generation, image recognition, machine learning, robotic process automation, or autonomous physical movement. That is a wide net. An enterprise running one supplier's OCR on invoices counts. Which makes the 19.95% floor, not a ceiling, and makes the distance to 88% harder to explain away as definitional narrowness at the official end.

The US data reproduces the size gradient from a completely different instrument: 37% of firms with 250+ employees, under 20% for firms with fewer than 20.

FINDING 01

The gap is manufactured by wording

The US survey asks whether the business used AI *in the past two weeks*. The consultancy survey asks whether the organisation uses AI *in at least one function*, with no reference window at all.

FINDING 02

The loudest source contradicts itself

The same survey reporting 88% adoption reports roughly 6% attributing over 5% of EBIT to AI, and nearly two-thirds not yet scaling. The 88% travels. The 6% does not.

FINDING 03

Six of twenty sell what they measure

Six of the twenty most-cited statistics in this field are published by organisations with a direct commercial interest in the answer.

FINDING 04

Productivity splits by task, not tool

Two randomised trials, opposite signs: 55.8% faster on a greenfield task, 19% slower for experienced developers on large mature codebases.

03 Grading the evidence

This paper does not produce a better adoption number. It argues the number is a property of the instrument, so it grades the instruments, and publishes what grading them cost.

The population rule, first

Twenty statistics were selected before any of them were read, on a rule published here so the grid cannot be a hit list. A statistic qualified if it **appears in at least three secondary sources**, or comes from an **official statistical agency**, or comes from a **peer-reviewed or preprint study with a stated design**. Nothing was added or dropped after grading began.

Each was then graded on six criteria: what instrument produced it, its sampling frame, whether n is stated, whether the question wording is knowable, whether the publisher has a commercial interest in the answer, and whether the underlying data is reproducible.

What the grid shows

TWENTY STATISTICS, SIX CRITERIA

MEASURE	RESULT
Official statistical instruments	7 of 20
Vendor, consultancy, practitioner surveys or platform telemetry	8 of 20
Randomised controlled trials	2 of 20
Published by an organisation that sells what it measures	**6 of 20**
Question wording knowable from the published source	**7 of 20**

Only seven of twenty publish enough for a reader to know what was actually asked. That is the deeper problem, and it is worse than the conflict-of-interest count. A figure whose question is unknown cannot be compared with anything, and is nonetheless compared with everything.

What opening them cost

Grading was not a paperwork exercise. Sixteen of the twenty are now verified: ten opened during this work, six carried forward from a previous audit. Three remain only partially opened. One was cut. Along the way:

- **Four figures were cut outright.** Two figures from a model provider, an automation-versus-augmentation split and a share-of-jobs figure, **do not appear in the report they are attributed to**. They are probably real and belong to an earlier edition; cited without one, they are unsupported.
- **One denominator error.** A widely-repeated occupational figure was carried as “37.2% of all queries”. The source says **30% of survey respondents**, against 4% of US employment. Different number, different population.
- **One national figure failed twice and was cut.** A Greek enterprise-adoption percentage appeared in no source that was supposed to contain it. The published lowest three in that dataset do not include Greece at all.
- **One sample-size correction.** A euro-area survey was reported everywhere as about 6,000 firms across 12 countries. **The ECB’s own account, opened at primary on 7 August 2026, says something different:** SAFE is described by the ECB as “our survey of over 5,000 firms across euro area countries”, with no country count given in that passage. The press figure is an overstatement of the Bank’s own description, and the country count appears to originate outside the source. The central bank’s own blog says **over 5,000 firms**; the country count comes from a different publication.
- **One reference-year correction.** An official European figure reported as a 2025 measurement is a **2024 reference year** collected in early 2025.

Five of those six corrections point the same way: the error was not in the underlying research. It was introduced between the research and the reader.

The research is mostly sound. The citation of it is not.

A complication that cuts against this paper

Honesty requires reporting the finding that weakens the argument as well as those that strengthen it.

The 88% figure this paper uses as its foil comes from a question about **regular** AI use in at least one business function. “Regularly” is not a loose word. It rules out the firm that experimented once, which is exactly the leniency the paper accuses the instrument of.

The difference between that survey and the official ones is therefore **not simply strictness**. It is the *unit*. The consultancy survey asks whether **at least one business function** regularly uses AI. The central bank asks about the **extent of the firm’s overall adoption**. An organisation where one team uses a chatbot every day is a clear yes to the first and, quite correctly, an early-stage answer to the second.

That makes the gap more interesting rather than less. Two instruments can both be well designed, both be honestly answered, and still produce 88% and 7%, because one is counting functions and the other is measuring firms.

And a defect in this grid

Three of the twenty rows, all from the same consultancy report, remain **only partially opened**. Field dates and sample size are confirmed. The adoption question wording is not: the publisher’s page timed out on three separate attempts, and the wording quoted above is itself consistent secondary reporting rather than something read at source.

That row carries the paper’s own headline comparison. It is recorded as partial, in the grid and here, rather than rounded up to verified.

04 One survey, three headlines

The clearest demonstration that adoption is a property of the threshold, not of the economy, comes from a single survey that produces roughly 71%, 38%, and single digits: depending only on where the line is drawn.

The European Central Bank added ad hoc questions on artificial intelligence to the fourth-quarter 2025 wave of its **Survey on the Access to Finance of Enterprises**, with fieldwork running October to December 2025. Firms were asked about the *extent* of their adoption of three technology families: predictive tools such as text mining, voice and image recognition and machine learning; generative tools such as chatbots and text or image generation; and robotic process automation.

That word: *extent*: is why one survey generates several incompatible headlines.

THE SAME FIRMS, SORTED BY HOW MUCH THEY USE IT

THRESHOLD	WHAT IT MEANS	SHARE OF EURO AREA FIRMS
Advanced stage	Significant or moderate use	38%
Early stage	Very infrequent or experimental use	33%
Any activity at all	Advanced plus early	71%
Intensive use	Responses indicating *significant* use	**7%**

Every “seven in ten firms now use AI” headline is the third row. Every sober account is the first. The last row is the one that almost never travels: **7%** of euro area firms describe their use as significant. Same respondents, same wave, same question: separated only by where the line is drawn.

The ECB states the definition plainly: intensive use means responses indicating *significant use* when assessing adoption, and “most firms report using AI only infrequently or moderately”. The survey is SAFE round 37, over 5,000 firms, October to December 2025.

A third of euro-area firms are using AI very infrequently or experimentally. Whether that is adoption is a definitional choice, not a measurement.

The survey also shows the gradients hold at every threshold: around **45%** of large firms and of listed or venture-capital-backed companies are at an advanced stage, and **56%** of young firms are, with significant use twice as common among the young.

Why this dissolves an apparent contradiction

Greek figures illustrate the problem in miniature, and they can now be stated exactly rather than approximately. The Hellenic Statistical Authority reports that **3,868 enterprises: 8.2%** use artificial-intelligence technologies for at least one business purpose. Coverage of the ECB's work reports around seven in ten Greek firms using AI. Those two numbers have been set against each other in Greek business media as though one of them must be wrong.

The 8.2% carries its conditions with it, and they are the whole point. The denominator is not all Greek companies: it is the **47,246 enterprises employing ten people or more that have internet access**, in NACE divisions 10-63, 68-82 and group 95.1. The reference period is **1 January to 15 September 2025**. The threshold is use for at least one of seven named purposes. A firm of four people using ChatGPT daily is not in this survey at all.

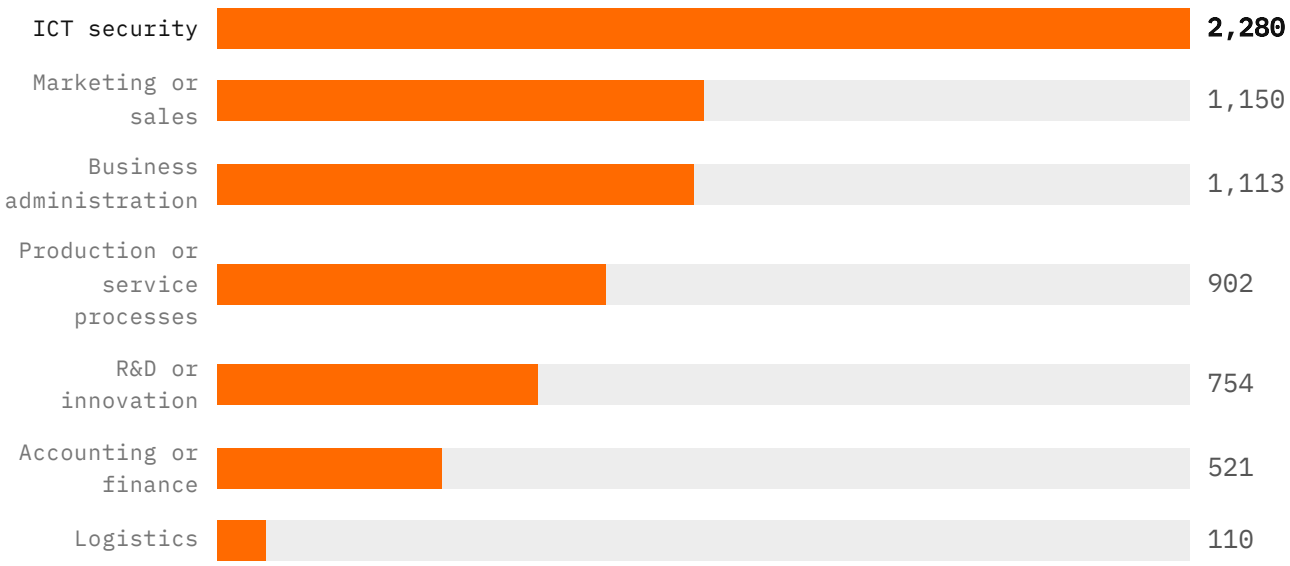

Number of Greek enterprises using AI, by purpose. ELSTAT, Survey on the Use of ICT and E-commerce in Enterprises 2025, released 17 December 2025, Graph 3. Enterprises employing 10+ persons with internet access; a firm may appear in more than one row.

The shape is worth a sentence of its own. The largest single use of AI in Greek enterprise is **ICT security**: roughly twice marketing and sales, and twenty times logistics. Where AI has arrived in Greek firms, it has arrived mostly as something bought inside a security product rather than something a company decided to do.

Neither is. They are different points on the same distribution: one near the strict end, one at the loosest. **An instrument that counts experimental use and an instrument that counts operating systems are not competing estimates; they are answers to different questions.**

The confirmation is in the numbers themselves. The ECB's **intensive-use** figure of **7%** and the Greek statistical office's **8.2%** are **1.2 points apart**. Two entirely independent instruments, a euro-area survey of firm finance and a national statistical collection built for European Business Statistics, run in overlapping periods of 2025 by different bodies, on different samples, for different purposes, **land within about a point of each other as soon as the thresholds are matched**. The nine-fold gap was never a disagreement about Greece. It was a disagreement about what counts.

That convergence is the strongest single piece of evidence in this paper, and it is worth being clear about what it does not prove. Two instruments agreeing does not make either correct; both could share a bias, and neither reaches firms under ten employees, which is the overwhelming majority of Greek businesses. What it establishes is narrower and still useful: **the spread of published adoption figures is a property of thresholds, not of measurement error**, because when the thresholds match, the measurements do.

> **What this paper will not do.** The ECB Economic Bulletin box carries **no country-level breakdown**. The Greek 70% figure does not appear in the primary source and remains press-reported. The national statistical office's figure is no longer in that position: it was opened at source on 6 August 2026 and is cited from the release itself. Its releases we could not open. The *structure* of the argument above is verified at primary. The Greek numbers illustrating it are not, and are described rather than quoted.

The practical consequence for anyone reading an adoption statistic is short. Before comparing two figures, establish whether either counts a firm that tried a chatbot once. If one does and the other does not, they cannot be compared, and the difference between them is not information about the economy.

05 Where Greece actually sits

Once the thresholds are matched, the Greek number stops being contested and starts being uncomfortable. Greece is fourth from the bottom of the European Union on enterprise AI adoption, and over the most recent year of data it is the only direction of travel that matters: while the European average rose by nearly half, Greece went backwards.

Everything in this section comes from a single harmonised instrument: the European Union's annual survey on ICT usage in enterprises, run in every member state under Regulation (EU) 2019/2152 and published by Eurostat.[6] It is the same collection whose Greek arm produced the 8.2% discussed above. Using it removes the comparability problem that defeats most cross-country claims about AI, because every country is asking the same question, of the same population, in the same window.

The population is enterprises employing ten or more people, in NACE divisions 10-63, 68-82 and group 95.1. The threshold is *use of at least one of eight named AI technologies*. The reference year is 2025, published 15 June 2026.[1]

1

ENTERPRISES USING AT LEAST ONE AI TECHNOLOGY, 2025. EUROSTAT ISOC_EB_AI, ENTERPRISES WITH 10 OR MORE PERSONS EMPLOYED.

#	COUNTRY	%	#	COUNTRY	%
1	Denmark	42.0	15	Ireland	19.6
2	Finland	37.8	16	France	18.2
3	Sweden	35.0	17	Slovakia	18.0
4	Belgium	34.5	18	Czechia	17.6
5	Luxembourg	33.6	19	Italy	16.4
6	Netherlands	33.2	20	Croatia	15.2
7	Austria	29.9	21	Latvia	12.2
8	Germany	26.0	22	Portugal	11.5
9	Estonia	23.4	23	Hungary	10.4
10	Slovenia	21.6	24	Cyprus	9.3
11	Malta	21.5	25	**Greece**	**8.9**
12	Lithuania	21.3	26	Bulgaria	8.6
13	Spain	20.3	27	Poland	8.4
n/a	**EU27 average**	**19.9**	**28**	Romania	**5.2**

Twenty-fifth of twenty-eight. Greece sits between Cyprus and Bulgaria, at **44% of the European average**, and closer to Romania than to Spain. There is no threshold argument available here: every country in the table answered the same question.

The year Greece went backwards

A single year's ranking can be an artefact of sampling. The direction of travel cannot.

1

ENTERPRISES USING AT LEAST ONE AI TECHNOLOGY, 2024 TO 2025. SAME INSTRUMENT, SAME POPULATION, CONSECUTIVE WAVES.

SIZE CLASS	GREECE 2024	GREECE 2025	EU 2024	EU 2025
10 or more employees	**9.8**	**8.9**	**13.5**	**19.9**
10-49	8.2	7.3	11.2	17.0
50-249	15.9	15.2	21.0	30.4
250 or more	**24.3**	**37.6**	**41.2**	**55.0**

The European average rose from 13.5% to 19.9%: a gain of nearly half in twelve months. Greece fell from 9.8% to 8.9%. The gap widened from 3.7 percentage points to 11.0 in a single wave. This is not a country moving slowly in the same direction as everyone else. Over this wave it moved the other way.

The gap between Greece and the European average nearly tripled in one year. It is not a lag. It is a divergence.

But the aggregate hides the more interesting fact, and it is the one a business reader should take away: **the fall is not evenly distributed. It is entirely a small-firm story.**

Greek enterprises of 250 or more employees went from 24.3% to **37.6%**: a gain of 13.3 points, faster in absolute terms than the EU's own large-firm gain of 13.8 and from a much lower base. Greek large firms are adopting AI at roughly the pace of large firms everywhere. Meanwhile Greek firms of 10 to 49 employees fell from 8.2% to 7.3%, against a European small-firm figure that rose from 11.2% to 17.0%.

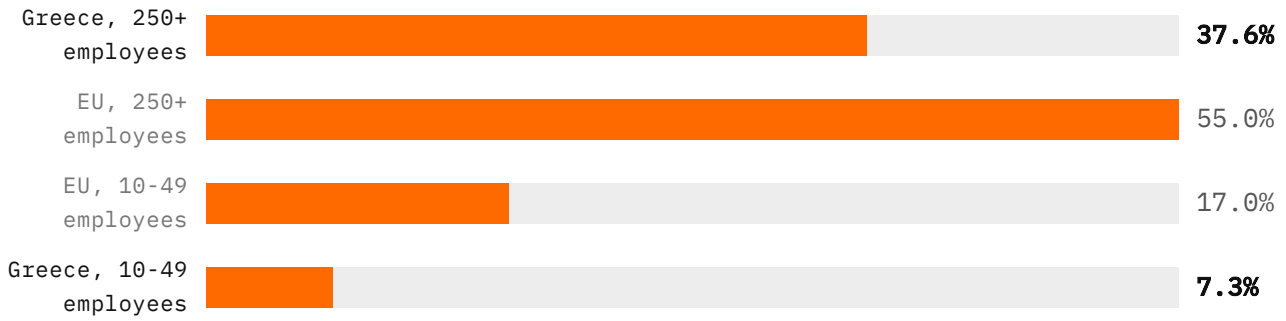

2025. Eurostat isoc_eb_ai. Full series in eurostat-ai-by-size.csv.

A Greek large enterprise is now closer to its European peer than at any point in the series. A Greek small enterprise is further away than it was a year ago. **Whatever is happening in Greece, it is happening to firms that already had the capacity to do it**, and the national average, which is dominated by small firms, records the result as decline.

That matters for how the number is used. “Greek business has not adopted AI” is a fair reading of 8.9%. “Greek business is not adopting AI” is not: one segment of it is adopting quickly. The aggregate is arithmetic about firm-size distribution as much as it is about technology.

The same survey, asked twice, gives two answers

One more property of this dataset deserves its own paragraph, because it is this paper's thesis appearing inside a single instrument rather than between two.

The Greek statistical office reported **8.2%** of enterprises using AI, on a question about seven business *purposes*.[5] Eurostat, publishing from the same national collection, reports **8.9%** on a question about eight named *technologies*, and separately reports **8.1%** for “using AI technologies for at least one of the purposes”.

THREE PUBLISHED FIGURES FOR GREEK ENTERPRISE AI USE IN 2025, ALL FROM THE SAME SURVEY.

PUBLISHED FIGURE	WHAT IT COUNTS	GREECE
Eurostat E_AI_TANY	Uses at least one of eight named AI *technologies*	8.9%
ELSTAT release, Graph 3	Uses AI for at least one of seven *purposes*, of firms with internet access	8.2%
Eurostat E_AI_P1ANY	Uses AI for at least one of the same *purposes*, of all firms 10+	8.1%

A spread of 0.8 points, or **a tenth of the value of the statistic**, produced by nothing except whether the questionnaire asked about technologies or about purposes, and whether the denominator was all firms or only those with internet access. Same respondents. Same fieldwork. Same year.

None of the three is wrong. Anyone quoting one of them without saying which is publishing a number they cannot defend, and the difference is large enough to matter: an 8.9% and an 8.1% will round into different sentences in a policy document.

The useful discipline is not to pick the best figure. It is to notice that a tenth of this statistic's value is a property of the questionnaire, and to treat every published adoption number, including the ones in this paper, as carrying that much play before any real-world disagreement has even begun.

06 It is not a skills gap

The standard explanation for a low national AI figure is talent: not enough people who can build it. For Greece the survey data contradicts that directly, and it does so three separate ways. Greek firms employ technical staff at above the European rate, report the skills barrier *less* often than the European average, and consider adopting AI at exactly the European rate. What they do not have is anywhere to run it.

Greece is above the European average on technical staff

Eurostat measures the share of enterprises employing ICT specialists, on the same survey and the same population.

4

ENTERPRISES EMPLOYING ICT SPECIALISTS, 2024. EUROSTAT ISOC_SKE_ITSPEN2, ENTERPRISES WITH 10 OR MORE PERSONS EMPLOYED. SELECTED COUNTRIES; FULL RANKING IN THE DATA FILES.

RANK	COUNTRY	%
1	Malta	34.1
2	Denmark	30.9
3	Finland	30.5
8	Cyprus	27.5
11	Germany	22.9
12	**Greece**	**22.5**
13	Sweden	21.6
n/a	**EU27 average**	**20.1**
24	France	15.9
26	Spain	14.6
28	Italy	12.4

Greece is **twelfth of twenty-eight and 2.4 points above the European average**: ahead of Sweden, ahead of France by nearly seven points, ahead of Italy by ten. A country ranked 25th on AI adoption is ranked 12th on employing the people who would build it.

Those two facts are difficult to hold together under a skills explanation. Whatever is stopping Greek firms, it is not an unusual scarcity of technical employees relative to their European peers.

Greek firms report the skills barrier less often than average

The survey does not leave this to inference. It asks non-adopters why, from a fixed list, and every country answers the same list.

2

REASONS ENTERPRISES GIVE FOR NOT USING AI TECHNOLOGIES, 2025. EUROSTAT ISOC_EB_AIN2, % OF ALL ENTERPRISES 10+. MULTIPLE ANSWERS PERMITTED.

STATED REASON	GREECE	EU27	DIFFERENCE
The costs seem too high	**6.0**	**4.2**	**+1.8**
Ethical considerations	3.4	2.7	+0.6
AI is not useful for the enterprise	2.4	2.0	+0.5
Incompatible with existing systems	4.8	4.6	+0.2
Availability or quality of the necessary data	4.1	4.8	−0.7
Data-protection and privacy concerns	5.1	5.8	−0.7
Lack of relevant expertise	**6.8**	**7.8**	**−0.9**
Lack of clarity about legal consequences	4.8	5.9	−1.1

Of eight possible reasons, Greek enterprises report **only one materially more often than the European average, and it is cost.** They report the expertise barrier *less* often than the EU. They report legal uncertainty less often, data-protection concerns less often, data-quality problems less often.

The picture the barrier question paints is not of firms that want AI and cannot staff it. It is of firms that have looked at the price.

They consider it at exactly the European rate

The survey asks a further question that is rarely quoted and is, for this argument, the most useful one in the dataset: whether the enterprise has *ever considered* using any of the AI technologies.

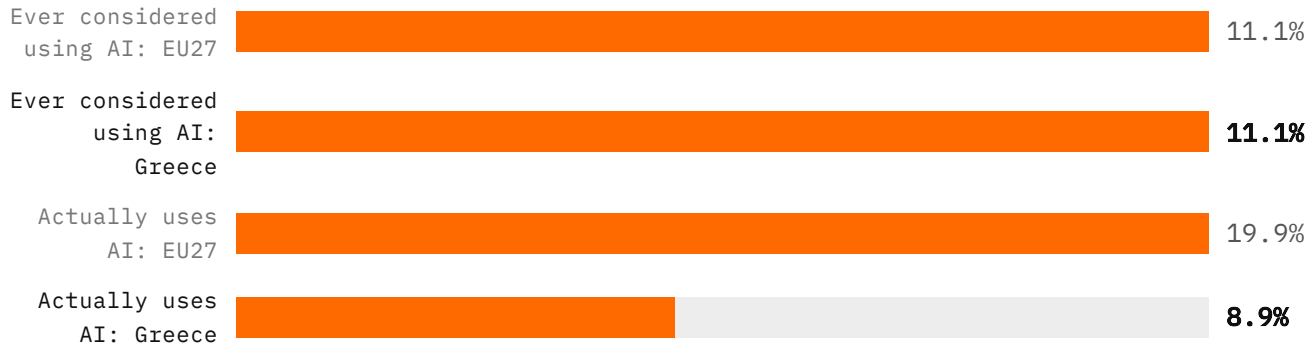

Eurostat isoc_eb_ain2, 2025. Greece 11.05% and EU27 11.10% on the consideration question. Bars are scaled within each pair.

11.05% against 11.10%. On the question of whether Greek management has thought about AI, Greece is indistinguishable from the European average: a five-hundredth of a percentage point apart. On the question of whether they are running it, Greece is at 45% of the European rate.

Greek firms are not less interested. They are less able to convert interest into operation, and the gap opens somewhere between considering and doing.

What is actually missing

The same survey measures the thing that sits between the two, and Greece's position on it is worse than its position on AI.

3

ENTERPRISES BUYING CLOUD COMPUTING SERVICES, 2014 TO 2025. EUROSTAT ISOC_CICCE_USE, ENTERPRISES WITH 10 OR MORE PERSONS EMPLOYED. SURVEY NOT RUN IN EVERY YEAR.

YEAR	GREECE	EU27	GAP
2014	7.6	17.8	10.2
2018	12.8	23.9	11.1
2020	16.7	36.1	19.4
2021	20.7	41.0	20.3
2023	23.6	45.3	21.7
2025	**24.3**	**52.7**	**28.4**

Greek enterprise cloud adoption is 24.3% against a European average of 52.7%. Only Bulgaria is lower. Greece is not merely behind; the gap has been widening continuously for a decade, from 10.2 points in 2014 to 28.4 in 2025. Between 2023 and 2025 Greece gained 0.7 points and the European average gained 7.4.

This is the missing term. Modern AI is, for almost every firm that is not building models itself, a cloud service: consumed through an API, a platform feature, or software that calls one. A company with no cloud footprint is not one step away from AI adoption. It is two, and the first step is the expensive one.

Eurostat publishes the joint distribution, which closes the argument without requiring anyone to believe an inference:

2

THE AI AND CLOUD CROSS-TABULATION, 2025. EUROSTAT ISOC_EB_AIN2, % OF ALL ENTERPRISES 10+.

COMBINATION	GREECE	EU27
Uses AI *and* buys any cloud service	**6.5**	**16.2**
Uses AI *and* buys sophisticated or intermediate cloud *and* does data analytics	4.7	11.1
Uses AI *and* performs data analytics	6.3	13.8
Performs data analytics but uses *no* AI and buys *no* advanced cloud	**16.6**	**11.7**
Uses at least two AI technologies	5.7	13.0
Uses AI for marketing or sales	2.4	6.9

The fourth row is the Greek economy in one line. **Greece over-indexes, by nearly five points, on firms that analyse their own data without advanced cloud and without AI.** These are companies doing the work, they have the data, they have the analyst, they are producing the reports, on infrastructure that will not carry the next step.

Greek firms consider AI at the European rate, employ technical staff above it, report the skills barrier below it, and rank second from bottom on cloud. The constraint is infrastructure, and it is priced as cost.

Read the four findings together and they form a chain rather than a list. Greek management considers AI as often as European management. It reports expertise as a barrier less often than European management. It employs ICT specialists more often than the European average. And it operates on a cloud base half the European level, which has been falling further behind every year since 2014, so when it prices an AI project, it is pricing a migration as well, and it answers “the costs seem too high” more often than anyone else in the table.

That reading is consistent with every barrier in the survey, and no other single explanation is. It is also, unlike a skills shortage, a problem with a known and purchasable solution, which is the only encouraging sentence in this section.

> **What this section does not establish.** These are cross-sectional associations from one survey wave, not a causal identification. Cloud adoption and AI adoption are both downstream of firm size, sector and capital access, and this analysis does not separate them. A firm-level panel would be needed to say that cloud absence *causes* AI absence rather than that both follow from something else.
>
> What the data does establish, and what the common explanation cannot survive, is narrower: the skills account is inconsistent with three independent measures in the same instrument: specialist employment above the EU average, the expertise barrier reported below it, and consideration identical to it. An explanation that contradicts the survey's own answers is not the leading candidate, whatever the causal story turns out to be.
>
> Note also that the whole of this section, like the whole of this survey, describes only enterprises employing ten or more people. In Greece that is a smaller share of the economy than in most of the countries it is being compared against, and every figure here should be read as a statement about that segment.

07

What Greek firms actually bought

The aggregate says Greek adoption is low. The breakdown says something more specific and more useful: of seven business purposes the survey asks about, there is exactly one where Greece is *ahead* of the European average, and it is the one you do not choose.

Seven purposes, one exception

2

ENTERPRISES USING AI, BY BUSINESS PURPOSE, 2025. EUROSTAT ISOC_EB_AIN2, % OF ALL ENTERPRISES 10+. GREECE AS A SHARE OF THE EU RATE IN THE FINAL COLUMN.

PURPOSE	GREECE	EU27	GREECE AS % OF EU
ICT security	**4.8**	**3.9**	**121%**
Production or service processes	1.9	4.1	45%
R&D or innovation	1.6	3.8	42%
Business administration	2.3	6.2	38%
Marketing or sales	2.4	6.9	35%
Accounting or finance	1.1	4.6	24%
Logistics	**0.2**	**1.2**	**19%**

ICT security is the only purpose on which Greek enterprises exceed the European average: 4.8% against 3.9%, or 121% of the EU rate. On every other purpose Greece runs between a fifth and a half of the European level.

The Greek statistical office's own breakdown says the same thing from the other direction: of the 3,868 Greek enterprises using AI, **2,280 use it for ICT security**: roughly twice as many as use it for marketing and sales, and twenty times as many as use it for logistics.

This is a distinctive shape and it is worth being precise about what it means. AI for ICT security is, in the overwhelming majority of cases, not a project. It is a property of a product: the endpoint protection, the mail filter, the firewall, the fraud-scoring layer in a payment provider. A firm acquires it by renewing a security subscription, and a great many of the firms that answer “yes” to this question did not make an AI decision at all.

Greece is above the European average on the one AI use that arrives inside something you were buying anyway, and below it on all six that require a decision.

Read that way, the Greek position is worse than the headline 8.9% suggests rather than better. A meaningful share of the measured adoption is *incidental*, embedded in vendor products, and the deliberate, operationally applied adoption that the productivity literature is about is thinner than the aggregate implies. Logistics, at 0.2% of enterprises, is the clearest case: in a country whose e-commerce economics turn on delivery, effectively nobody is applying AI to it.

The technology mix says the same thing

2

ENTERPRISES USING EACH AI TECHNOLOGY, 2025. EUROSTAT ISOC_EB_AIN2, % OF ALL ENTERPRISES 10+. A FIRM MAY USE SEVERAL.

TECHNOLOGY	GREECE	EU27	GREECE AS % OF EU
Text mining	6.1	11.8	52%
Machine learning for data analysis	3.1	5.1	60%
Generating pictures, video or speech	4.0	9.6	42%
Natural-language generation	3.7	8.8	42%
Speech recognition	3.2	7.2	45%
Image recognition	2.1	3.8	54%
Process automation	**1.9**	**5.4**	**35%**
Autonomous robots or vehicles	**0.4**	**1.4**	**27%**

Greece is least far behind on the technologies you consume, text mining at 52% of the EU rate, machine learning at 60%, image recognition at 54%, and furthest behind on the two that require changing how work is organised: **process automation at 35% and autonomous systems at 27%.**

The pattern across both tables is consistent, and it is the same pattern the productivity evidence in the next section describes at the level of an individual task. Adoption is easiest where AI is a feature of something already being purchased, and hardest where it requires a firm to alter a process. Greece's deficit is concentrated exactly where the alteration is largest.

Depth, not just breadth

One more cut separates a firm experimenting from a firm operating. Eurostat publishes how many technologies each enterprise uses.

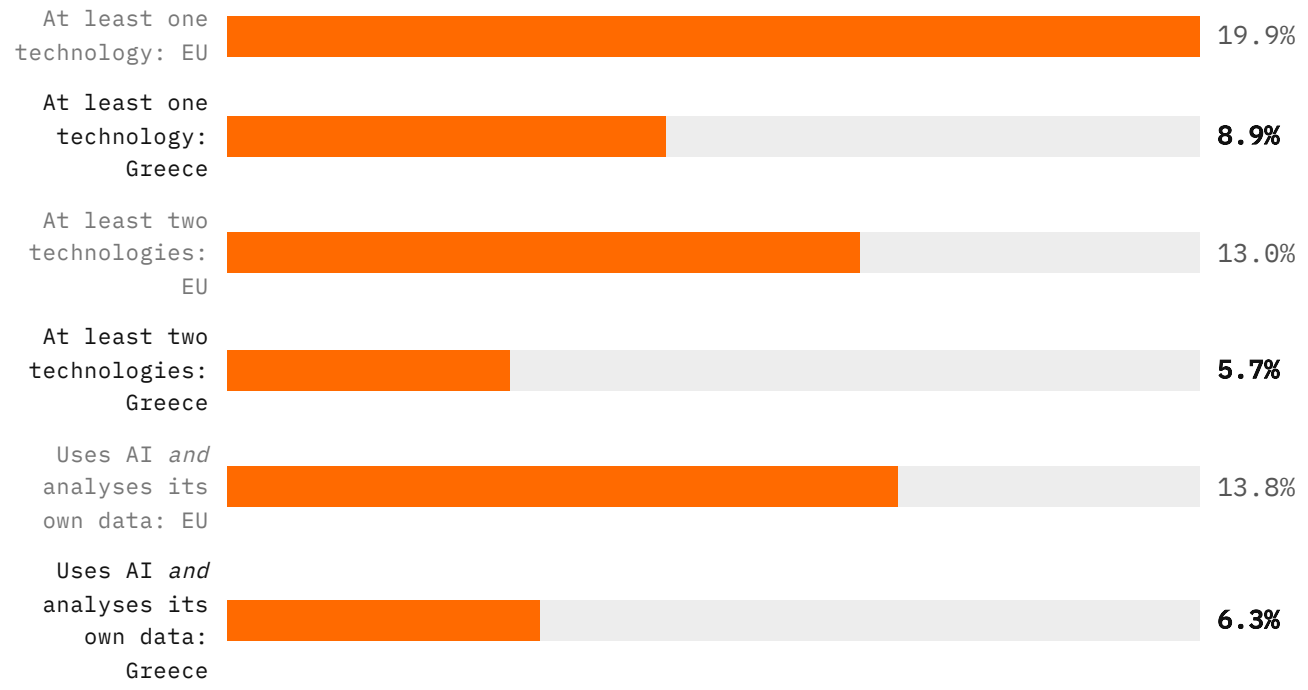

2025, all bars scaled to the EU “at least one” figure. Eurostat isoc_eb_ain2.

The Greek ratio holds at roughly 45% of the European level whichever depth threshold is applied: 8.9 against 19.9 for any use, 5.7 against 13.0 for two or more technologies, 6.3 against 13.8 for AI combined with the firm's own data analysis. Greece is not disproportionately full of shallow adopters; it has proportionally fewer firms at every depth.

That is a mildly reassuring finding for the quality of the Greek number and an unwelcome one for anybody hoping the gap is an artefact. The adopters Greece has appear to be about as serious as European adopters. There are simply half as many of them.

08 Does it make anyone more productive?

The two best-designed studies in the field disagree, and the disagreement is the finding.

TWO RANDOMISED CONTROLLED TRIALS, OPPOSITE SIGNS

STUDY	SETTING	RESULT
Peng et al. (2023)	Greenfield task: implement an HTTP server	**55.8% faster**
METR (2025)	Experienced developers, mature repos, 1M+ lines	**19% slower**

Same technology, opposite signs, and the variable is not the tool. It is the codebase. Writing something new from nothing is where the gains are; working inside a large, mature, high-standard system is where they reverse.

A third trial, and the pattern underneath all three

A field study of customer support gives the clearest view of the mechanism. Brynjolfsson, Li and Raymond followed **5,179 support agents** given access to a generative assistant, measuring issues resolved per hour.

GENERATIVE AI AT WORK: EFFECT BY WORKER EXPERIENCE

GROUP	PRODUCTIVITY EFFECT
All agents	+14%
Novice and low-skilled	**+34%**
Experienced and highly skilled	**minimal**

The authors' explanation is that the model **disseminates the best practices of more able workers**, moving newer staff down the experience curve faster. That is not a claim about the technology being clever. It is a claim about where the knowledge in an organisation already sits.

Put the three trials together and one pattern runs through all of them:

Gains concentrate where the worker or the task is least established. They shrink, and can reverse, exactly where expertise already exists.

A greenfield task has no incumbent expertise to compete with, and the gain is large. A novice agent has little, and the gain is large. An experienced agent has a lot, and the gain vanishes. A senior engineer inside a mature codebase has a great deal, and the measured effect turns negative.

This reframes the deployment question. The useful thing to ask is not “will this make my team faster”. It is **where in this organisation is expertise thin, and how expensive is it to verify the output there**, because those two conditions, not the choice of tool, decide the sign of the result.

Two practitioner instruments point the same way as the second result. 66% of developers report spending more time fixing AI output that is “almost right, but not quite”: the single most-cited frustration at 45%. And the DevOps research programme finds a negative relationship between AI adoption and delivery stability absent disciplined automated testing.

AI moves work from producing to verifying. Whether that is a gain depends entirely on how expensive verification is in your setting.

Three statistics that are routinely misread

- **“Costs fall one to two orders of magnitude per year.”** True at *fixed capability*. Quoted as though AI is getting cheaper generally. What gets cheaper is last year's capability.
- **Occupational shares from model-provider telemetry.** That 37.2% of queries come from computer and mathematical occupations is a fact about who uses that product, read as a fact about which occupations are exposed to AI.

- **The closed-to-open capability gap**, narrowing from 8.04% to 1.70%, is a leaderboard measurement. It describes benchmark scores, not business outcomes.

09 What to conclude

01 **Stop quoting adoption percentages without the question.** A number between 5% and 88% is available for any argument. Ask what was asked, of whom, over what window.

02 **Weight the instrument, not the headline.** An official probability survey and a self-selected vendor panel are not two estimates of the same quantity. The official one is the estimate; the vendor one is the sentiment.

03 **Expect gains where verification is cheap and reversals where it is expensive.** The trial evidence is not ambiguous once you read it as a verdict on the setting rather than on the technology.

04 **Assume you are in the 94%, not the 6%.** Nearly two-thirds of organisations in the most optimistic survey available have not begun scaling.

05 **If you are Greek and small, the binding constraint is probably not talent.** Three independent measures in the same survey say otherwise. Price the infrastructure before you price the model.

06 **Check whether your AI is incidental.** A firm whose only AI is in its security stack has bought a product, not adopted a capability, and the distinction is invisible in every adoption statistic including the ones above.

The Greek finding, stated plainly

This paper began as an argument about measurement and ended with a substantive result it did not set out to find. Both are worth separating.

The measurement result is that the published spread of AI adoption figures, the one that produces a 5% and an 88% for the same year, is almost entirely a property of thresholds, populations and question wording rather than of genuine disagreement. The proof is that the figures converge as soon as thresholds are matched: the ECB's intensive-use 7% and the Greek statistical office's 8.2% land 1.2 points apart despite being produced by different institutions, on different samples, for different purposes. And the residual play is not small: within one national survey, asking about technologies rather than purposes moves the Greek number from 8.1% to 8.9%.

The substantive result is about Greece, and it is uncomfortable enough that it deserves to be stated without softening. On the single harmonised European instrument, Greek enterprise AI adoption is **8.9% against a European average of 19.9%**, twenty-fifth of twenty-eight member states. In the most recent wave the European average rose by nearly half and the Greek figure fell. Almost all of the decline sits in firms under fifty employees; Greek firms of 250 or more grew adoption by 13.3 points in the same period.

The usual explanation for that shape is a shortage of talent, and the survey contradicts it three ways: Greek enterprises employ ICT specialists at 22.5% against a European average of 20.1%, they report lack of expertise as a barrier *less* often than the European average, and they report having considered AI at a rate identical to the European average to within a twentieth of a point. The one barrier they report materially more often is cost, and they sit second from bottom in the Union on cloud adoption: 24.3% against 52.7%, a gap that has widened every year since 2014.

What this adds up to is not a country that cannot do the work. It is a country whose firms are thinking about it at the European rate, employing the people at above the European rate, and operating on infrastructure that turns every AI project into an infrastructure project first. That is a more tractable problem than a talent shortage, and a

more expensive one to ignore, because the gap compounds: each year that the cloud base does not move, the AI step gets larger relative to peers who took it earlier.

The final observation is the one that should worry a Greek operator most, and it comes from the purpose breakdown rather than the headline. Of seven business purposes, the only one where Greece exceeds the European average is **ICT security**: the purpose that arrives inside a product renewal rather than out of a decision. Greek firms are, on this evidence, above average at receiving AI and below average at applying it. An adoption statistic cannot tell those two things apart, which is a reasonable note on which to end a paper about what adoption statistics cannot tell you.

TWO INSTRUMENTS THAT WERE NOT BUILT

A Greek job-postings corpus. Adoption surveys record what companies *say*. Job advertisements record what they are *willing to pay for*. Where the two disagree, the second is the better evidence, and nobody publishes it for Greece.

A task-level study. Baseline time and rework against post-AI time and rework, with the column most case studies omit: which tasks were tried and reverted to manual, and why. A case study with no failures in it is marketing.

Neither was built. Both need inputs this project did not obtain, a defined set of target sectors for the first, and a real task list including the abandoned experiments for the second, and both were left unbuilt rather than approximated from what was available. **They are described here so that a reader can see the shape of the paper this is not.** What this paper does instead is grade the evidence that already exists, which turned out to be the more useful job: the central finding is that the disagreement between published adoption figures is an artefact of thresholds, and that finding required no new fieldwork, only opening the sources.

10 Sources, and how to re-derive every figure

Every European figure in this paper comes from Eurostat's public dissemination API and can be reproduced by anyone with a browser. This section gives the dataset codes, the exact filters, and the publication dates, so that no number here has to be taken on trust.

The base endpoint is https://ec.europa.eu/eurostat/api/dissemination/statistics/1.0/data/<dataset>. It requires no key and no registration. Appending ?format=JSON&lang=EN plus the dimension filters below returns the figures used here.

EUROSTAT DATASETS USED, WITH THE PUBLICATION DATE OF THE EXTRACT TAKEN ON 6 AUGUST 2026.

DATASET	TITLE	FILTERS APPLIED	PUBLISHED
isoc_eb_ai	Artificial intelligence by size class of enterprise	indic_is=E_AI_TANY, unit=PC_ENT, size classes GE10 / 10-49 / 50-249 / GE250	15 Jun 2026
isoc_eb_ain2	Artificial intelligence by NACE Rev. 2 activity	nace_r2=C10-S951_X_K, unit=PC_ENT; technology, purpose, barrier and cloud cross-tab indicators	15 Jun 2026
isoc_cicce_use	Cloud computing services by size class of enterprise	indic_is=E_CC, unit=PC_ENT, size_emp=GE10	27 Feb 2026
isoc_ske_itspen2	Enterprises that employ ICT specialists by NACE Rev. 2 activity	nace_r2=C10-S951_X_K, unit=PC_ENT	15 Jun 2026

The national release used alongside them is the Hellenic Statistical Authority's announcement *Survey on the Use of Information and Communication Technologies and E-commerce in Enterprises, 2025*, published 17 December 2025, reference period 1 January to 15 September 2025. It is a five-page PDF on statistics.gr and is the source of the 8.2% figure, the 3,868 enterprises, the 47,246 denominator and the purpose breakdown.

The indicator codes, in plain language

EUROSTAT INDICATOR CODES CITED IN THIS PAPER.

CODE	WHAT IT COUNTS
E_AI_TANY	Uses at least one of eight AI technologies: text mining, speech recognition, natural-language generation, image recognition, machine learning for data analysis, process automation, autonomous robots or vehicles, generating pictures/video/speech
E_AI_TGE2	Uses at least two of those technologies
E_AI_P1ANY	Uses AI for at least one of seven purposes: marketing or sales, production or service processes, business administration, logistics, ICT security, accounting or finance, R&D or innovation
E_AI_EC	Has ever *considered* using one of the AI technologies
E_AI_CC	Uses AI technologies *and* buys any cloud computing service used over the internet
E_AIX_CC1SIX_DA	Performs data analytics but uses no AI technologies and buys no sophisticated or intermediate cloud
E_AI_BLE / E_AI_BCST	Does not use AI because of a lack of relevant expertise / because the costs seem too high
E_CC	Buys cloud computing services used over the internet

Derived files published with this paper

The extracts were saved as tidy CSVs rather than left as API calls, so that a reader checking this paper in a year's time sees the figures as they stood on the date of writing rather than as Eurostat has since revised them. Revisions to these series are routine.

FILE	CONTENTS
eurostat-ai-ranking-2025.csv	All 28 territories ranked on AI adoption, 2025
eurostat-ai-by-size.csv	Greece and EU27 by size class, 2024 and 2025
eurostat-ai-mechanism-2025.csv	All 48 AI indicators for Greece against EU27, including technologies, purposes, barriers and the cloud cross-tabulation
eurostat-cloud-trend.csv	Cloud adoption, Greece and EU27, 2014 to 2025
eurostat-ict-specialists.csv	ICT-specialist employment ranking, 2024

Two conventions apply throughout. Percentages are reported to one decimal place where the source gives more; the underlying values are in the files. And where a figure appears in the text and in a table, the table is authoritative: the prose was reconciled against the data files by script rather than by reading.

11 References

Every source below was opened by the author at the URL shown. Where a source could not be opened, that is stated in the entry rather than omitted. Dates in brackets are the dates of access.

Official statistics

1. **Eurostat. Artificial intelligence by size class of enterprise.** Dataset isoc_eb_ai. Luxembourg: Statistical Office of the European Union. Data published 15 June 2026.
   `https://ec.europa.eu/eurostat/databrowser/view/isoc_eb_ai/default/table`
   Extracted via the Eurostat dissemination API, 6 August 2026. Used for the country ranking and the size-class series.
2. **Eurostat. Artificial intelligence by NACE Rev. 2 activity.** Dataset isoc_eb_ain2. Data published 15 June 2026.
   `https://ec.europa.eu/eurostat/databrowser/view/isoc_eb_ain2/default/table`
   Extracted 6 August 2026. Used for technologies, purposes, barriers, the consideration question and the AI-cloud cross-tabulation.
3. **Eurostat. Cloud computing services by size class of enterprise.** Dataset isoc_cicce_use. Data published 27 February 2026.
   `https://ec.europa.eu/eurostat/databrowser/view/isoc_cicce_use/default/table`
   Extracted 6 August 2026. Used for the 2014-2025 cloud series.
4. **Eurostat. Enterprises that employ ICT specialists by NACE Rev. 2 activity.** Dataset isoc_ske_itspen2. Data published 15 June 2026.
   `https://ec.europa.eu/eurostat/databrowser/view/isoc_ske_itspen2/default/table`
   Extracted 6 August 2026.

5 **Hellenic Statistical Authority (ELSTAT). Survey on the Use of Information and Communication Technologies and E-commerce in Enterprises, 2025.** Piraeus, 17 December 2025. Reference period 1 January to 15 September 2025.
https://www.statistics.gr/documents/20181/4d030e6d-89df-280f-e3ae-7471cf769816
Five-page release, opened at source 6 August 2026 after four earlier retrieval failures. Source of the 8.2% figure, the 3,868 enterprises, the 47,246 denominator and the purpose breakdown.

6 **Regulation (EU) 2019/2152 of the European Parliament and of the Council of 27 November 2019 on European business statistics.** *Official Journal of the European Union* L 327, 17 December 2019.
https://eur-lex.europa.eu/eli/reg/2019/2152/oj
The legal basis under which the enterprise ICT survey is conducted in every member state, and therefore the reason the country comparison is valid.

7 **Council Regulation (EEC) No 696/93 of 15 March 1993 on the statistical units for the observation and analysis of the production system in the Community.** *Official Journal* L 76, 30 March 1993.
https://eur-lex.europa.eu/eli/reg/1993/696/oj
Defines the statistical unit "enterprise", the change to which makes ELSTAT's 2025 results non-comparable with 2024.

Central bank and institutional research

1 **European Central Bank. Adopting and investing in AI: evidence from euro area firms.** *ECB Economic Bulletin*, Issue 2/2026, Box 6.
https://www.ecb.europa.eu/press/economic-bulletin/focus/2026/html/ecb.ebbox202602_06~c9f485c478.en.html
Source of the threshold ladder: 71% any use, 38% advanced, 7% intensive. Survey on the Access to Finance of Enterprises (SAFE), round 37, over 5,000 firms, fielded October-December 2025. Opened at source; contains no country-level breakdown.

2 **European Central Bank. What separates firms that use AI intensively from the rest?** The ECB Blog, 24 June 2026.
https://www.ecb.europa.eu/press/blog/date/2026/html/ecb.blog20260624~44f70da110.en.html
Opened at source. Used for the definition of intensive use.

3 **Stanford Institute for Human-Centered Artificial Intelligence. AI Index Report 2026.** Stanford University.
https://aiindex.stanford.edu/
Used in the evidence-appraisal grid.

4 **United States Census Bureau. Business Trends and Outlook Survey (BTOS), AI use supplement.**
https://www.census.gov/hfp/btos/
Cited as the probability-sample counter-example to vendor panels. Used in the appraisal grid.

Peer-reviewed and preprint research

1 **Peng, S., Kalliamvakou, E., Cihon, P. and Demirer, M. (2023). The Impact of AI on Developer Productivity: Evidence from GitHub Copilot.** arXiv preprint arXiv:2302.06590.
https://arxiv.org/abs/2302.06590
Randomised controlled trial. Greenfield task: implement an HTTP server. Result: 55.8% faster.

2 **METR (2025). Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity.**
https://metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study/
Randomised controlled trial. Experienced developers, mature repositories exceeding one million lines. Result: 19% slower.

3 **Brynjolfsson, E., Li, D. and Raymond, L. (2023). Generative AI at Work.** NBER Working Paper No. 31161. National Bureau of Economic Research.
`https://www.nber.org/papers/w31161`
Field study of customer-support agents; gains concentrated among less experienced workers.

Commercial and industry research, graded in the appraisal grid

1 **McKinsey & Company. The State of AI.** QuantumBlack.
`https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai`
Partially verified. *Field dates and sample confirmed; the verbatim questionnaire wording for the headline adoption item could not be retrieved after four attempts by two different methods, and is recorded as unresolved rather than assumed. The construct is reported by McKinsey as regular AI use in at least one business function.*

2 **Stack Overflow. Developer Survey 2025.**
`https://survey.stackoverflow.co/2025/`
Self-selected respondent panel. Graded accordingly in the appraisal grid.

3 **Anthropic. Anthropic Economic Index.**
`https://www.anthropic.com/economic-index`
Platform-usage data, not a population survey. Graded accordingly. The author uses Anthropic products; see the interest statement below.

Data published with this paper

1 **Broikos, N. (2026). Evidence-appraisal grid: twenty cited AI statistics graded on six criteria.** Dataset.
`https://broikos.gr/research/data/paper-03/appraisal.csv`

2 **Broikos, N. (2026). Eurostat extracts, claim ledger and source log for “The adoption gap”.** Datasets.
`https://broikos.gr/research/data/paper-03/`
Contains the five Eurostat extracts frozen at the date of writing, the 61-row claim ledger recording every figure’s verification status including the four that were cut, and the source log.

12 Limitations

LIMITATION	HANDLING
“AI adoption” has no agreed definition: the paper's subject, not a fixable flaw	Every cited figure carries its question wording where one is published, and a note where it is not
Three of twenty graded statistics remain only partially opened	All three are the same consultancy’s rows. Field dates and sample are confirmed; the adoption question wording behind this paper’s headline comparison is not, after three failed retrieval attempts. Recorded as partial rather than rounded up
Four figures were cut from the grid during verification	Two attributed to a report edition that does not contain them, one with a wrong denominator, one that failed twice. All recorded in the claim ledger with reasons rather than quietly dropped
A Greek 9% enterprise-adoption figure was carried in an earlier draft and failed verification twice	**Cut, then replaced.** The 9% was not in the ECB material and not in Eurostat's published country data. The figure now in the paper is **8.2%**, taken from the ELSTAT press release of 17 December 2025 opened at source, with its denominator, reference period and threshold printed alongside it. The original claim stays in the ledger as cut
The Greek figure covers only enterprises employing ten or more people	Stated wherever the figure appears. Greek business is overwhelmingly smaller than that, so 8.2% describes the measured segment and not the economy
ELSTAT states its 2025 results are **not directly comparable with 2024**, the statistical unit having changed from legal unit to enterprise[7]	No year-on-year change is computed from this source anywhere in this paper
Official statistics lag 6 to 18 months	Publication dates on every figure
The cross-country comparison is cross-sectional. It cannot show that low cloud adoption causes low AI adoption	Stated in the section itself. The claim made is the weaker and defensible one: that the skills explanation is inconsistent with three separate measures in the same instrument. A firm-level panel would be required to establish direction
Every Eurostat figure here covers only enterprises employing ten or more people	Stated wherever a figure appears. Greek business skews smaller than most comparators, so the omitted segment is larger in Greece than in the countries it is compared with: which, if anything, means the national picture is worse than these numbers show, not better
The 2024-to-2025 movement is two waves of a sample survey, not a measured trend	Reported as the movement between two published waves, with both values printed. Sampling variation could account for part of a 0.9-point fall; it cannot account for a 6.4-point European rise in the opposite direction
Reasons for non-adoption are self-reported and permit multiple answers	Reported as what firms say rather than as diagnosis. The argument does not rest on the barrier question alone: it rests on the barrier question agreeing with two independent structural measures
ELSTAT's own release cites two different implementing regulations for the same survey: (EU) 2024/1883 in the body and (EU) 2022/1883 in the explanatory notes	Noted, not resolved. It does not affect any figure used here, but it is recorded because this paper grades other people's source discipline

LIMITATION	HANDLING
The two instruments that would make this original are not yet built	Marked blocked, not estimated
The author builds AI products and benefits if AI looks important	The central finding is that adoption is *over-reported*, which cuts against that interest

The materially weakening limitation is the first. With no agreed definition there is no clean total, and this paper deliberately does not produce one. It reports ranges by instrument type and explains why they differ.

13 Citation, licence and interests

HOW TO CITE THIS PAPER

> Broikos, N. (2026) *The adoption gap: what businesses say about AI, what the statistics measure, and why the two disagree.* Version 1.0, 6 August 2026. Athens. Available at: https://broikos.gr/research/the-adoption-gap.pdf (Accessed: DD Month YYYY).

VERSION

Version 1.0, published 6 August 2026. This is the first public version. Substantive changes will increment the version number and be listed in the change log published with the data.

DATA AVAILABILITY

All data underlying this paper is published at https://broikos.gr/research/data/paper-03/ under the same licence as the text. It contains the five Eurostat extracts frozen at the date of writing, the twenty-statistic evidence-appraisal grid, the 61-row claim ledger and the source log.

The claim ledger records the verification status of every figure considered for this paper, **including the four that failed verification and were cut**, with the reason for each cut. Nothing was removed from the record.

No confidential or personal data was used in this paper. Every figure comes from a published statistical release or a published study.

DECLARATION OF INTERESTS

The author builds and sells services involving AI, and therefore benefits commercially if AI is perceived as either more or less capable than it is. That interest cuts both ways in this paper, which argues that adoption is lower and more poorly measured than the loudest sources claim.

Two specific disclosures. The author is a paying user of Anthropic products, and Anthropic's Economic Index appears in the appraisal grid, where it is graded as platform-usage data rather than as a population survey. The second coder used for the reliability check in the companion paper was a commercial

language model. Neither vendor had any involvement in, or sight of, this research.

LICENCE

FUNDING AND INDEPENDENCE

Unfunded. No sponsor, client, trade body or commissioning party paid for, commissioned, reviewed or approved this research. No organisation named in it was given sight of it before publication. This matters because a substantial share of the sources graded here *were* commissioned, and the paper argues that a commissioning interest should be stated wherever it exists.

CORRECTIONS

If a figure here is wrong, or an organisation is described incorrectly, write and it will be corrected with the date and the reason recorded in a public change log rather than silently edited. That log is published at https://broikos.gr/research/corrections.md and already records every figure cut, corrected or downgraded during production. Every correction made during production is already recorded in the claim ledger published with this paper, including the figures that were cut.

Contact: https://broikos.gr/contact · https://broikos.gr

ABOUT THE AUTHOR

Nikolaos Broikos operates e-commerce businesses in Greece, works in web development and digital strategy, and builds agent harnesses. He writes from Athens.

This programme is the record of that work rather than a commentary on it. The e-commerce operations supply the transaction data behind the landed-cost model in *What a Greek online order actually costs*; the harnesses are the instruments measured in *What actually makes an agent harness work*. Where the author's own systems are the subject, that is stated in the first paragraph of the paper concerned as well as in its declaration of interests.

No academic affiliation, no institutional backing, no funding, and no client commissioned any of this. The papers therefore ask to be judged on their published instruments, data and corrections rather than on credentials: every dataset, every claim ledger including the claims that failed, and every retraction is published alongside the text, so a reader who distrusts the author can check the work instead.

THE OTHER PAPERS IN THIS PROGRAMME

Independent, unfunded and published free under CC BY 4.0, each with its underlying data. Read separately; they share a method, not an argument.

1. Who publishes the numbers: a census of corporate research in Greece

   `https://broikos.gr/research/who-publishes-the-numbers.pdf`

   68 Greek publication programmes coded on a published schema: who produces the numbers business decisions rest on, what motivates them, how far the press carries them, and which of them answer engines actually cite.

2. What a Greek online order actually costs

   `https://broikos.gr/research/what-a-greek-online-order-costs.pdf`

   What the €36.1bn e-commerce figure counts and what it does not, the convergence of Greek and European online buying, and a landed-cost model built from published tariffs rather than quoted rates.

3. What actually makes an agent harness work

   `https://broikos.gr/research/agent-harness.pdf`

   Techniques, prompt engineering and measured results from three agent harnesses, two built by the author and one not, across 45 graded sessions on a single model, including three faults found in the measuring apparatus itself.

4. The Greek e-shop technical audit

   `https://broikos.gr/research/greek-eshop-audit.pdf`

   900 Greek domains measured on one day: what Greek online retail actually sends to a browser. 36.8% send no security headers, 11.1% fail a validating TLS handshake, and 5 shops out of 505 pass five elementary checks.

5. The Greek SME digital bill

   `https://broikos.gr/research/greek-sme-digital-bill.pdf`

   What it costs per year to sell online in Greece, priced entirely from published pages: EUR 157.50 at the floor. Of 25 cost lines, 7 have no published price at all, and they are disproportionately the mandatory ones.

6. What an agent actually costs

   `https://broikos.gr/research/what-an-agent-costs.pdf`

   Seven versions of one production agent on the same benchmark and the same model, including the two versions that got worse, the failure taxonomy, and the chart that would have shown a cost explosion that never happened.

STATUS

Version 1.0, August 2026. Complete as a study of the evidence base; two originally planned instruments were not built.

SOURCES OPENED

Six statistics opened at primary source, six carried forward already verified, seven graded from summaries, one not found at the attributed source.

PRIMARY SOURCES

Eurostat's 2025 enterprise AI release and the US Census Bureau's BTOS AI-use analysis, both opened 6 August 2026.

broikos.gr