Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

No need, or no category? Recording refusal of generative AI in European official statistics

Jonas A. Mandalunes

jonas.mandalunes@perpetual.edu.ph

University of Perpetual Help System DALTA

Research Article

Keywords: Generative AI, Non-use and refusal, Official statistics, Survey measurement, Disability and digital exclusion, Classification

Posted Date: August 28th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10694433/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Official European statistics report that most people who do not use generative AI tools simply see no need. We show this conclusion is constrained by the answer categories on offer. Using the 2025 EU survey on ICT use in households and by individuals, the first wave carrying artificial-intelligence questions, we compare two non-use items put to the same population in the same year. The internet item lets respondents tick all that apply from seven reasons, including disability, equipment cost and opposition to the internet in general. The AI item permits one main reason from five, and offers none of those three. Across nineteen countries the three categories with no AI counterpart sum to a median 26 percentage points of the internet non-user base, a multiple-response total, while the AI item's entire residual absorbs a median 4.5%. The gap is positive in every country, retains roughly two-thirds of its size under conservative rescaling, persists in all five overlap strata, and is larger among higher-educated respondents, which composition alone does not explain. We find no evidence that suppressed reasons accumulate in the residual or in "no need". The variable list adopted for 2026 drops the single-answer rule and adds an ethical-concerns category, but still records neither disability nor cost as a reason for refusing AI. We argue that an apparatus unable to record disability or cost as grounds for refusing AI cannot see them, and that policy built on such categories risks mistaking a measurement artefact for public indifference.

1 Introduction

In December 2025 Eurostat reported that 32.7% of people aged 16–74 in the European Union had used generative artificial-intelligence tools in the previous three months (Eurostat 2026e). Attention understandably fell on the two-thirds who had not. Eurostat's own commentary offered a reassuring account of them: asked why, most non-users said that they "simply didn't feel the need to" (64%), with smaller shares reporting that they were unsure how to use the tools (14%), unaware they existed (9%), or concerned about privacy and security (7%).

Read at face value, this is a story about indifference. It implies that Europe's AI divide is not really a divide at all, but an aggregate of individually unremarkable preferences. It implies, further, that there is not much for policy to do: awareness campaigns and skills training might address the 23% who report unawareness or incapacity, and the 7% with privacy concerns present a design problem, but a population that sees no need requires no intervention.

This article argues that the reassurance is partly manufactured by the instrument that produced it, and it argues so on the basis of a comparison that the same instrument makes available. The 2025 EU survey on ICT use in households and by individuals asks two structurally similar questions. One asks people who have not used the internet why not. The other asks people who have not used generative AI why not. The first invites respondents to tick all that apply from seven reasons, among them a physical or mental disability, the cost of connection or equipment, and being opposed to the internet in general. The second permits a single main reason from five, and offers none of those three.

The consequence is a difference not in what Europeans think but in what the European statistical system is built to hear. A person who does not use generative AI because she objects to it on principle has no box. A person who cannot use it because of a visual impairment or a cognitive disability has no box. A person who cannot afford the device or the subscription has no box. Each of them must either choose a reason that does not describe them or fall into a residual category labelled "Other".

We quantify what this costs. Across nineteen European countries, the three categories that have no AI counterpart sum to a median 26 percentage points of the internet non-user base. Because the internet item permits multiple responses, that total over-counts individuals and is not the share of all endorsed reasons; Sect. 5.4 states precisely what it is and what it is not. The AI item's entire residual category absorbs a median 4.5%. The gap is positive in every single country we examine. It survives a conservative rescaling that removes the arithmetic advantage of the multiple-response format. Most importantly, it persists when we restrict the comparison to the population strata in which internet non-users and AI non-users most overlap — two elderly strata, the low-educated, the economically inactive and the rural — and it is larger, not smaller, among the high-educated. Whatever produces it, it is not simply that internet non-users are an older and poorer population than AI non-users.

Two things this article does not do. It does not estimate an instrument effect: the two questions are answered by different subpopulations at very different stages of diffusion, and no design on published aggregates can separate those influences. And it does not show where the suppressed reasons go. We tested two candidate destinations — the residual category and the "no need" option — and found support for neither. The claim we defend is narrower and, we think, more consequential: the apparatus has no place to record certain grounds for refusing AI, and what a statistical system cannot record, a policy system cannot see.

One development belongs at the outset, because it bears directly on how this article should be read. The instrument has since been revised, and the revision is on the public record. Eurostat reports that the single-main-reason rule has become "tick all that apply" (Eurostat 2026g), and the variable list adopted for the 2026 survey year adds one further reason for not using generative AI: ethical concerns, glossed as the perpetuation of stereotypes, absence of accountability, and unauthorised use of intellectual property (European Commission 2025). That is close to the volitional ground we show the 2025 item omitted, and we read it as corroboration of the diagnosis rather than as a rebuttal of it.

Two things the revision does not do. It does not make disability or cost recordable: the same 2026 list keeps an accessibility category and a cost category for internet non-use and gives the AI item neither, so two of the three asymmetries documented here survive into the next wave. And it does not revise the 2025 figures. The 2025 wave remains the only published European measurement of why people do not use generative AI, and it is the baseline against which AI policy will be argued until a comparable wave appears. Section 7.5 sets out what the revision settles and what it leaves standing.

The argument proceeds as follows. Section 2 develops the theoretical frame: classification as infrastructure, and the residual category as the place where a category system stores its own failures.

Section 3 situates the argument in three decades of work on technology non-use, which has consistently insisted that abstention is a social position rather than an absence. Section 4 sets the two instruments side by side. Section 5 describes data and methods, including the two response structures we verify independently. Section 6 presents results, including two failed predictions. Section 7 discusses what follows, and Section 8 states what does not.

2 Classification as infrastructure: the politics of a residual category

2.1 Categories that act

Official statistics are not a window onto a population; they are an apparatus that makes a population governable by making it countable. Desrosières (1998) traced how statistical categories become simultaneously descriptions and instruments of the state that produces them, and Porter (1995) showed how quantification acquires authority precisely by appearing to remove judgement from the act of measurement. Scott (1998) named the mechanism state simplification: to administer a domain, a state must first render it legible, and legibility requires that the domain be recoded into a small set of standardised categories that the administrative apparatus can process. What does not fit the categories is not thereby abolished; it becomes administratively invisible.

Bowker and Star (1999) give this process its most careful treatment. Classification systems, they argue, are infrastructure: they are built, contested, embedded, and then forgotten. Once naturalised, a classification does work that no longer looks like work. It determines what can be said, by whom, in what form, and it does so silently, because the categories present themselves as the natural joints of the world rather than as the outcome of a design process with winners and losers. Star (1999) adds the methodological corollary: because infrastructure disappears into use, it must be studied by deliberately attending to its breakdowns, its exceptions, and its residues.

Two of Bowker and Star's concepts do the heavy lifting here.

The first is *torque*: the friction generated when a person's biography does not align with the classification applied to it, and the person is bent to fit. Torque is not merely inconvenience. It is the experience of being recorded as something one is not, with consequences that follow from the record rather than from the life. A person who has thought carefully about generative AI and declined it on ethical grounds, and who is then recorded as having "no need", experiences a small instance of torque; the statistical system registers her considered position as an absence of interest.

The second is the *residual category* — the "other", the "not elsewhere classified", the box that absorbs what the system was not built to name. Bowker and Star insist that residual categories are diagnostic rather than incidental: they mark the boundary of a classification's imagination. Where a category system is well matched to its domain, the residual is small because little escapes. Where it is poorly matched, the residual swells. Reading residuals is therefore a way of reading the limits of a category system from the inside.

2.2 Why response architecture is a classification decision

Survey methodology has long known that question form shapes answers. Schuman and Presser (1981) demonstrated that offering or withholding an explicit response option changes its endorsement substantially; Tourangeau et al. (2000) formalised the cognitive stages at which such effects arise; and Krosnick (1991) showed that respondents facing cognitive demand satisfice, selecting an acceptable rather than an optimal answer. Smyth et al. (2006) established that check-all-that-apply and forced-choice formats yield systematically different endorsement rates for the same underlying content. Behind all of these sits the trade-off between error and cost that constrains every survey design (Groves 1989), and that in a legally mandated multi-country instrument is unusually tight.

This literature is usually read as a warning about measurement error — a matter of getting the instrument right. We read it differently. When the instrument in question is not a research survey but a legally mandated European statistical instrument, the choice of response architecture is not a technical nuisance. It is a governing decision about which accounts of a citizen's conduct the state will accept. A forced choice does not merely risk error; it enacts a theory of action in which a person has one reason for what she does. A check-all format enacts a different theory, in which conduct is overdetermined. A category list that includes principled opposition concedes that refusal can be a considered stance; a list that omits it does not.

Espeland and Stevens (1998) call the general operation commensuration: the transformation of qualitative difference into a common metric, which necessarily discards the differences that the metric cannot carry. The concern is live in current work on AI measurement specifically: Sioumalas-Christodoulou and Tympas (2025) show that the global metrics used to assess AI are weighted towards technical and economic properties and sit awkwardly beside the societal and ethical priorities that national AI strategies articulate, and they argue for recalibrating the frameworks accordingly. Our case is a narrower instance of the same problem, one level further down: not a mismatch between an indicator suite and a policy agenda, but a mismatch between the answer categories of a single survey item and the range of positions the item is used to summarise. Our contribution is to show that commensuration is not uniform even within a single instrument. The 2025 EU ICT survey commensurates internet refusal generously and AI refusal narrowly, and it does so in the same questionnaire, in the same fieldwork period, to the same population. That asymmetry is the object of this article.

2.3 The expected direction, and why the observed direction is surprising

One might expect a mature measurement domain to be measured more crudely than a new one: settled questions attract less methodological attention. The internet non-use item has been iterated across roughly two decades of European ICT surveys, and its seven categories bear the marks of that iteration — the disability option, the cost option, and the general-opposition option each read as accretions from successive rounds of cognitive testing and stakeholder consultation.

The AI item, by contrast, entered official European statistics for the first time in 2025. It is the newest, most contested, most politically salient technology in the instrument, and it is measured with the thinnest vocabulary of refusal in it. This is the opposite of what a naive account of measurement development would predict, and it is what makes the asymmetry worth explaining rather than merely reporting. A new module under time pressure, with limited questionnaire space and no accumulated cognitive testing, will tend to inherit a minimal category set — and that minimal set then becomes the baseline against which AI-related policy is calibrated for years. That extending ICT usage surveys to artificial intelligence raises questions of category design under constraint is not a new observation; it has been the subject of dedicated methodological attention in the international statistical community (OECD 2021).

This account no longer rests on our inference. In March 2026, after the data analysed here were released, Eurostat published its own statement of why the item took the form it did:

Because of space constraints, the question on reasons for not using generative AI was limited to a single main reason; it has since been expanded to a 'tick all that apply' format, with an additional response option added. The wording of the first question has also been adapted slightly to better reflect how generative AI is currently used. (Eurostat 2026g, p 12)

Two consequences follow, and they pull in opposite directions. The producer of the statistic confirms the constraint we posited, which removes a layer of conjecture from the argument: the narrow vocabulary of AI refusal followed from questionnaire space, not from a judgement that the omitted grounds do not exist. The same sentence, however, records that the response rule has already been revised, so this article's diagnosis is retrospective.

The report does not say which reason was added, but the legal instrument does. The variable list adopted for the 2026 survey year specifies six separate variables for reasons for not using generative AI, one of which is new: alongside not knowing the tools existed, not needing them, not knowing how to use them, concerns about data protection, privacy, security or safety, and a residual, it adds *ethical concerns, such as perpetuating of stereotypes, no accountability, unauthorised use of intellectual property* (European Commission 2025). Two features of that list matter for what follows. The reasons are now separate variables rather than mutually exclusive categories of one, which is the form a tick-all-that-apply item takes and which corroborates Eurostat's statement. And the added category is a recognisable form of principled objection: the ground our analysis identifies as the most consequential omission is the ground the next wave restores. Section 7.5 takes up what that settles, and what it does not.

3 Non-use as a social position

The proposition that abstention deserves explanation in its own right is not new. Wyatt (2003) argued that non-users are constitutive of technologies rather than residual to them, and distinguished four positions that are routinely collapsed into one: resisters who never adopted by choice, rejecters who discontinued voluntarily, the excluded who never had access, and the expelled who lost it involuntarily.

Selwyn (2003) pressed the corollary that non-use is frequently active, informed and rational, and that treating it as deficit misdescribes both the people and the problem.

Digital-inequality research supplied the structural counterpart. Van Dijk (2005) theorised access as a sequence of resource-dependent stages, so that inequalities displaced from physical access reappear as inequalities of skill and usage. Helsper (2012) added that the translation of social resources into digital outcomes is field-dependent, so that identical individual resources yield different outcomes in different institutional contexts. More recently, a systematic review has consolidated the literature on voluntary digital disconnection as a distinct and internally varied practice rather than a residual of adoption (Nassen et al. 2023). Work on online-news non-use has gone further, combining behavioural tracking with automated classification and so treating the measurement of non-use as itself the analytic difficulty (Reiss 2023).

Two implications organise our analysis.

First, the volitional–deprivational distinction is the analytically load-bearing one. Whether someone abstains because she has decided to or because she cannot is the difference between a preference to be respected and an exclusion to be remedied. Any instrument that cannot represent both sides of that distinction cannot support the policy inference that is routinely drawn from it.

Second, the capacity to occupy the volitional side is itself unequally distributed. Declining a technology on principle presupposes knowing it exists, understanding roughly what it does, and having enough standing to treat one's refusal as legitimate. Framing abstention as a considered judgement is, in this sense, a resourced act. This is why our finding that the recording gap is larger among the high-educated is theoretically coherent rather than anomalous: the category that the AI item omits is disproportionately the category that resource-rich non-users would otherwise choose.

The AI item collapses this structure. Of its four substantive options, one ("no need") is volitional, one ("privacy, security or safety concerns") is volitional-but-risk-based, and two ("didn't know they existed", "didn't know how to use them") are deprivational. Absent are the deprivational grounds of disability and cost, and absent is the volitional ground of principled objection. The instrument thus offers a truncated version of exactly the distinction that thirty years of scholarship identifies as the one that matters.

3.1 The present wave of generative-AI divide research, and our departure from it

Research on generative AI has begun to map adoption gaps in familiar terms. Giner-Pérez (2026) examines how financial digital engagement, age and trust structure generative-AI uptake in Spanish household microdata; Henseke (2026) traces the distance between exposure to these tools in European workplaces and their actual adoption; and Jamal Eddine et al. (2026) map the emerging literature connecting generative AI to digital-divide phenomena in higher education.

The closest neighbour to this article is Nakagomi et al. (2026), who analyse an emerging generative-AI divide in terms of personal, positional and resource-based factors, and who, unlike most of this literature, take reported *reasons* for non-use as an object of analysis rather than only adoption itself. The proximity is worth stating precisely, because it marks where our contribution begins. They ask which people give which reasons; we ask which reasons an official instrument is built to accept. Their design treats the reason categories as given and reads variation across respondents; ours treats the categories as the variable and reads them against a sibling item in the same questionnaire. The two are complementary and neither substitutes for the other, but an analysis of reported reasons inherits whatever the category set permits, which is the prior question we take up.

These studies share an analytic posture: non-adoption is an outcome to be explained by individual and contextual characteristics. The posture is productive, and it is not ours. Our question is prior to it. Studies of this kind take their dependent variable from instruments like the one we examine, and we ask whether those instruments are capable of representing the range of positions that non-adopters actually occupy. Where adoption research asks who does not use generative AI and why, we ask what the measurement apparatus permits "why" to mean. The two questions are complementary, but the second constrains the first: no amount of analytic sophistication downstream can recover a distinction the instrument declined to record.

4 Two instruments, one questionnaire

Both items come from the 2025 model questionnaire of the EU survey on ICT use in households and by individuals (Eurostat 2025a), conducted under Regulation (EU) 2019/1700 (European Parliament and Council 2019). We quote both verbatim.

Item B7 (artificial intelligence). Asked of respondents who had not used generative AI tools in the previous three months:

What is the main reason for not using generative AI tools in the last 3 months?

(tick one)

1. a) There was no need
2. b) I didn't know they existed
3. c) I didn't know how to use them
4. d) Concerns about privacy, security, or safety
5. e) Other

Item B10 (internet). Asked of respondents who had not used the internet in the previous three months:

What were the reasons for not using the internet in the last 3 months? **[optional]**

(tick all that apply)

1. a) No need (because the internet is not useful, not interesting, etc.)
2. b) Use of the internet is too difficult
3. c) Physical or mental challenges to using the internet (including physical

disability, visual impairment, dementia, special needs, or other disability)

1. d) Concerns about security or privacy, fear of being scammed or harassed online
2. e) Costs of the internet connection or equipment too high
3. f) I am opposed to the internet in general; I don't trust the information on the

internet

g) Other

4.1 The authoritative specification, and a scope restriction that matters

Eurostat's European compilers' manual for the 2025 survey (Eurostat 2025b, catalogue number KS-01-25-066, 346 pages) specifies both items formally, and its scope statements are consequential:

B7 — *[Scope: individuals who used the internet in the last 3 months, i.e. 1st*

option in B1, and who ticked 'no' in B5] [Type: one single answer needed, i.e. Tick only

one]

B10 — *[Scope: individuals who did not use the internet in the last 3 months]*

[Type: multiple answers allowed, i.e. Tick all that apply]

Two things follow. First, both response rules are officially specified, not inferred; Eurostat's 2026 statistical report reproduces the same five B7 options and the same single-main-reason rule (Eurostat 2026g, p 12), so the item text used here is corroborated by two independent official documents. Second, and more importantly, the two items are asked of mutually exclusive respondent sets: B7 goes only to internet users, B10 only to internet non-users. No individual in the survey answers both. The comparison in this article is therefore, by the survey's own construction, a comparison between two instruments rather than between two accounts given by the same people. We return in Sect. 7.4 to what this filter does and does not justify.

The manual also glosses each option, and three glosses bear on our argument. B7a is defined as *"The respondent was aware of the existence of generative AI tools but did not have the need or desire to use them"* — so the modal category is officially a statement about desire, not ignorance. B7d, printed as "privacy, security, or safety", is glossed only as *"concerns surrounding the security of generative AI tools"*, a narrower reading we note as a comparability caveat. B10f is glossed as *"a general skepticism or*

opposition to the internet, and/or a distrust of the information found online", confirming that this category bundles opposition with distrust.

Table 1 sets out the differences.

Table 1

Response architecture of the two non-use items in the 2025 EU survey on ICT use in households and by individuals. Rows compare the response rule, the framing of the question stem, the scope filter that determines who is asked, the number and content of substantive answer categories, the presence of a residual category, and the number of countries publishing each item. A tick indicates that the category is offered to respondents; "absent" indicates that it is not

Feature	**Item B7 (AI)**	**Item B10 (internet)**
Response rule	Forced single choice ("tick only one")	Multiple response ("tick all that apply")
Framing	"the **main** reason"	"the reasons"
Scope	internet users only	internet non-users only
Respondent sets	mutually exclusive by construction	mutually exclusive by construction
Substantive categories	4	6
Volitional: no need	✓	✓
Volitional: privacy / security	✓	✓ (broader: adds fear of scams and harassment)
Volitional: opposed in principle / distrust	**absent**	✓
Deprivational: difficulty / lack of skill	✓	✓
Deprivational: unaware of existence	✓	absent
Deprivational: disability or impairment	**absent**	✓
Deprivational: cost	**absent**	✓
Residual "Other"	✓	✓
Country coverage in published data	36	24[a]

[a] *The internet non-use item is flagged optional in the model questionnaire, which is why fewer countries publish it. Counts exclude the European aggregates present in each table's country dimension. Only 21 of the 24 pair with the AI item, and 19 support country-level analysis;* Table 2 *documents each step*

Source: Eurostat (2025a, 2025b, 2026a, 2026b*)*

Three points deserve emphasis. First, the asymmetry is not a single design choice but three compounding ones: the response rule, the framing of a singular "main" reason, and the content of the category list. Second, the omissions are not symmetric across the volitional–deprivational axis; the AI item drops one volitional ground and two deprivational ones, so it is not simply a shorter version of the internet item but a differently shaped one. Third, the internet item's privacy option is materially broader than the AI item's, which matters for any direct comparison of that specific category and is a limitation we return to.

5 Data and methods

5.1 Sources

We use two published Eurostat tables, both for reference year 2025 and both released on 17 April 2026:

- isoc_ai_iaiuxr, individuals' reasons for not using generative AI tools (Eurostat 2026a), analysed at unit PC_IND_IUAIX, the percentage of AI non-users;
- isoc_iiu_iuxr, individuals' reasons for not using the internet (Eurostat 2026b), analysed at unit PC_IND_IU3X, the percentage of internet non-users.

Both derive from the same survey, whose 2025 EU fieldwork covered approximately 172,000 households and 330,000 individuals aged 16–74, with reference periods falling predominantly in the first quarter of 2025 (Eurostat 2026f). We use isoc_ai_iaiu (Eurostat 2026c) for national adoption context and isoc_sk_dskl_i21 (Eurostat 2026d) for digital-skill context. The two non-use tables share 103 population-subgroup codes and 24 country codes once European aggregates are set aside, permitting a matched comparison at the level of published aggregate cells.

5.2 Verifying the two response structures

Rather than assume the response structures implied by the questionnaire, we verified them arithmetically across all 960 paired cells. For the AI item, the five category shares sum to a median of exactly 100.0, with a maximum of 100.02 and no cell exceeding 105: a closed composition consistent with a forced single choice. For the internet item, the seven shares sum to a median of 148.0 with a maximum of 346.5: a multiple-response structure in which respondents endorse on average roughly 1.5 reasons, and in some countries close to three. Both figures are computed over the same 960 cells, so they are directly comparable.

The AI distribution has a modest low tail, which we report rather than summarise away. Its minimum is 90.00; 37 of the 960 cells sum to less than 99 and 20 to less than 95. At country level the shortfall is confined to one case: Sweden's national total sums to 92.47, against 99.33 for the next lowest and 99.99 or above for the remaining seventeen. Published two-decimal rounding cannot produce departures of this size, so the likelier explanation is a category value withheld as statistically unreliable, though the published tables do not let us attribute individual shortfalls. This matters less for our design than it

would for most: the AI-side quantity we use is a single published category, the residual "Other", and not a derived composition, so an incomplete AI sum does not propagate into it. It does mean the AI item's closure should be described as near-complete rather than exact.

This verification is not incidental. It establishes that the two tables cannot be compared share-by-share as though they were the same kind of quantity, and it dictates the conservative rescaling reported below.

5.3 Unit of analysis, sample and exclusions

The unit is the published aggregate cell: country × population subgroup, 2025. European aggregates (EU27_2020, EA) are excluded throughout as they are not countries, and no value is imputed at any point: cells are either present or absent.

The sample reduces in three documented steps, set out in Table 2. The AI item is published for 36 countries and the internet item for 24, so 24 is the binding constraint. Of those 24, three do not survive pairing. Ireland and Luxembourg appear in the table's country dimension but publish no internet reason cells for the subgroups we require. Norway publishes six of the seven internet categories but omits the disability category altogether, so no Norwegian cell is complete on both items; we record this as a gap rather than impute it. That leaves 21 countries and 960 paired cells across 80 subgroup codes. The panel is unbalanced, with between 5 and 71 cells per country, because national statistical institutes publish different breakdowns.

Country-level analyses use the IND_TOTAL cells, and two further countries drop out. The Netherlands publishes no IND_TOTAL cell for these items, though it does publish subgroup cells, so 20 of the 21 countries have a country-level row. Finland is then excluded because it reports zero on every internet reason category. Nineteen countries remain for all country-level results.

Table 2
Sample reduction from published coverage to the analysed set. Each step is a documented data availability constraint, not an analytic choice. No value is imputed

Step	Countries	Reason for reduction
AI non-use item published	36	—
Internet non-use item published	24	item is optional, so fewer countries field it
Both items pairable	21	Ireland and Luxembourg publish no internet reason cells; Norway omits the disability category
Country-level (IND_TOTAL) row present	20	the Netherlands publishes subgroup cells but no country total
Analysed at country level	19	Finland reports zero on every internet reason category

Source: authors' calculations from Eurostat (2026a, 2026b*)*

5.4 Constructed quantities

Our primary constructed quantity is the *reason-space with no AI counterpart*: the sum of the internet item's principled-opposition, disability and cost shares.

Its interpretation requires care, and we state the difficulty rather than gloss it. Because the internet item is multiple-response, a respondent who cites two of the three categories contributes to two shares. The sum therefore over-counts people: it is an upper bound on the share of internet non-users citing at least one of the three, and it is not that share. The quantity we would prefer — the proportion of non-users for whom at least one AI-unrecordable reason applies — requires the joint distribution across categories, which published aggregates do not contain. It is therefore not identified from these data, and no rescaling recovers it.

What we report instead is a bracket. The raw sum is an upper bound. The modal-reason rescaling described below is a conservative approximation that removes the arithmetic inflation under an explicit assumption, and we treat it as the figure a sceptical reader should prefer. Both are reported for every country in Sect. 6.5, and no conclusion in this article depends on the raw figure alone.

The *recording gap* is that quantity minus the AI item's "Other" share. It is a comparison of two published quantities from two differently architected instruments, not a difference in a common measure, and we treat it as descriptive throughout.

For the conservative robustness analysis we compute a *modal-reason approximation*, in which each country's internet shares are rescaled proportionally to sum to 100 before the no-AI-counterpart share is recomputed. This is the operation known as closure in the analysis of compositional data (Aitchison 1986). It assumes that a respondent's single main reason would be drawn in proportion to her endorsements — a strong assumption, and one that yields a transparent and deliberately conservative approximation rather than a proven bound.

5.5 Statistical procedures

Given nineteen countries, we report medians with percentile bootstrap confidence intervals (10,000 replicates, resampling countries, seed 20260813), interquartile ranges, exact two-sided sign tests on paired differences, Spearman rank correlations with bootstrap intervals, and leave-one-country-out refits. All procedures were implemented directly rather than through a statistical package, so that every reported number is reproducible from the deposited code without library-version dependence.

We pre-specified six tests. H1: the no-AI-counterpart reason-space is substantial. H2: the AI residual is too small to absorb it. H3: AI "no need" exceeds internet "no need". H4: countries recording more principled internet opposition show larger AI residuals. H5, the decisive test: the gap persists within strata where the two non-user populations overlap. R3: the gap survives modal-reason rescaling. We committed in advance to withdrawing the interpretive claim if H5 failed.

5.6 Claim class

Our claims are *documentary* — what categories the instrument offers, established from the official questionnaire and compilers' manual — and *descriptive* — what each category records, at the level of published aggregate cells. We do not estimate an instrument effect, and the reason is stronger than sampling difference: as Sect. 4.1 shows, the two items are put to mutually exclusive respondent sets by design, so no individual answers both. Internet non-users are roughly 6% of European adults and AI non-users roughly 67%, and the two technologies stand at very different stages of diffusion. No analysis of these published aggregates — and, given the disjoint scopes, no analysis of this survey's microdata either — can separate architecture from population from technology. We use no causal language for any estimated quantity and draw no individual-level inference.

6 Results

6.1 The reason-space with no AI counterpart is large (H1)

Across the 19 analysed countries, the three internet categories with no AI counterpart — principled opposition, disability and cost — sum to a median 26.4 percentage points of the internet non-user base (95% CI [21.6, 44.0]; IQR [21.2, 45.1]). The quantity is a multiple-response total and must be read as one: each of the three is a percentage of internet non-users, and because a respondent may endorse several reasons the sum over-counts individuals, as Sect. 5.4 sets out. It is not the share of all endorsed reasons that these three categories represent, which would require dividing by the total endorsement sum, and it is not the share of non-users citing at least one of them, which the published aggregates do not identify. Eighteen of nineteen countries exceed 10 points and sixteen exceed 20. The range is wide: Greece 7.5, Portugal 90.6.

Figure 1 *Reason-space with no AI counterpart against the whole AI residual category, by country, 2025. Open squares give the summed share of the internet item's three categories that have no AI counterpart, namely principled opposition or distrust, disability or impairment, and cost, expressed as a percentage of internet non-users. Filled circles give the entire AI item's residual category "Other" as a percentage of AI non-users. A grey horizontal line joins the two values for each country, and its length is the recording gap. Countries are ordered from the largest to the smallest gap. Nineteen countries are shown; the Netherlands is absent because it publishes no country-level cell for these items, and Finland is excluded because it reports zero on every internet reason category. Data from* Eurostat (2026a, 2026b*)*

Portugal and Spain are the extremes, and both reflect unusually heavy multiple-response endorsement (Portugal's internet shares sum to 287%). We therefore report the rescaled column alongside the raw one throughout, and note that both countries retain substantial gaps after rescaling (Portugal + 27.6, Spain + 23.9 points).

6.2 The residual cannot hold it (H2)

The AI item's "Other" category has a median of 4.5% across the same countries (95% CI [2.6, 6.2]; range 0.4–12.9). The recording gap therefore has a median of + 21.4 percentage points (95% CI [19.5, 34.6]),

and is positive in 19 of 19 countries without exception (exact two-sided sign test, $p = 3.8 \times 10^{-6}$, the smallest value attainable with 19 concordant pairs).

In none of the nineteen countries analysed is the AI item's residual category remotely large enough to hold what the internet item records under the three missing headings.

6.3 The gap persists in matched strata (H5, decisive)

The principal threat to any substantive reading is that internet non-users and AI non-users are different kinds of people. Internet non-use in Europe is concentrated among the very old and the least educated; AI non-use is the majority condition. If the gap simply reflects that composition, it tells us about populations rather than instruments.

Table 3 and Fig. 2 restrict the comparison to strata where the two non-user populations most overlap, and add three contrast strata where they overlap least.

Table 3

Recording gap by population stratum, 2025. The first five rows are overlap strata, chosen because internet non-users and AI non-users are both well represented in them; the last three are contrast strata, where the two populations overlap least. Column *n* is the number of countries with complete data for that stratum. "No AI counterpart" is the median summed share of the internet item's principled-opposition, disability and cost categories, as a percentage of internet non-users. "AI Other" is the median share of the AI item's residual category, as a percentage of AI non-users. "Gap" is the median paired difference in percentage points, with a 95% percentile bootstrap confidence interval, and *p* is from an exact two-sided sign test

Stratum	**n**	**No AI counterpart**	**AI "Other"**	**Gap**	**95% CI**	**Sign test *p***
Overlap strata						
Aged 65–74	19	26.6	3.6	+ 22.9	[17.3, 36.9]	< 0.0001
Aged 55–74	19	28.5	3.7	+ 22.2	[17.9, 36.5]	0.0001
Low or no formal education	18	38.9	3.5	+ 29.4	[16.2, 40.3]	< 0.0001
Retired / not in labour force	19	31.4	3.4	+ 29.0	[19.2, 38.1]	< 0.0001
Rural areas	17	30.3	4.5	+ 25.4	[14.9, 38.9]	0.0003
Contrast strata						
High formal education	8	43.5	6.2	+ 38.9	[21.9, 74.6]	0.0078
Aged 25–54	14	40.3	5.1	+ 32.3	[20.7, 42.1]	0.0001
Cities	18	34.7	4.6	+ 27.2	[16.3, 40.9]	0.0001

Medians are across countries within each stratum. Source: authors' calculations from Eurostat (2026a, 2026b)

In all five overlap strata the gap is large and its confidence interval excludes zero. The gap therefore remains positive within each of the strata in which the two non-user populations are both well represented, which is not the pattern a composition effect alone would produce. It does not follow that composition contributes nothing: stratifying on published breakdowns controls only the characteristics Eurostat publishes, and residual compositional difference within a stratum cannot be ruled out from aggregate cells.

More instructively, the contrast strata show larger gaps, not smaller: +38.9 among the high-educated and + 32.3 among the working-age. If the gap were an artefact of internet non-users being unusually old and deprived, it should shrink among the young and educated. It grows. This is consistent with the

theoretical expectation of Sect. 3: the category the AI item most conspicuously omits is principled objection, and the capacity to occupy that position is concentrated among the resource-rich.

6.4 The gap survives conservative rescaling (R3)

Rescaling internet shares to sum to 100 reduces the no-AI-counterpart share from 26.4% to 20.8% and the gap from + 21.4 to + 14.7 points (95% CI [11.7, 19.1]), still positive in every country ($p = 3.8 \times 10^{-6}$). Removing the arithmetic inflation of the multiple-response format therefore costs the gap about a third of its size and leaves it positive everywhere. Because closure is an assumption rather than a measurement, this partition should be read as the effect of one specific correction, not as a decomposition of the gap into artefactual and genuine components. Figure 3 shows both versions side by side for every country.

6.5 Two predictions fail

H3: AI "no need" does not reliably exceed internet "no need". We predicted that a forced choice would inflate the single available volitional category. It does not. AI "no need" has a median of 66.6% against the internet item's 61.2%, a median paired difference of + 9.3 points with a confidence interval of [− 12.3, 15.2] and an eleven-to-eight country split (sign test p = 0.65). The prediction is not supported.

H4: national opposition does not predict the AI residual. Austria records both the highest principled internet opposition (36.0%) and the highest AI "Other" (12.9%), which invites a story about tech-sceptical publics. The association does not survive examination: Spearman $\rho = +0.106$ with a bootstrap interval of [− 0.40, 0.56], and leave-one-country-out shows the correlation turning negative ($\rho = -0.052$) when Austria is removed. Figure 4 plots the relationship and identifies Austria as the influential point. This was a single-country pattern and we report it as such.

Table 4 presents the country-level results in full.

Table 4

Recording gap by country, 2025. The first three numeric columns give the internet item's shares for the three categories that have no counterpart in the AI item, each as a percentage of internet non-users. "No AI counterpart" is their sum, which over-counts individuals because the item is multiple-response and is therefore an upper bound. "AI Other" is the AI item's residual category as a percentage of AI non-users. "Gap" is "No AI counterpart" minus "AI Other". "Rescaled gap" repeats that difference after each country's internet shares are proportionally rescaled to sum to 100. All figures are percentages or percentage points, computed in exact decimal arithmetic from the published two-decimal source values and rounded once to one decimal place, half away from zero. Countries are ordered by the raw gap

Country	Opposed	Disability	Cost	No AI counterpart[a]	AI "Other"	Gap	Rescaled gap[b]
Portugal	27.6	23.1	39.9	90.6	4.0	+ 86.6	+ 27.6
Spain	29.7	27.8	13.9	71.4	6.2	+ 65.2	+ 23.9
Bosnia and Herzegovina	16.4	9.1	23.5	49.0	0.9	+ 48.1	+ 30.3
Malta	16.5	21.3	9.9	47.6	3.6	+ 44.1	+ 17.2
Cyprus	25.0	6.0	13.0	44.0	0.4	+ 43.6	+ 21.4
Hungary	21.3	7.0	10.9	39.1	4.6	+ 34.6	+ 18.5
Austria	36.0	5.7	4.6	46.3	12.9	+ 33.3	+ 11.7
Lithuania	7.4	10.4	13.6	31.5	2.6	+ 28.9	+ 19.4
France	12.0	9.8	13.2	35.0	7.9	+ 27.1	+ 15.4
Kosovo	2.5	20.3	0.7	23.5	2.1	+ 21.4	+ 13.6
Poland	8.7	11.4	2.6	22.7	1.7	+ 21.1	+ 13.7
Czechia	4.8	13.9	6.8	25.5	4.5	+ 20.9	+ 14.2

[a] *Upper bound: the internet item permits multiple responses, so a respondent citing two of the three categories contributes to two shares. The share of non-users citing at least one of the three is not identified from published aggregates*

[b] *Sweden's rescaled gap exceeds its raw gap because Sweden's internet shares sum to less than 100, so closure raises rather than lowers them*

Source: authors' calculations from Eurostat (2026a, 2026b*)*

Country	Opposed	Disability	Cost	No AI counterpart[a]	AI "Other"	Gap	Rescaled gap[b]
Slovakia	9.6	2.8	14.0	26.4	6.7	+ 19.6	+ 11.6
Bulgaria	4.5	4.8	12.3	21.6	2.2	+ 19.5	+ 14.7
Belgium	5.7	10.0	4.5	20.2	7.2	+ 13.0	+ 10.9
Romania	4.5	5.3	8.3	18.1	5.1	+ 13.0	+ 9.8
Slovenia	6.4	2.3	6.8	15.5	3.6	+ 12.0	+ 5.6
Sweden	6.7	3.3	10.9	20.9	9.8	+ 11.1	+ 19.1
Greece	3.7	2.9	1.0	7.5	5.0	+ 2.5	+ 0.5

[a] *Upper bound: the internet item permits multiple responses, so a respondent citing two of the three categories contributes to two shares. The share of non-users citing at least one of the three is not identified from published aggregates*

[b] *Sweden's rescaled gap exceeds its raw gap because Sweden's internet shares sum to less than 100, so closure raises rather than lowers them*

Source: authors' calculations from Eurostat (2026a, 2026b*)*

6.6 The disability finding

One column of Table 4 deserves separate statement. A median 9.1% of internet non-users across these 19 countries report that a physical or mental challenge — explicitly including physical disability, visual impairment and dementia — is a reason they do not use the internet. In Spain the figure is 27.8%, in Portugal 23.1%, in Malta 21.3%, in Kosovo 20.3%.

The AI item contains no such category. A person who cannot use generative AI tools because of an impairment cannot say so. In an instrument that will inform European AI policy through the remainder of the decade, disability is not a recordable reason for not using artificial intelligence.

7 Discussion

7.1 A residual that cannot do its job

Bowker and Star's account of residual categories supplies a natural expectation: where a classification is poorly matched to its domain, the residual absorbs the overflow and grows. Our results show something more troubling. The AI item's residual does not grow. It sits at a median of 4.5% while the comparison

instrument records more than a quarter of its reason-space in categories the AI item lacks, and it does so in every country we examine.

A residual category that stays small in the presence of large unnamed content is not performing the function residuals are supposed to perform. Several mechanisms could produce this, and our data cannot distinguish them. Respondents facing a forced choice may satisfice towards a plausible listed option rather than reach for "Other" (Krosnick 1991). Interviewers may steer towards substantive codes. "Other" may carry an implicit pragmatic instruction that one's reason is unusual, which a respondent with a common objection will decline. Or the AI non-user population may genuinely differ.

What we can say is that the standard reassurance about residuals — that whatever the categories miss will show up in "Other" — does not hold here. This is a general lesson for reading official statistics on emerging technologies, and it is not confined to this instrument. Where a new module is fielded with a minimal category set, a small residual should not be read as evidence that the categories are adequate. It may equally indicate that respondents are being absorbed into categories that do not describe them.

7.2 Why the failed predictions matter

Our two failed hypotheses discipline the argument in a way we want to make explicit rather than bury.

We predicted that suppressed reasons would surface as inflated "no need" (H3). They do not, at least not detectably. We predicted that countries with more recorded principled opposition would show larger residuals (H4). They do not. Taken together, these failures mean that we cannot trace the displacement. We cannot point to where a principled objector to generative AI ends up in the published table.

This is a real limitation, and it also sharpens the claim. The article is not about where refusal goes; it is about the fact that the apparatus has nowhere to put it. That is a claim about institutional capacity, which the documentary evidence establishes, not a claim about respondent behaviour, which would require microdata and cognitive testing we do not have. Section 8 states what would be needed.

The H4 failure carries a further methodological warning that we take seriously. The Austrian pattern — highest opposition, highest residual — was exactly the kind of observation that invites a narrative. It dissolved under leave-one-out. Nineteen countries is a small base for cross-national association, and the appropriate reading of our interval [− 0.40, 0.56] is that the question remains open, not that the association is absent.

7.3 Legibility, refusal, and the standing to object

The deeper implication concerns the relationship between measurement and political standing. To be counted as objecting is a precondition of being treated as an objector. Where a statistical instrument records that a population sees no need, the policy response is education or nothing. Where it records that a population objects on principle, the policy response must engage the objection. The categories therefore do not merely describe a distribution of attitudes; they allocate a form of political standing.

This is where the contrast strata bear on the argument, and also where we must be careful about how much they can carry. The gap is largest among the high-educated. That is consistent with the possibility that the omitted volitional category is particularly consequential among groups with the resources to form and articulate a considered position on generative AI, which is what Sect. 3 would lead one to expect. It does not establish it. The stratum gap is a difference between a sum of three internet shares — opposition or distrust, disability, and cost — and a single AI share, so published aggregates do not allow us to attribute the high-education pattern to the opposition component rather than to the other two, and they say nothing about how considered any individual respondent's position is. We therefore offer this reading as hypothesis-generating rather than as a result. The instrument comparison is consistent with a second-order inequality, in which the capacity to refuse is unequally distributed and the category that would record the most articulate form of refusal is also the category omitted; establishing it would require the component-level and individual-level evidence set out in Sect. 8.

For a journal concerned with the shift from algorithmic governance shaping society to society shaping algorithms, this is not a peripheral observation. A public's ability to shape a technology depends in part on its ability to be counted as having a position on it. The 2025 instrument gives European publics four ways to explain why they do not use generative AI, and none of them is that they have thought about it and declined. The 2026 instrument adds one, which is the point rather than a refutation of it: the category was treated as omissible, and then it was not.

7.4 A partial rationale, and its limits

The scope restriction in Sect. 4.1 supplies a defensible rationale for one of the three omissions, and we concede it. Because B7 is asked only of internet users, every respondent has already cleared the most obvious cost barrier: they have a device and a connection. Omitting a cost category is therefore not arbitrary. It is incomplete rather than unjustified — generative AI carries costs of its own, in paid subscriptions and access tiers, which an internet connection does not cover — but a designer could reasonably have judged the category redundant.

No comparable rationale is available for the other two.

Disability. Being an internet user does not remove impairment as a barrier to generative AI. Voice interfaces, dense prompt-and-response layouts, long generated text and rapid interface change each raise accessibility questions that are specific to these tools and independent of internet access. A median 9.1% of internet non-users cite a physical or mental challenge as a reason for not using the internet; the AI item provides no way for an internet-using respondent to say the same of generative AI. The filter removes no part of this barrier.

Principled objection. Here the filter does not merely fail to justify the omission; it aggravates it. Internet users are, on average, the population most exposed to public debate about generative AI, most likely to have encountered its outputs, and therefore most likely to hold a considered position. By restricting B7 to internet users, the survey concentrates its sample in exactly the group where informed objection is most

probable, and then offers that group no category in which to record it. Eurostat's own gloss makes the gap explicit: "no need" is defined as being aware of the tools but lacking "the need or desire" to use them (Eurostat 2025b). The 2025 instrument thus recognises disinclination and withholds recognition from objection. This is the one omission the next specification repairs, as Sect. 7.5 sets out; the wave that produced the figures analysed here does not.

The asymmetry, in other words, is not uniformly indefensible. It is defensible in one place, weak in a second, and self-undermining in the third.

7.5 What should change, and what has already changed

Three changes follow from the analysis, addressed to the statistical system rather than to the population. One of them has already been made, which alters what the other two are for.

The response rule: already revised. Eurostat reports that the forced single choice has been replaced by a "tick all that apply" format (Eurostat 2026g). We would have recommended exactly this, and we record it as agreement rather than advice. It was the most tractable of the three design problems and it has been resolved. Worth noting for other new modules facing the same space constraint: where a full check-all format is not affordable, a "main reason" item followed by a single "any other reasons" probe recovers much of the lost structure at minimal length.

The category set: one of three restored. We identify three grounds with no AI counterpart — principled objection or distrust, disability or impairment, and cost. The 2026 variable list adds exactly one of them. Ethical concerns, glossed as stereotyping, absent accountability and unauthorised use of intellectual property, is a recognisable form of principled objection, and its addition means the volitional omission we treat as the most consequential is no longer an omission from 2026 onward (European Commission 2025).

The other two are not restored, and the same document shows this is not an oversight of drafting economy. The 2026 list continues to offer internet non-users a cost category and an accessibility category, the latter reworded and strengthened into explicit interface terms: no compatibility with screen readers or other assistive technology, cluttered interfaces, no alt text for images. The AI item receives neither. So the asymmetry this article documents is not a transitional feature of a first wave that has since been tidied up. For disability and for cost it is reproduced, deliberately and with the internet wording freshly revised, in the very next specification.

That makes the disability recommendation sharper rather than weaker, and we restate it plainly. Generative-AI tools raise accessibility questions of their own — voice interfaces, dense prompt-and-response layouts, long generated text, rapid interface change — and the 2026 internet wording shows the statistical system is fully capable of expressing them. A category asking whether accessibility barriers are a reason for not using generative AI would require almost no drafting invention, because the sibling item now contains one. Cost remains the weakest of the three priorities for the reasons given in Sect. 7.4, though generative-AI subscriptions and access tiers would justify it.

Design-time benchmarking: unaffected, and reinforced. When official statistics extend to a novel and contested technology, the category set should be benchmarked against the most mature comparable item in the same instrument before fieldwork. Had the AI item been read against the internet item in 2025, the three omissions and the response-rule asymmetry would have been visible at design time rather than after publication. That two of the three were addressed by subsequent revision rather than by design review is the transferable lesson, and the two left unaddressed are precisely the two a benchmarking step would have caught, since the comparison item already contained both and still does.

The sequence deserves stating as a finding rather than a postscript, because what it describes is a trajectory in measurement governance and not a static complaint. The first European measurement of a contested technology used the narrowest vocabulary of refusal in its instrument. Within one cycle the response architecture was opened and an ethical-objection category was added. The two grounds that speak to material exclusion rather than to considered dissent were left outside the AI category set while remaining inside its sibling. On the evidence of these two specifications, what the apparatus learned to hear first was objection, not incapacity.

None of this retires the finding. Revising an item does not revise the data it has already produced. The 2025 wave is the first and, at the time of writing, the only European measurement of generative-AI non-use reasons, and it is the evidence base from which AI-related policy will be argued until a comparable wave is published. A partly corrected instrument ahead and an uncorrectable baseline behind are compatible states of affairs, and the second is what this article documents.

8 Limitations

Disjoint respondent sets. The two items are administered to mutually exclusive groups by design: B7 to internet users, B10 to internet non-users. Internet non-users are approximately 6% of European adults and AI non-users approximately 67%. Our matched-stratum analysis (H5) shows the recording gap is not produced by gross demographic differences between the two sets, but the comparison is between different people in the same demographic stratum, never the same people. Because the scopes are disjoint by construction, no analysis of this survey — including its microdata — can establish how a given individual would answer under the other architecture.

One omission is partly defensible. As Sect. 7.4 concedes, the internet-user filter supplies a reasonable rationale for omitting a cost category. Our critique rests principally on the disability and principled-objection omissions, for which no comparable rationale exists.

Definitional narrowing of the privacy category. The AI item is printed as "privacy, security, or safety" but glossed in the compilers' manual only as concerns about security. Direct comparison of that specific category across the two items should be treated cautiously.

Different diffusion stages. The internet is a mature technology in Europe; generative AI is at an early-mass stage. Reason profiles plausibly differ by stage independently of architecture. We therefore claim

no instrument effect.

No displacement mechanism identified. H3 and H4 both failed. We cannot say where unrecordable reasons go.

Arithmetic non-comparability and double-counting. Multiple-response and forced-choice shares are not the same kind of quantity. Because respondents may cite several internet reasons, our raw headline figure over-counts people and is an upper bound on the share citing at least one AI-unrecordable reason. That share is not identified from published aggregates, which lack the joint distribution across categories. The modal-reason rescaling is a conservative approximation under an explicit proportionality assumption, not a proven lower bound. Both figures are reported throughout, and every substantive conclusion holds under the rescaled version, which is the one a sceptical reader should prefer.

Coverage. The internet item is optional, so only 24 countries publish it, 21 survive pairing with the AI item and 19 are analysable at country level, through the documented steps in Table 2. Two of the losses are themselves telling: Norway omits the disability category from the internet item, and the Netherlands publishes no country total. The Western Balkans appear only through Bosnia and Herzegovina and Kosovo, and the Nordic countries only through Sweden, so coverage leans towards southern and central Europe.

Category non-equivalence. The internet item's privacy option adds fear of scams and harassment, so it is broader than the AI item's. The opposition option conflates principled objection with distrust of information.

Aggregate data. All analyses use published cells. No individual-level inference is available, and ecological interpretation must be avoided.

What the confidence intervals cover. Our intervals resample countries, so they describe how the median gap would vary if a different set of countries had been drawn from the same pool. That pool is not a probability sample: the nineteen countries are those that publish both items in a comparable form, which is closer to a census of the available evidence than to a sample of Europe. The intervals should therefore be read as a summary of between-country dispersion, not as inference to a European population of countries. They also treat each published share as fixed. Eurostat's own first-stage sampling error, which is not published cell by cell for these tables, is not propagated into them, so the true uncertainty around every figure we report is wider than the interval we print. The direction of our results does not turn on this: the gap is positive in all nineteen countries individually, which is a property of the point estimates rather than of the intervals.

Leverage. Portugal and Spain exert strong influence on H1; both are always reported with the rescaled column, and leave-one-country-out is reported for H4.

What would settle the open questions. Microdata access would permit direct examination of who selects "Other" for AI and how those respondents answer other items. A split-ballot experiment offering

the three missing categories to a random half-sample would identify the architecture effect directly. Cognitive interviewing with self-identified AI refusers would establish whether "no need" is being used to carry principled objection. None of these is available from published aggregates, and each is a tractable next step.

9 Conclusion

Europe's official account of why people do not use generative AI rests on an instrument that cannot record three of the reasons its own sibling item records for a comparable behaviour. The three categories the AI item lacks sum to a median 26 percentage points of the internet non-user base, a multiple-response total; the AI item's residual absorbs 4.5%; the gap is positive in every country examined, survives conservative rescaling at roughly two-thirds its size, and persists in every stratum where the two non-user populations overlap — indeed it is larger among the better educated.

We cannot say where the missing reasons go, and we do not claim that the architecture causes what we observe. We claim that principled objection, disability and cost were not recordable grounds for refusing artificial intelligence in the first European measurement of the question, and that this was a design decision rather than a fact about Europeans.

The next specification supports that reading in the most direct way available. The forced-choice rule has been replaced (Eurostat 2026g), and the 2026 variable list adds an ethical-concerns category covering stereotyping, absent accountability and unauthorised use of intellectual property (European Commission 2025). The ground we identify as the most consequential omission is the ground the next wave restores. The other two are not restored: the same list keeps a cost category and a newly strengthened accessibility category for internet non-use and extends neither to generative AI, so a person who cannot use these tools because of an impairment still has no box, and the sibling item shows that the box could be written. None of this alters the 2025 record. For as long as that record is the evidence base, an aggregate finding that most people simply see no need should be read as a statement about the instrument at least as much as about the public.

Declarations

Funding

No funding was received for conducting this study or for the preparation of this manuscript.

Author Contribution

J.A.M. conceived and designed the study, retrieved and analysed the data, implemented and tested the analysis code, produced all figures and tables, wrote the manuscript, and read and approved the final version.

Acknowledgement

The author acknowledges Eurostat for providing public access to the statistical data used in this study. The analysis and interpretation presented in this article are solely those of the author and do not represent the views of Eurostat or any other European Union institution.

Data Availability

The data used in this study are publicly available from Eurostat. The primary datasets were Individuals: reasons for not using generative AI tools (isoc_ai_iaiuxr; reference year 2025) and Individuals: reasons for not using the internet (isoc_iiu_iuxr; reference year 2025), available at https://doi.org/10.2908/isoc_ai_iaiuxr and https://doi.org/10.2908/isoc_iiu_iuxr, respectively. Two additional Eurostat datasets, isoc_ai_iaiu (Individuals: use of generative AI tools) and isoc_sk_dskl_i21 (Individuals' level of digital skills), were used only for contextual information and are available at https://doi.org/10.2908/isoc_ai_iaiu and https://doi.org/10.2908/isoc_sk_dskl_i21. The derived analysis dataset, country-level results, subgroup results, and accompanying codebook are provided with this article as Online Resources 1–4 (ESM 1–4).

References

1. Aitchison J (1986) The statistical analysis of compositional data. Chapman and Hall, London. https://doi.org/10.1007/978-94-009-4109-0
2. Bowker GC, Star SL (1999) Sorting things out: classification and its consequences. MIT Press, Cambridge. https://doi.org/10.7551/mitpress/6352.001.0001
3. Desrosières A (1998) The politics of large numbers: a history of statistical reasoning. Harvard University Press, Cambridge
4. Espeland WN, Stevens ML (1998) Commensuration as a social process. Annu Rev Sociol 24:313–343. https://doi.org/10.1146/annurev.soc.24.1.313
5. European Commission (2025) Annex to the Commission delegated regulation supplementing Regulation (EU) 2019/1700 by specifying the number and titles of the variables for the use of information and communication technologies domain for the reference year 2026. C(2025) 4755 final, Brussels, 18 July 2025
6. European Parliament and Council (2019) Regulation (EU) 2019/1700 establishing a common framework for European statistics relating to persons and households. Off J Eur Union L 261I:1–32
7. Eurostat (2025a) EU survey on the use of ICT in households and by individuals 2025: model questionnaire. Eurostat, Luxembourg
8. Eurostat (2025b) European compilers' manual for statistics on the use of ICT in households and by individuals, 2025 survey. Publications Office of the European Union, Luxembourg. Catalogue number KS-01-25-066

9. Eurostat (2026a) Individuals: reasons for not using generative AI tools (isoc_ai_iaiuxr). Online data code, reference year 2025, released 17 April 2026, accessed 13 August 2026. Eurostat, Luxembourg. https://doi.org/10.2908/isoc_ai_iaiuxr
10. Eurostat (2026b) Individuals: reasons for not using the internet (isoc_iiu_iuxr). Online data code, reference year 2025, released 17 April 2026, accessed 13 August 2026. Eurostat, Luxembourg. https://doi.org/10.2908/isoc_iiu_iuxr
11. Eurostat (2026c) Individuals: use of generative AI tools (isoc_ai_iaiu). Online data code, reference year 2025, accessed 13 August 2026. Eurostat, Luxembourg. https://doi.org/10.2908/isoc_ai_iaiu
12. Eurostat (2026d) Individuals' level of digital skills (isoc_sk_dskl_i21). Online data code, reference year 2025, accessed 13 August 2026. Eurostat, Luxembourg. https://doi.org/10.2908/isoc_sk_dskl_i21
13. Eurostat (2026e) Use of artificial intelligence by individuals. Statistics explained, accessed 13 August 2026. Eurostat, Luxembourg
14. Eurostat (2026f) ESMS metadata: ICT usage in households and by individuals (isoc_i). Eurostat, Luxembourg. Accessed 13 August 2026
15. Eurostat (2026g) The use of artificial intelligence (AI) technologies in the European Union: key results, 2026 edition. Statistical report. Publications Office of the European Union, Luxembourg. Catalogue number KS-01-26-009-EN-N. https://doi.org/10.2785/9221093
16. Giner-Pérez JM (2026) Who adopts generative AI? Financial digital engagement, the age divide and trust as a barrier: evidence from Spanish household microdata. FinTech 5:61. https://doi.org/10.3390/fintech5030061
17. Groves RM (1989) Survey errors and survey costs. Wiley, New York. https://doi.org/10.1002/0471725277
18. Helsper EJ (2012) A corresponding fields model for the links between social and digital exclusion. Commun Theory 22:403–426. https://doi.org/10.1111/j.1468-2885.2012.01416.x
19. Henseke G (2026) From exposure to adoption: generative AI in European workplaces. arXiv preprint. https://doi.org/10.48550/arXiv.2604.18849. arXiv:2604.18849
20. Jamal Eddine R, Gide E, Al-Sabbagh A (2026) Systematised evidence mapping of generative artificial intelligence (GenAI) and digital divide phenomena in higher education. Discov Comput 29:157. https://doi.org/10.1007/s10791-026-10044-w
21. Krosnick JA (1991) Response strategies for coping with the cognitive demands of attitude measures in surveys. Appl Cogn Psychol 5:213–236. https://doi.org/10.1002/acp.2350050305
22. Nakagomi A, Abe N, Tabuchi T (2026) Emerging generative AI divide: personal, positional, and resource-based factors associated with use and reasons for non-use. Telemat Inf 104:102360. https://doi.org/10.1016/j.tele.2025.102360
23. Nassen ML, Vandebosch H, Poels K, Karsay K (2023) Opt-out, abstain, unplug: a systematic review of the voluntary digital disconnection literature. Telemat Inf 81:101980. https://doi.org/10.1016/j.tele.2023.101980

24. OECD (2021) AI measurement in ICT usage surveys. OECD digital economy papers. OECD Publishing, Paris. https://doi.org/10.1787/72cce754-en
25. Porter TM (1995) Trust in numbers: the pursuit of objectivity in science and public life. Princeton. Cited from the first edition; the 2020. Princeton University Press. https://doi.org/10.23943/princeton/9780691208411.001.0001reissue is registered at
26. Reiss MV (2023) Dissecting non-use of online news: systematic evidence from combining tracking and automated text classification. Digit J 11:363–383. https://doi.org/10.1080/21670811.2022.2105243
27. Schuman H, Presser S (1981) Questions and answers in attitude surveys. Academic, New York
28. Scott JC (1998) Seeing like a state: how certain schemes to improve the human condition have failed. Yale University Press, New Haven. Cited from the first edition; the 2020 reissue is registered athttps://doi.org/10.12987/9780300252989
29. Selwyn N (2003) Apart from technology: understanding people's non-use of information and communication technologies in everyday life. Technol Soc 25:99–116. https://doi.org/10.1016/S0160-791X(02)00062-3
30. Sioumalas-Christodoulou K, Tympas A (2025) AI metrics and policymaking: assumptions and challenges in the shaping of AI. AI Soc 40:4655–4670. https://doi.org/10.1007/s00146-025-02181-5
31. Smyth JD, Dillman DA, Christian LM, Stern MJ (2006) Comparing check-all and forced-choice question formats in web surveys. Public Opin Q 70:66–77. https://doi.org/10.1093/poq/nfj007
32. Star SL (1999) The ethnography of infrastructure. Am Behav Sci 43:377–391. https://doi.org/10.1177/00027649921955326
33. Tourangeau R, Rips LJ, Rasinski K (2000) The psychology of survey response. Cambridge University Press, Cambridge. https://doi.org/10.1017/CBO9780511819322
34. Van Dijk JAGM (2005) The deepening divide: inequality in the information society. Sage, Thousand Oaks. https://doi.org/10.4135/9781452229812
35. Wyatt S (2003) Non-users also matter: the construction of users and non-users of the internet. In: Oudshoorn N, Pinch T (eds) How users matter: the co-construction of users and technology. MIT Press, Cambridge, pp 67–80. https://doi.org/10.7551/mitpress/3592.003.0006

Figures

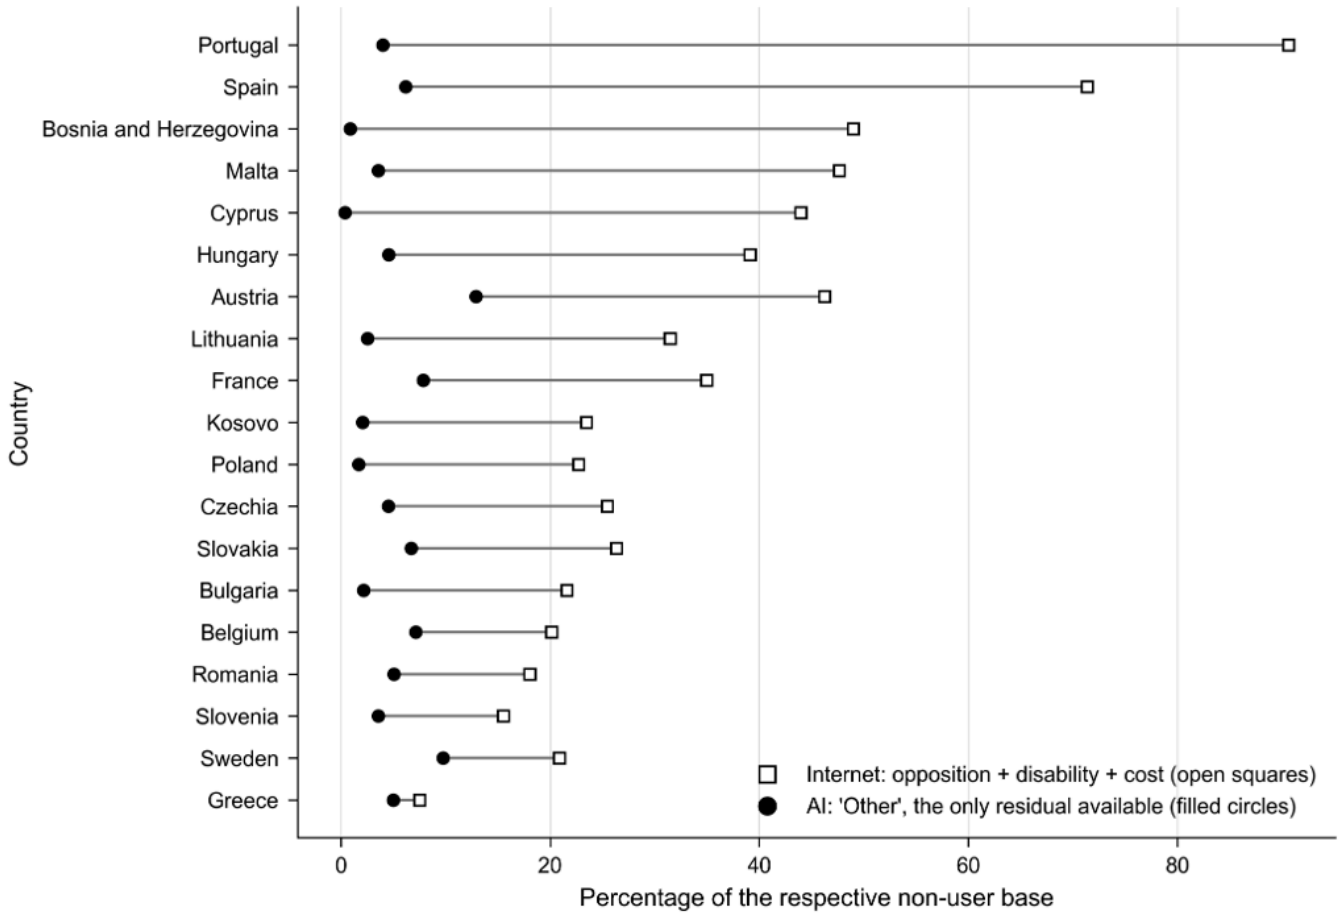


Figure 1

Reason-space with no AI counterpart against the whole AI residual category, by country, 2025. Open squares give the summed share of the internet item's three categories that have no AI counterpart, namely principled opposition or distrust, disability or impairment, and cost, expressed as a percentage of internet non-users. Filled circles give the entire AI item's residual category "Other" as a percentage of AI non-users. A grey horizontal line joins the two values for each country, and its length is the recording gap. Countries are ordered from the largest to the smallest gap. Nineteen countries are shown; the Netherlands is absent because it publishes no country-level cell for these items, and Finland is excluded because it reports zero on every internet reason category. Data from Eurostat (2026a, 2026b)

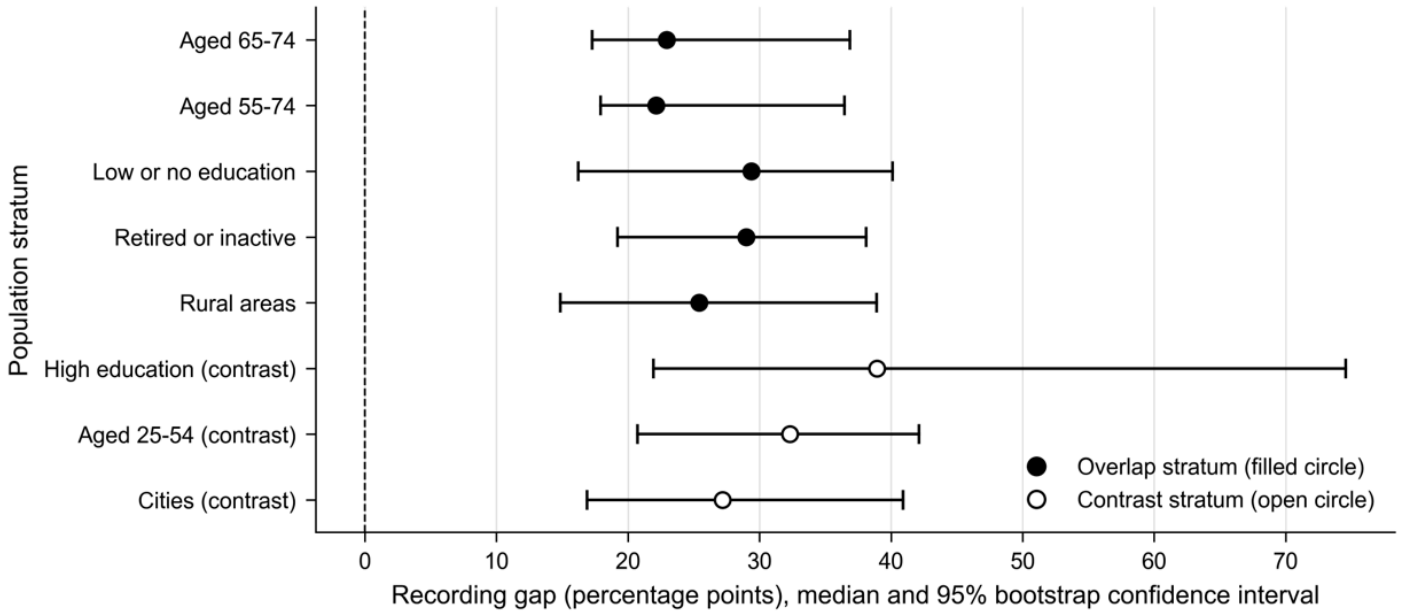


Figure 2

Recording gap by population stratum, 2025, with 95% percentile bootstrap confidence intervals from 10,000 replicates resampling countries. Circles give the median gap in percentage points across countries within the stratum, and horizontal lines with end caps give the confidence interval. Filled circles mark the five overlap strata, where internet non-users and AI non-users are both well represented. Open circles mark the three contrast strata, where the two populations overlap least, and these are labelled "(contrast)". The dashed vertical line marks a gap of zero. Every interval lies wholly to the right of it. Data from Eurostat (2026a, 2026b)

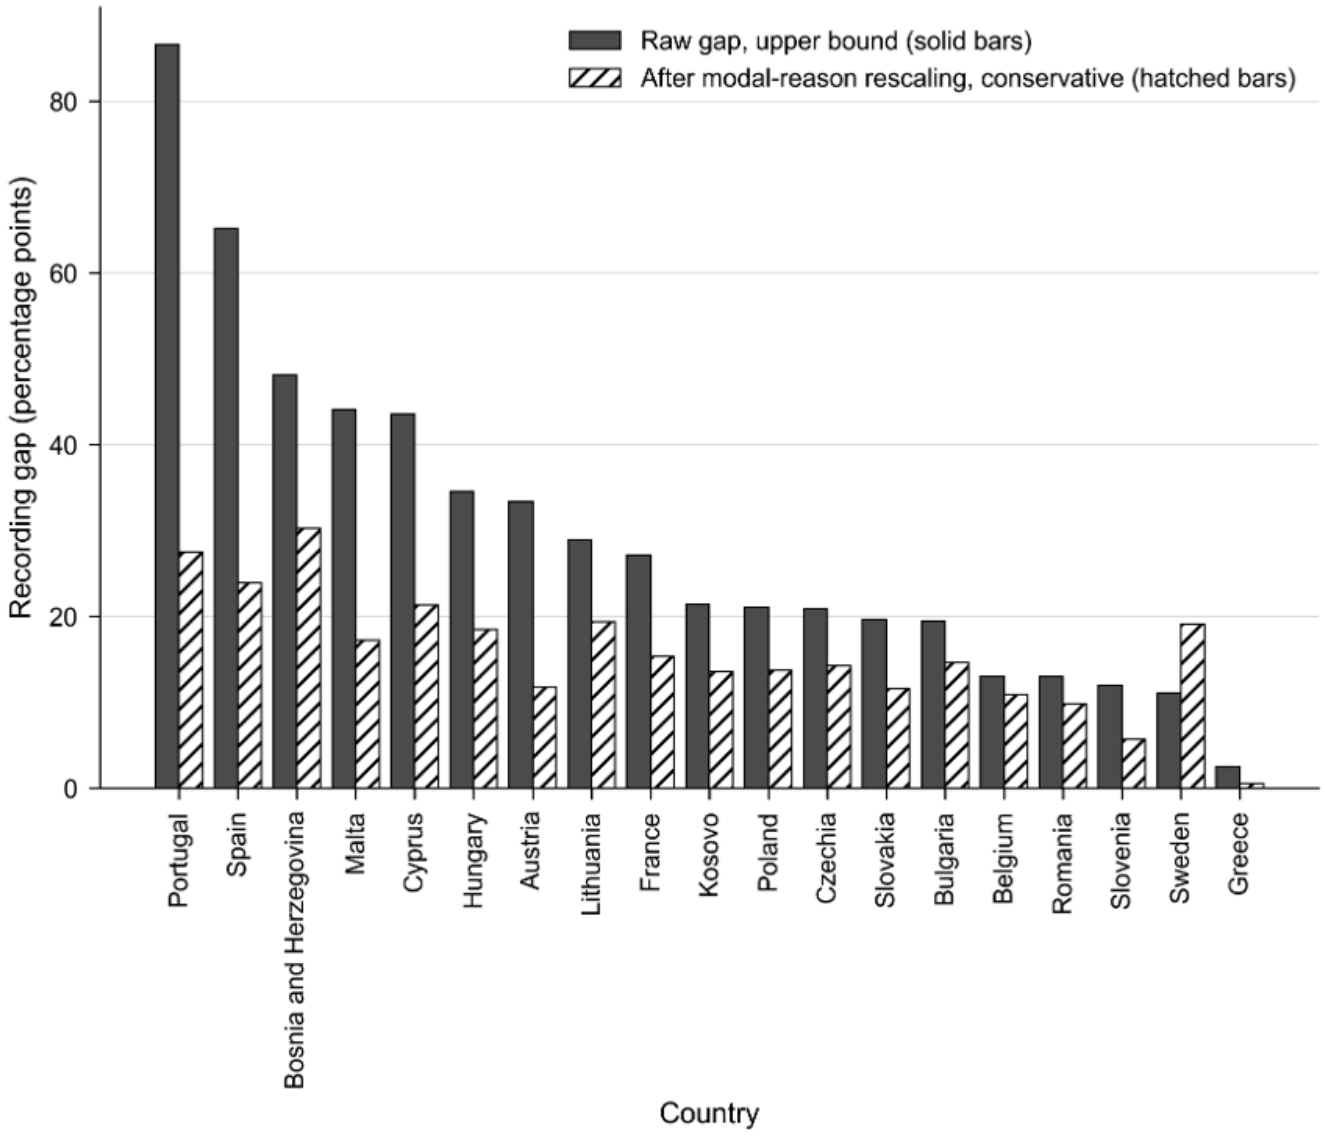


Figure 3

Recording gap before and after conservative rescaling, by country, 2025. Solid bars give the raw gap in percentage points, which is an upper bound because the internet item permits multiple responses. Hatched bars give the gap after each country's internet shares are proportionally rescaled to sum to 100, the closure operation described in Sect. 5.4. Countries are ordered from the largest to the smallest raw gap. Both versions are positive for all nineteen countries. Sweden is the one case where the rescaled gap exceeds the raw gap, because its internet shares sum to less than 100. Data from Eurostat (2026a, 2026b)

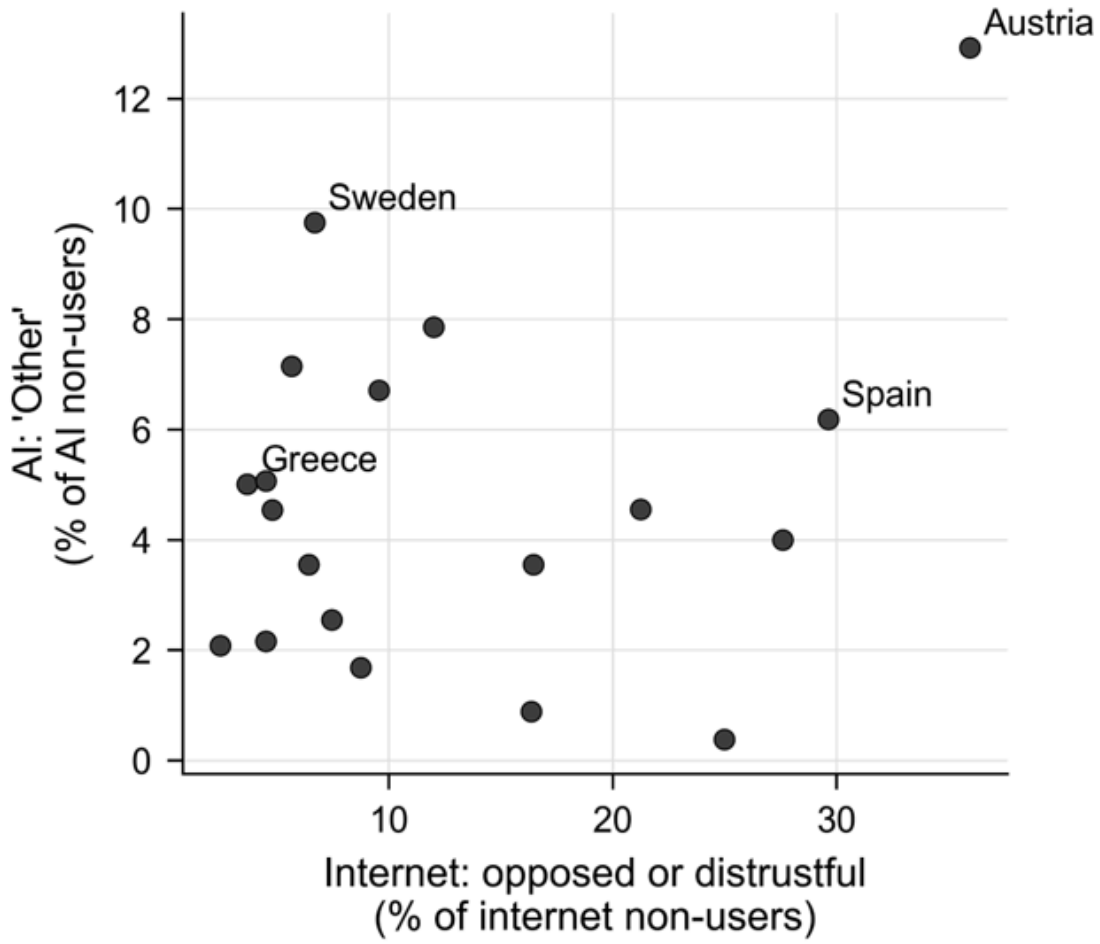


Figure 4

The cross-national association that failed (H4), 2025. Each point is one of the nineteen countries. The horizontal axis gives the internet item's principled-opposition or distrust share as a percentage of internet non-users; the vertical axis gives the AI item's residual "Other" share as a percentage of AI non-users. Austria, Spain, Sweden and Greece are labelled. Spearman ρ = +0.106 with a bootstrap interval of [−0.40, 0.56]; removing Austria turns the correlation negative at ρ = −0.052. Data from Eurostat (2026a, 2026b)

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- ESM3.csv
- ESM2.csv
- ESM4.pdf
- ESM1.csv
- Appendix.docx