



EVIDENCE REVIEW

AI Use Across Generations, Professions and Disciplines

What thirty-three data sources establish about who uses artificial intelligence, who does not, and what the difference actually measures.

September 2026

Version 1.0

Scope, method and a warning about the numbers

This review synthesizes publicly available research published between September 2024 and September 2026 on the distribution of artificial intelligence use across age cohorts, occupations, industries, firm sizes and professional disciplines. It draws on four categories of source: nationally representative government and academic surveys, large-panel opinion research, platform telemetry from AI providers themselves, and professional association and sector research.

The figures that follow are not directly comparable to one another. They differ in the population sampled, in the definition of artificial intelligence supplied to respondents, in the reference period, and in whether they measure any use, weekly use or daily use. Reported adoption on materially similar questions ranges from roughly 18 percent to roughly 89 percent depending on instrument design. As a rule, the highest figures come from vendor and marketing surveys with self-selected samples and should be treated as unreliable; the lowest come from firm-level government instruments that ask a narrow question about institutional deployment. Where divergence is material it is flagged below rather than reconciled.

A second caution applies specifically to the generational material. Almost every widely cited generational statistic is a bivariate comparison. Age is correlated with occupation, education, income, firm size and job level, each of which independently predicts AI use, and several of which predict it more strongly than age does. Any finding reported as a generational difference should be read as a generational association until controls are applied.

1. The population baseline

Adoption is now majority behavior in the United States and has arrived faster than any comparable general purpose technology. Pew Research Center, surveying 5,119 U.S. adults in February 2026, found 49 percent using AI chatbots, up from 33 percent in mid-2024. The Real-Time Population Survey maintained by Bick, Blandin and Deming reports 61.8 percent of adults aged 18 to 64 using generative AI as of May 2026, with 45.2 percent of employed respondents using it for work. Stanford's 2026 AI Index places global population-level adoption at 53 percent within three years of ChatGPT's launch, faster than either the personal computer or the internet reached comparable levels.

The measurement spread is itself a finding. The U.S. Census Bureau's Business Trends and Outlook Survey, which asks firms whether they used AI in a business function during the prior two weeks, returned 19.8 percent as of May 2026. Gallup's second quarter 2026 workforce survey of 22,573 employed adults found 52 percent of workers using AI in their role. These are not contradictory readings. The Census figure measures sanctioned institutional deployment. Gallup measures individual behavior, a substantial share of it unsanctioned. The distance between them is roughly the size of the shadow adoption problem. Direct measures of unsanctioned use are scarce;

the best attributable one is Thomson Reuters’ finding that 34 percent of law firm professionals use AI tools their firm has not authorized.

Intensity remains far below reach. Against 61.8 percent population use, only 6.3 percent of total work hours are spent using generative AI, producing an estimated 2.2 percent time saving across all work hours. Gallup finds 15 percent of workers using AI daily against 52 percent using it at all. Stanford reports 88 percent organizational adoption but fewer than 10 percent of organizations fully scaled in any single business function. The gap between adoption and integration is the dominant structural feature of the current moment.

2. The adoption timeline

Adoption has a shape, and the shape carries information that no single reading does. Three series in this literature are long enough to be read as trajectories rather than snapshots: population use by age cohort, workplace use by quarter, and profession-level use by year. Read together they show a technology that is now growing faster in the cohorts and sectors that started late than in the ones that started early, while the leading edge approaches saturation.

The diffusion benchmark

Stanford’s 2026 AI Index places generative AI at 53 percent global population adoption within three years of mass-market release. The personal computer took roughly fifteen years to reach comparable penetration, the internet about ten, and streaming video about seven. The Real-Time Population Survey makes the same comparison on the narrower measure of workplace adoption and finds that generative AI reached roughly 44 percent work adoption in one third of the time the personal computer required.

These comparisons are directionally sound and numerically soft. They are sensitive to what counts as a launch date and to what counts as adoption, and they compare a technology requiring no hardware purchase, no installation and no infrastructure against two that required all three. The speed is real. The precise multiple is not worth defending.

The cohort series, 2023 to 2026

Pew’s ChatGPT question is the only four-point annual series broken out by age band, which makes it the most useful trajectory data available for cohort analysis.

% of U.S. adults who ever use ChatGPT	All adults	18-29	30-49	50-64	65+
2023	18%	33%	21%	13%	4%
2024	23%	43%	27%	17%	6%
2025	34%	58%	41%	25%	10%

% of U.S. adults who ever use ChatGPT	All adults	18-29	30-49	50-64	65+
2026	44%	61%	55%	37%	19%
Change, 2023 to 2026	+26 pts	+28 pts	+34 pts	+24 pts	+15 pts
Multiple, 2023 to 2026	2.4x	1.8x	2.6x	2.8x	4.8x
Change, 2025 to 2026	+10 pts	+3 pts	+14 pts	+12 pts	+9 pts

Source: Pew Research Center. Note that prior to 2026 the question was asked only of respondents who had heard at least a little about ChatGPT, and in 2026 the question also asked about other chatbots. The 2026 point is therefore not perfectly continuous with the earlier three and will slightly overstate the final-year increase.

The single most important line in that table is the last one. Between 2025 and 2026, adults under thirty grew three points. Adults in their thirties and forties grew fourteen, adults 50 to 64 grew twelve, and adults 65 and over grew nine. The youngest cohort has effectively stopped growing. Every point of national growth from here has to come from older cohorts, from less-educated ones, or from nowhere.

The measurement trap inside that table

The gap between adults 18-29 and adults 65 and over was 29 points in 2023 and 42 points in 2026. Measured as a ratio it was 8.3 to 1 in 2023 and 3.2 to 1 in 2026. The same four numbers show a gap that has widened by 45 percent and a gap that has narrowed by 61 percent, depending entirely on which statistic is reported.

This is not a curiosity. It is a structural property of diffusion curves: absolute point differences widen through the middle of adoption and only close as the leading cohort saturates, while proportional differences narrow from the beginning. A signed point difference is the natural way to express a cohort gap and it is also the statistic that will show the gap growing for several more years in exactly the cases where convergence is underway. Anyone reporting a signed delta needs to have decided in advance whether that is the intended behavior, and needs to be able to answer the reader who arrives with the ratio.

The middle bands show the other half of the pattern. The distance between adults 18-29 and adults 30-49 was twelve points in 2023 and six points in 2026. That gap closed on both measures, because the leading cohort saturated. It is a reasonable model for what the older bands will do next, on a longer timetable.

The workplace series, quarterly

Gallup has tracked U.S. workplace AI use on a consistent instrument since the second quarter of 2023, with sample sizes above 20,000 employed adults in recent waves. It is the densest time series in the literature.

Gallup wave	Any use	A few times a week or more	Daily	Org. has integrated AI
Q2 2023	21%	11%	4%	not asked
Q2 2024	~22%	~12%	4%	33%
Q2 2025	40%	19%	8%	44%
Q3 2025	45%	23%	10%	37%
Q4 2025	46%	rising	rising	38%
Q1 2026	50%	28%	13%	41%
Q2 2026	52%	30%	15%	47%

Organizational integration figures are drawn from differently framed questions across waves and should be read as a rough trend rather than a clean series. The 2024 figures are approximate; Gallup reported that use was essentially unchanged between 2023 and 2024.

Four features of this series are worth naming. Use was flat from 2023 to 2024, which is the strongest available evidence that adoption is not automatic and depends on organizational action rather than on tool availability. Acceleration began in the first half of 2025. Growth stalled to a single point between the third and fourth quarters of 2025 and then resumed, which should discipline any confident extrapolation. And frequent use has grown faster than any use throughout: any use roughly two and a half times its 2023 level, frequent use nearly three times, daily use nearly four times. Adoption is deepening faster than it is broadening.

Organizational integration lags individual use by roughly five points and moved six points in the single quarter to Q2 2026, the sharpest jump in the series. Twenty percent of employees still say they do not know whether their employer has adopted AI at all.

The profession and sector series

Population	Earlier reading	Later reading	Source
U.S. physicians using AI professionally	38% (2023)	81% (2026)	AMA
U.S. physicians using AI in clinical practice	47% (spring 2025)	63% (early 2026)	Doximity
K-12 teachers using generative AI for work	25% (2023-24)	53% (2024-25)	RAND
Lawyers using AI-based tools	11% (2023)	30% (2025)	ABA
Legal organizations actively using generative AI	14% (2024)	26% (2025)	Thomson Reuters
Law firms reporting generative AI use	28% (2025)	41% (2026)	Thomson Reuters
Corporate legal departments	23% (2025)	47% (2026)	Thomson Reuters
EU enterprises using at least one AI technology	8.0% (2023)	20.0% (2025)	Eurostat
U.S. adults 50 and over using generative AI	9% (2023)	30% (2025)	AARP
U.S. district AI training for teachers	23% (2023)	48% (2024)	RAND

The ordering is informative. Medicine moved fastest, roughly doubling in three years, and it is the profession with the clearest task fit, the strongest administrative pain point and institutional procurement of specialized tools. Law moved fast on research and drafting and slowly on judgment, which matches its liability exposure. Education doubled in a single school year on teacher use while district training doubled from a very low base. Enterprise adoption in the EU grew two and a half times in two years and remains at one fifth of firms.

The fastest proportional growth in this entire review belongs to adults 50 and over. AARP recorded 9 percent in 2023, 18 percent in 2024 and 30 percent in 2025, a tripling in two years. Pew’s 65-plus band shows the same shape at 4.8 times its 2023 level. Whatever else the cohort data shows, the trajectory does not describe a population that is failing to adopt. It describes a population that started late.

What the timeline suggests, and what it does not

- The leading edge is saturating. Three points of growth in the youngest cohort in the most recent year, against nine to fourteen in every older band. Absolute cohort gaps should begin closing within one to two years on current trajectories.
- Leading sectors have plateaued. Neither the Information nor the Finance and Insurance sector showed a statistically significant change in AI use between December 2025 and May 2026.
- Depth is outpacing breadth. Daily workplace use has grown nearly four times since 2023 while any use has grown two and a half times.
- The integration gap has not closed at all. Eighty-eight percent organizational adoption against fewer than 10 percent fully scaled in any single function is the same picture as a year earlier.

Against all of that: every series here is short. The longest is four annual observations. One already contains a stall. None spans a full economic cycle or a recession, and none predates the tools it measures by enough to establish a counterfactual. These trajectories are worth reading as descriptions of what has happened. They will not support extrapolation beyond about a year, and any instrument that builds normative bands on them should expect to recalibrate annually.

3. The generational pattern

The direction of the age gradient is consistent across every source reviewed. Its magnitude is contested, and its shape is more complex than the usual reporting suggests. Pew’s February 2026 survey provides the cleanest single view.

Measure (U.S. adults, Feb 2026)	18-29	30-49	50-64	65+
Ever use AI chatbots	66%	61%	42%	23%
Use daily or more	31%	34%	19%	7%

Measure (U.S. adults, Feb 2026)	18-29	30-49	50-64	65+
Ever use ChatGPT	61%	55%	37%	19%
Use chatbots for work tasks (employed)	41%	43%	29%	18%
Extremely or very confident using chatbots	31%	23%	12%	6%
Ever read AI search summaries	72%	68%	59%	38%
Have heard a lot about chatbots	66%	52%	35%	20%

Source: Pew Research Center, *Americans and AI 2026*, n = 5,119, fielded February 17 to 23, 2026.

Three observations follow, none of which survives the usual summary.

The steepest break is not between young and old. Between the 18-29 and 30-49 bands the gap in chatbot use is five points. Between the 50-64 and 65+ bands it is nineteen points. On daily use, adults in their thirties and forties now exceed adults under thirty. Pew notes explicitly that the increase among those in their thirties and forties puts them on par with younger adults. A finding routinely reported as a young-versus-old divide is more accurately a divide at retirement age, with a shallower gradient across the entire working-age population.

Convergence is running from the middle outward. Between 2024 and 2026, chatbot use rose eleven points among adults 18-29, twenty points among adults 30-49, seventeen points among adults 50-64 and eleven points among adults 65 and over. The fastest growth occurred in the two middle bands. On ChatGPT specifically, Pew's longer series shows use among adults 50-64 rising from 13 percent in 2023 to 37 percent in 2026, and among adults 65 and over from 4 percent to 19 percent. Both cohorts have grown faster in proportional terms than the youngest band, which is approaching saturation.


Confidence falls faster than use. Six percent of adults 65 and over describe themselves as extremely or very confident using chatbots, against 31 percent of adults under thirty. But no age group is confident in the aggregate: even among adults under thirty, fewer than a third reach that threshold. Low confidence is a general condition of the current moment, not a cohort-specific deficiency.

The same data mapped to named generations

Pew, like most major survey houses, reports age bands rather than generational labels, and it does so deliberately: generational categories have no agreed boundaries, they shift as cohorts age, and they invite explanation by cohort character when the underlying driver is life stage or occupation. Mapping bands onto labels is therefore a presentational convenience, and the correspondence is close but not exact.


For adults surveyed in 2026, the 18-29 band corresponds almost exactly to adult Gen Z. The 30-49 band is predominantly Millennials but includes the youngest Gen X, those born between 1977 and 1980. The 50-64 band is predominantly Gen X but includes the oldest Boomers. The 65-plus band covers Boomers and the Silent Generation together. The charts below use those labels with the

source bands shown alongside them, and any figure drawn from them should carry the band, not the label alone.



Chatbot use by generational equivalent, U.S. adults, February 2026. Work-task figures are among employed adults only. Source: Pew Research Center, n = 5,119. Labels are approximate mappings of Pew age bands.

Two things in that chart run against the standard telling. Millennials, not Gen Z, are the heaviest daily users, at 34 percent against 31 percent, and they also lead on use for work tasks, at 43 percent against 41 percent. Gen Z leads only on ever-use and on self-reported confidence. The gap between Gen X and Boomers on ever-use, nineteen points, is nearly four times the Gen Z to Millennial gap of five points. And confidence collapses faster than use across the whole distribution: Gen X uses chatbots at 64 percent of the Gen Z rate but reports confidence at 39 percent of the Gen Z rate.



ChatGPT use by generational equivalent, 2023 to 2026. Source: Pew Research Center. Prior to 2026 the question was asked only of respondents who had heard at least a little about ChatGPT, so the final year is not perfectly continuous with the earlier three.

The trajectory chart is the more useful of the two for planning purposes. Gen Z added three points in the most recent year and has visibly flattened. Millennials added fourteen and have nearly closed on Gen Z, from a twelve-point deficit in 2023 to six points in 2026. Gen X added twelve and Boomers nine, both from far lower bases, and neither shows any sign of flattening. The bars that are still growing are the older ones.

What the generation-labeled sources actually say

Where researchers do use generational labels, the instruments are smaller, the definitions vary, and the figures are not comparable with one another or with Pew. They are worth reading for direction rather than level.

Finding	Gen Z	Millennials	Gen X	Boomers
Offered AI training by employer (Randstad)	45%	43%	28%	22%
Offered opportunities to use AI in role (Randstad)	34%	44%	28%	17%
Eager to use AI at work (Randstad)	49%	48%	43%	34%
Report high AI expertise (McKinsey, by age)	50% (18-24)	62% (35-44)	not reported	22% (65+)
Drop in technology confidence, 2025 to 2026 (ManpowerGroup)	not reported	not reported	down 25%	down 35%
Using AI in day-to-day work (Deloitte, 2026)	74%	74%	not surveyed	not surveyed

Sources differ in population, geography, definition and field date and cannot be read across rows as a single series. Deloitte surveyed 14,384 Gen Z and 8,211 millennial respondents across 44 countries and does not cover older generations at all. Randstad figures come from a 2025 release that may have been superseded. McKinsey reports age bands, shown here in their nearest generational position.

The pattern across those rows is the one worth carrying. On every measure of what the employer provides, training and opportunity, the ordering is steeply generational and the drop from Millennials to Gen X is the sharpest break. On measures of what the worker wants, eagerness to use AI, the spread is far narrower: 49 percent for Gen Z against 34 percent for Boomers, a fifteen-point range against a twenty-seven-point range on training provision. Appetite varies less across generations than access does.

The 50-plus population is not a cohort

AARP’s Technology Trends research found generative AI use among adults 50 and over rising from 9 percent in 2023 to 18 percent in 2024 to 30 percent in 2025. Within that band the spread is wide: 47 percent of respondents in their fifties know about and use AI chatbots, against 25 percent of those over seventy. The University of Michigan National Poll on Healthy Aging found 55

percent of adults 50-plus had used AI they speak or type to, while 46 percent reported little or no trust in AI-generated information.

The internal variation within the 50-plus population is larger than the variation between a fifty-two year old and a thirty-four year old. Any instrument or campaign that treats 50-plus, or older workers, as a single analytic unit is aggregating away most of its own signal.

The workplace confidence reversal

ManpowerGroup's 2026 Global Talent Barometer, drawn from interviews with nearly 14,000 workers across nineteen countries, found regular AI usage rising thirteen percent to reach 45 percent of workers while confidence in using the technology fell eighteen percent. The decline was not evenly distributed: baby boomers reported a 35 percent drop in technology confidence and Gen X a 25 percent drop. More than half of the global workforce reported receiving no recent training, and 57 percent reported no access to mentorship. The pattern the study documents is adoption outrunning support, with the confidence cost falling most heavily on the most experienced workers.

4. What the age gap actually measures

The most consequential finding in this literature is that age is not the strongest predictor of AI use, and in several datasets it is not close. The Federal Reserve Bank of New York’s November 2025 Survey of Consumer Expectations, fielded to a representative sample of the U.S. population, found 39 percent of employed respondents using AI tools at work or having done so in the prior twelve months. The variation underneath that figure is instructive.

Cut (employed U.S. respondents, Nov 2025)	Lower group	Higher group	Spread
Education (no degree vs. college graduate)	22.9%	58.7%	35.8 pts
Income (under \$50K vs. over \$200K)	15.9%	66.3%	50.4 pts
Hours (part-time vs. full-time)	24.7%	42.7%	18.0 pts

Source: Federal Reserve Bank of New York, Survey of Consumer Expectations, November 2025 supplement.

A fifty-point income gradient and a thirty-six point education gradient sit alongside an age gradient that, between adjacent working-age bands, is frequently single digits. Age is correlated with all three. An instrument that reports a generational gap without simultaneously reporting education, function, level and earnings will be met, correctly, with the objection that it has measured occupation and called it age.

Access and permission

The distribution data is more consistent than the usage data, and it points in one direction.

- Randstad found AI training opportunities extended to 22 percent of baby boomers and 28 percent of Gen X, against 45 percent of Gen Z and 43 percent of millennials. On opportunities to use AI in role: millennials 44 percent, Gen Z 34 percent, Gen X 28 percent, boomers 17 percent.
- AARP’s survey of workers 50 and over found 12 percent had taken AI training at work, while 49 percent said they were interested in it and 63 percent said employers were not doing enough. Among employers that offer AI training, 97 percent report taking steps to ensure it is accessible and inclusive across ages. Employer-reported availability and worker-reported participation do not reconcile.
- Generation, the global employment nonprofit, found roughly 15 percent of U.S. workers aged 45 and over using AI at work, and that those who do are typically self-taught. Its Age-Proofing AI research reported that 32 percent of U.S. employers would likely consider a candidate over 60 for a role involving AI, against 90 percent for candidates under 35, with a comparable European figure of 33 percent. This last figure was retrieved secondhand and should be verified against the primary report before external use.

Provision is thin across the workforce generally, not only for older workers. The New York Fed found 37 percent of employed respondents saying their workplace does not offer AI tools at all, a further 11 percent saying their employer actively prohibits them, and only 15.9 percent saying their employer offers any AI training, against 38 percent who say such training matters to them. Gallup found 70 percent of employees saying their organization has no guidance or policy for AI use, improving only to 60 percent in white-collar industries. Among teachers, Gallup and the Walton Family Foundation found 68 percent reporting no training on AI tools, and a later wave found roughly six in ten reporting no guidance on applying AI to specific instructional tasks.

The reasonable conclusion is that the organizational failure is general and the generational skew is a second-order effect layered on top of it. Older cohorts are not being excluded from a well-functioning distribution system. They are receiving the smallest share of a system that is failing nearly everyone.

5. Occupation, industry and firm size

The largest and most stable differences in this literature are occupational, not generational.

By collar

Gallup finds 27 percent of white-collar employees reporting frequent AI use, defined as a few times a week or more, an increase of twelve points in a single year. Among production and frontline workers the equivalent figure is 9 percent, and it has moved slightly downward since 2023. Earlier Gallup waves found 81 percent of production and frontline employees never using AI against 54 percent of white-collar workers. This gap is the one visible structural divide in the workforce data that is not narrowing.

By industry and firm size

Sector or size class	AI use rate	Source
Technology (frequent use, employees)	50%	Gallup 2025
Professional services (frequent use)	34%	Gallup 2025
Finance (frequent use)	32%	Gallup 2025
Information sector (firms)	39.7%	Census BTOS, May 2026
Finance and insurance (firms)	33.9%	Census BTOS, May 2026
All U.S. businesses	19.8%	Census BTOS, May 2026
Firms with 250+ employees	37%	Census BTOS, May 2026
Firms with 100 to 249 employees	32%	Census BTOS, May 2026
Firms with 4 or fewer employees	under 20%	Census BTOS, May 2026

Between December 2025 and May 2026, business AI use increased among firms with at least twenty employees and did not change significantly among firms with fewer than twenty. Neither leading sector showed a significant shift over the same period, which suggests an early plateau at the top rather than continued acceleration. The Census working paper on the 2026 AI supplement puts firm use at 18 percent, rising to 32 percent on an employment-weighted basis, and reaching 50 to 60 percent (60 to 70 percent employment-weighted) among very large firms in Information, Professional Services and Finance.

Scope of use is narrow even among adopters. Fifty-seven percent of adopting firms use AI in three or fewer business functions, most commonly Sales and Marketing (52 percent), Strategy and Business Development (45 percent) and IT (41 percent). In 23 percent of firms, 41 percent employment-weighted, workers use AI in work-related tasks.

By job level

Gallup finds one third of leaders, defined as managers who manage other managers, using AI a few times a week or more, against 16 percent of individual contributors. Boston Consulting Group research places the figure for leaders and managers using generative AI several times a week above three quarters. This runs directly counter to a simple age story, since leadership skews older than the workforce it directs. Where organizational position confers both access and permission, the age gradient inverts.

By task

Anthropic's Economic Index, which maps anonymized conversations to the Department of Labor O*NET task taxonomy, found usage concentrated in software development and technical writing. Roughly 36 percent of occupations show AI use in at least a quarter of their associated tasks; approximately 4 percent show it across three quarters. Use is highest for mid-to-high wage occupations such as programmers and data scientists and lower at both the bottom and the very top of the wage distribution. The March 2026 report finds early adoption still uneven, concentrated in high-income countries, in U.S. locations with more knowledge workers, and in a small set of specialized tasks. Geographic concentration inside the United States is converging: the top five states' share of per-person usage fell from 30 percent to 24 percent between August 2025 and February 2026, a pace the Index estimates at roughly ten times faster than historical technology diffusion.

OpenAI's analysis of a representative sample of consumer ChatGPT conversations found work-related messages falling from 47 percent to 27 percent of total volume between June 2024 and June 2025, not because work use declined but because non-work use grew faster. Work usage is concentrated among educated users in highly paid professional occupations, with professional and technical occupations showing both the highest work share and the highest share of advice-seeking rather than task-delegation. Practical Guidance, Seeking Information and Writing account for nearly 80 percent of all conversations. Computer programming is 4.2 percent.

6. Profession-level patterns

Medicine

The American Medical Association's 2026 Physician Survey on Augmented Intelligence (n = 1,342) found 81 percent of physicians using AI professionally, more than double the 38 percent recorded in 2023, with an average of 2.3 use cases per physician against 1.1 in 2023. The most common applications are medical research summarization and clinical documentation. Doximity's 2026 State of AI in Medicine report (n = 3,151 across fifteen specialties) found 94 percent using AI or interested in doing so, 54 percent currently using it in clinical practice, and adoption rising from 47 percent in spring 2025 to 63 percent by early 2026. Thirty-seven percent report daily use.

The age gradient inside medicine is markedly flatter than in the general population. Doximity found current clinical adoption above 61 percent among physicians thirty and under, 57 percent among those in their thirties and forties, and 55 percent among those in their fifties. Adoption is lower among those sixty and over, but only 11 percent of that group reported no interest. Doximity's compensation report puts at-least-weekly use at 73 percent for physicians in their thirties against 52 percent for those in their sixties: a twenty-one point spread across a thirty-year age range within one profession, set against a forty-three point spread between the youngest and oldest bands in the general adult population. When occupation, education and access are effectively held constant by professional membership, most of the generational gap disappears.

Law

The Thomson Reuters Institute's 2026 AI in Professional Services report found 41 percent of law firms and 47 percent of corporate legal departments reporting generative AI use, up from 28 percent and 23 percent respectively a year earlier. Corporate legal departments now outpace the firms that serve them. The American Bar Association's Legal Technology Survey Report is widely cited as showing 30 percent of lawyers using AI-based tools, up from 11 percent in 2023, with a size divide of 46 percent at firms of 100 or more attorneys against 18 percent among solo practitioners. These figures could not be verified against the ABA publication, which is paywalled, and every accessible citation traces to secondary summaries. They are reported with that limitation and should be confirmed before use. Use concentrates in document review (77 percent), legal research (74 percent) and summarization (74 percent), with drafting of briefs and memos at 59 percent.

Governance lags adoption. Thirty-four percent of law firm professionals report using AI tools their firm has not authorized. Only 17 percent of legal professionals say they are ethically comfortable allowing AI to give legal advice, and the share calling AI a major threat to the unauthorized practice of law rose from 36 percent in 2025 to 50 percent in 2026. The profession is adopting rapidly on research and drafting while holding the line on judgment and advice, which is the same boundary the research community draws.

Education

Gallup and the Walton Family Foundation, in a web survey of 2,232 U.S. public school teachers fielded in April 2025, found 60 percent using an AI tool for their work during the 2024-25 school year, 32 percent at least weekly and 28 percent monthly or less. Weekly users reported saving an average of 5.9 hours per week. RAND found teacher use doubling from 25 percent to 53 percent between the 2023-24 and 2024-25 school years. Adoption varies by subject and setting: RAND found English language arts and science teachers roughly twice as likely to use AI as mathematics or general elementary teachers, and Gallup analysis shows high school teachers at 66 percent, early-career teachers at 69 percent, suburban schools at 65 percent, urban at 58 percent and rural or town at 57 percent.

The binding constraint is instructional rather than technological. Sixty-eight percent of teachers reported receiving no training on AI tools, and 82 percent reported receiving no formal guidance on applying AI to their work across multiple task types. RAND separately reports that most students were never explicitly taught how to use AI, a figure widely quoted as over 80 percent but not verified here against the RAND publication. Adoption of AI detection tools among K-12 teachers rose from 38 percent to 68 percent in a single school year, which is a measure of institutional posture as much as of practice.

Research and the sciences

Nature Research Intelligence and Fudan University surveyed 10,480 researchers worldwide in March 2026, with invitations deliberately not mentioning AI in order to reduce selection bias. Two findings bear directly on the questions in this review.

First, AI use concentrates on information-intensive and text-intensive work and thins out sharply where professional judgment and accountability are at stake. Nearly half of respondents use AI every time or most times for information gathering (43.8 percent combined) and for improving or editing papers (44.6 percent combined). Fewer than a quarter do so for reviewing papers (23.5 percent). The report's own conclusion is that AI is reinforcing the exploratory and generative stages of research while leaving tasks resting on accountability and professional judgment largely untouched.

Second, frequent use declines with seniority, and the report attributes this specifically to senior researchers opting out rather than to lighter use among senior researchers who do engage. For paper editing, the frequent-user share falls from 46.8 percent among researchers with fewer than three years of experience to 36.1 percent among those with twenty or more years. For information gathering, it falls from 47.8 percent to 33.4 percent. The magnitude, roughly eleven to fourteen points across a twenty-year experience range, is comparable to the physician finding and again much smaller than the general population gap.

Funding is instructive on the access question. Only 11.1 percent of researchers report that their institution pays for their primary AI tool. Forty-five percent use it free of charge and 41.3 percent

pay for it personally. Heavy users remain more likely to pay for themselves even after controlling for research experience, sector and discipline. In this population, institutional provision is not what drives intensive use.

Discipline-level differences are real but less dramatic than task-level differences. A separate Anthropic survey of 1,260 quantitative social scientists in early 2026 found 81 percent had tried AI chatbots in research, principally for writing code and editing prose, while only 20 percent had adopted autonomous coding agents. The gap between trying a tool and delegating work to it is the same gap visible in the enterprise data.

7. What adoption costs, and what it returns

Cost enters this subject at four layers, and the layer that gets discussed is the smallest one. The license price is the visible number. Total cost of ownership is roughly half again as much. Training is a larger cost than either, and organizational redesign is larger than all three combined. The evidence quality declines sharply as the layers get more expensive, which is the opposite of what a buyer needs.

Layer one: the license

Tier	Reported price	Basis
Consumer free tier	\$0	Published by all major vendors
Individual paid tier	about \$20 per month	Published; consistent across vendors
Individual premium tier	\$100 to \$200 per month	Published
Team or business seat	about \$20 to \$30 per user per month	Published; seat minimums apply
Microsoft 365 Copilot, enterprise add-on	\$30 per user per month, annual term	Published list price
Same, all-in with required base license	reported \$66 to \$90 per user per month	Third-party analysis
ChatGPT Enterprise	not published; reported \$45 to \$75	Procurement analyses; quote-only
ChatGPT Enterprise entry commitment	reported near \$108,000 per year	Implied by reported 150-seat minimum

Reliability warning. Only the Microsoft list price and the published consumer tiers are firm. Enterprise pricing in this market is negotiated and unpublished, and nearly every figure circulating publicly originates from vendor-adjacent blogs, licensing consultancies or competitor comparison pages rather than from the vendors. Treat the enterprise figures as orders of magnitude and confirm any of them directly with the vendor before they appear in a budget or a proposal.

One published total cost of ownership model, for a 200-seat enterprise deployment, puts the effective year-one all-in cost at roughly \$89 per user per month against a \$60 license, about 48 percent above the license alone. The additional categories are integration, administration, security review and support. That ratio is worth carrying as a rule of thumb even though the underlying model is a single vendor-adjacent estimate.

Layer two: who is actually paying

For most individual users today, the tool cost is zero or about twenty dollars a month, and it is not being paid by an employer. Among the 10,480 researchers in the Nature and Fudan survey, 45.4 percent use their primary AI tool free of charge and 41.3 percent pay for it personally. Only 11.1 percent report institutional funding, and heavy users remain more likely to pay for themselves even after controlling for experience, sector and discipline.

The workplace picture matches. The New York Fed found 37 percent of employed respondents saying their workplace does not offer AI tools at all and a further 11 percent saying their employer actively prohibits them, which together means close to half of American workers have no sanctioned employer-provided tool. Stanford estimates U.S. consumer surplus from generative AI at \$172 billion annually, with most tools still free.

This reframes the access question. The barrier is not license affordability. At twenty dollars a month, and with capable free tiers available in every market, price is not what separates a user from a non-user in an advanced economy. What separates them is time, permission, instruction and a task worth applying the tool to. An access argument built on cost will be easy to rebut. An access argument built on permission and training will not.

Layer three: training, and what workers say it is worth

The New York Fed put a price on AI training by asking workers to trade salary for it, which makes it the most rigorous cost estimate in this entire review.

Measure (employed U.S. respondents, Nov 2025)	Value
Workers saying AI training is important to them	38%
Workers whose employer offers any AI training	15.9%
Average salary workers without training would give up to get it	11.4%
Median willingness to pay among that group	0%
Share of that group with zero willingness to pay	61%
Share of that group willing to give up more than 10% of salary	19%
Average raise workers with training would require to give it up	24.2%
Median raise required among that group	15%
Share of that group requiring more than 10%	53%

Read the two averages together. Workers who have employer-provided AI training price it at roughly a quarter of their salary. Workers who do not have it price it at roughly a ninth. Fewer than one in six employers provide it. The New York Fed attributes part of the gap to loss aversion and part to selection, since AI tools are more valuable in the kinds of jobs that also train people to use them, and both explanations are plausible. Either way, a benefit that its recipients value at a 24 percent salary premium, and that costs a fraction of that to deliver, is not being delivered.

The distributional detail matters for any cohort argument. Willingness to pay for training is significantly higher among younger workers, non-white workers, workers without a college degree and workers with less than a year of tenure. Workers without a college degree also require a significantly larger salary premium to give training up. The people who value training most are, in

the New York Fed’s own summary, among those with the lowest usage rates and the lowest access to employer-provided training.

The honest counterweight sits in the same dataset. Sixty-one percent of workers without access to training would give up nothing at all for it, and among the roughly 60 percent of employed respondents who do not consider AI training important, the most common reason by far is that they do not expect to use AI in their jobs, at 48.5 percent. A further 21.1 percent do not think AI will change their industry. This is the same relevance judgment that dominates the European non-use data, appearing here in an American workplace sample. Demand for training is real, concentrated and unmet. It is not universal, and a business case that assumes it is will overstate its market.

Layer four: the return, which is genuinely contested

Published estimates of what organizations get back range from near-total failure to nearly five dollars per dollar spent. They are not measuring the same thing, and the spread is mostly definitional rather than empirical.

Finding	Source	What it measures
95% of organizations saw no measurable P&L impact	MIT Project NANDA, 2025	P&L impact within 6 months of pilot
89% report no impact on labor productivity	Yotzov et al., NBER WP 34836, Feb 2026, n = ~6,000	Sales per employee, prior 3 years
80% report no tangible enterprise-level EBIT impact	McKinsey State of AI	EBIT attribution
56% of CEOs report no financial benefit	PwC 2026	Self-reported financial benefit
39% attribute any EBIT impact; 6% attribute 5% or more	McKinsey State of AI	EBIT attribution
~74% of firms that measure ROI report it positive	Wharton Human-AI Research and GBK Collective, Oct 2025, n = 800+	Among the 72% that track ROI
\$3.50 to \$4.60 returned per dollar spent	IDC/Microsoft, IBM, Accenture	Vendor-sponsored; discount accordingly
15% productivity gain in customer support	Brynjolfsson, Li and Raymond, QJE 2025, n = 5,172	Issues resolved per hour
26% productivity gain in software development	Stanford AI Index 2026	Task-level

The 95 percent figure deserves specific correction, because it is the most quoted number in this field and it does not say what it is quoted as saying. It comes from a preliminary working paper produced by a group building agentic AI infrastructure, based on roughly 52 executive interviews, 153 leader surveys collected at industry conferences, about 350 employee surveys and

analysis of 300 public deployments. Its own funnel reports that 60 percent of organizations investigated enterprise AI tools, 20 percent reached pilot and 5 percent were successfully implemented, which means roughly three quarters of organizations that actually ran a pilot did not clear the bar, not 95 percent. The bar was deliberately demanding: deployment beyond pilot, measurable KPIs, and ROI assessed six months after the pilot, a window that will undercount any enterprise implementation with a normal payback period. Wharton professor Kevin Werbach and others published detailed methodological objections within days. Anyone citing the 95 percent in a client-facing document should expect to be corrected by someone who read the paper.

The counterweight most often cited is Accountable Acceleration: Gen AI Fast-Tracks Into the Enterprise, the third annual AI Adoption Report from Wharton Human-AI Research and GBK Collective, released 28 October 2025. It surveyed more than 800 enterprise decision-makers at U.S. companies with revenue above \$50 million and found roughly three in four reporting positive returns on generative AI investment. Three qualifications belong with that figure. The denominator is narrow: only 72 percent of respondents track structured ROI metrics at all, so the 74 percent describes returns among those measuring rather than among all firms. The sample is confined to large U.S. enterprises and is not representative of the business population. And GBK Collective is a commercial analytics consultancy, which places the study nearer the sponsored end of the spectrum than a purely academic one. The divergence from MIT is also definitional: Wharton counted productivity, speed and quality effects, MIT counted profit and loss. Secondary accounts of the Wharton seniority split contradict one another, reporting the VP-to-manager gap as both 81 against 69 percent and 45 against 27 percent on differently worded items, so no figure for it is given here.

The reconciliation is straightforward and unsatisfying. Task-level experiments measure whether an individual gets faster, and they consistently find yes. Firm-level surveys measure whether the profit and loss statement moved, and they consistently find not yet. Both results can be true simultaneously, and the integration gap documented throughout this review is the reason: Stanford records 88 percent organizational adoption against fewer than 10 percent fully scaled, and Microsoft finds only 16 percent of its sample of existing AI users have redesigned their workflows around the technology. Microsoft's own modelling attributes roughly two thirds of the variance in AI impact to organizational factors such as culture, manager behavior and talent practices, against roughly one third to individual mindset or capability, though this is vendor research and the underlying model is not published. What no one has published is a credible firm-level return estimate with a clean counterfactual. Vendor-sponsored figures in the three to five dollar range should be read as marketing until an independent replication exists.

Costs the license price does not capture

- Unsolicited use. Thomson Reuters found 34 percent of law firm professionals using AI tools their firm has not authorized, and the MIT NANDA report described a broader shadow AI economy of employees using consumer tools without employer approval. No representative

measure of the practice exists. The confidentiality and data-governance exposure it creates is priced at zero on the license and is not zero in fact.

- Absent governance. Gallup found 70 percent of employees saying their organization has no guidance or policy for AI use, improving only to 60 percent in white-collar industries. Policy work is a real cost that arrives after procurement, not before.
- Verification labor. Among more than 7,500 open-ended researcher responses in the Nature survey, accuracy and hallucination concerns drew 4,935 mentions, more than any single hope. Output that must be checked transfers time from generation to verification rather than eliminating it.
- The learning trough. The New York Fed found 19 percent of workplace AI users saying tasks actually take longer because they are still learning, against 40 percent saying tasks finish faster and 22 percent saying they complete more. Early-period negative productivity is a real and predictable cost, and it is one reason a six-month ROI window is the wrong instrument.
- Judgment risk. Deloitte identified a distinct overconfidence risk among younger, technically stronger cohorts. Errors that pass unreviewed are a cost that appears in a different line item than the one the tool was bought under, if it appears at all.

What the cost evidence implies

The cheapest intervention has the strongest evidence behind it. Licenses are inexpensive and widely available free. Training is valued by its recipients at up to a quarter of salary and is provided by fewer than one employer in six. Redesign is expensive and is where the unrealized return sits. Spending sequenced in that order, permission and training before broad licensing, and licensing before redesign, matches the Index's own diagnostic ordering, which puts Access and Permission ahead of Fluency, and it now has an economic justification as well as a diagnostic one.

An instrument should measure conditions, not promise returns. The return literature spans from 95 percent failure to 460 percent return, and every point on that range has a methodological objection attached to it. Any diagnostic that attaches a financial promise to its output inherits the credibility problem of both extremes. Measuring the conditions that the evidence does support, access, permission, training provision and integration depth, is a defensible position. Forecasting the dollar return is not, and the temptation to do it in a sales conversation should be resisted precisely because the buyer has probably already been burned by someone who did.

8. Findings that complicate the obvious readings

Two narratives dominate commentary on this subject. The first holds that older cohorts are behind. The second holds that older cohorts are excluded rather than deficient. The evidence complicates both, and an objective report should say so.

Against the deficit reading

Technical skill and judgment separate. Deloitte Access Economics, in research for RMIT Online covering the Australian workforce, found that 84 percent of workers have used at least one AI tool but only 7 percent reach advanced literacy and over half remain at beginner level. Gen Z and millennials scored higher on technical AI skills. The report identifies a distinct overconfidence risk in those cohorts. Baby boomers performed relatively better on judgment-based skills while remaining more hesitant to use AI at all. The report's framing is that the gap is not principally about frequency of use but about whether a user can deploy the tool safely, effectively and strategically.

Skepticism runs young. Pew found adults under thirty the most likely of any age group to expect AI to affect society negatively (48 percent) and to affect them personally negatively (37 percent). They are evenly split on whether chatbots help or hurt their own creativity, 25 percent against 20 percent, while every older band tilts positive. The heaviest users are the most pessimistic. Separately, Pew reported in August 2026 that young adults in the U.S. are increasingly wary of AI and concerned about job displacement, a concern the Stanford AI Index gives empirical grounding: employment for software developers aged 22 to 25 has dropped nearly 20 percent since 2024.

Seniority correlates positively with use in organizational settings. Gallup's finding that leaders use AI at roughly twice the rate of individual contributors runs opposite to the population age gradient. Where position confers access and permission, older workers use AI more, not less.

Demand exceeds supply among older workers. Forty-nine percent of workers 50 and over say they want AI training. Twelve percent have received it. That is a provision constraint, not a motivation constraint.

Against the pure exclusion reading

Some of the gap is elective. AARP research finds trust, privacy and data security cited by roughly half of adults 50-plus as reasons for hesitating to try AI tools, and the University of Michigan poll found 46 percent of adults 50-plus reporting little or no trust in AI-generated information. Comfort varies sharply by application: among adults 50-plus, 54 percent are comfortable with AI handling administrative billing documentation and 53 percent with an AI scribe, but only 38 percent with AI assisting a diagnosis or treatment recommendation. That is a discriminating pattern, not an undifferentiated aversion, and it is difficult to characterize as a capability deficit or as pure exclusion. The Nature finding that senior researchers opt out is consistent with a considered judgment as well as with structural exclusion. The published data do not adjudicate between them.

Measured productivity gains have been largest for the least experienced. The most cited field experiment is Brynjolfsson, Li and Raymond, published in the Quarterly Journal of Economics in May 2025. It studied the staggered rollout of a generative AI assistant to 5,172 customer support agents at a Fortune 500 firm and found issues resolved per hour rising 15 percent on average. The distribution matters more than the average: less experienced and lower-skilled workers improved on both speed and quality, by roughly 30 percent on issues resolved, while the most experienced

and highest-skilled workers saw small gains in speed and small declines in quality. The Stanford AI Index cites this work at 14 percent, the figure in the 2023 working paper rather than in the published version, and pairs it with a reported 26 percent for software development. If generative AI compresses the value of accumulated experience in some task categories, an argument that experience is the advantage cannot rest on productivity evidence alone. It has to rest on judgment, restraint and accountability, which is precisely where the measurement is thinnest.

Within-profession gaps are real, if modest. The twenty-one point spread between physicians in their thirties and physicians in their sixties, and the eleven to fourteen point spread across researcher experience levels, are much narrower than the general population gap but they are not zero. Controlling for occupation and access shrinks the generational difference substantially. It does not eliminate it.

9. What the evidence does not establish

- No dataset reviewed here demonstrates a generational difference in the capacity to learn or apply AI tools. Every generational finding measures use, self-reported confidence, self-assessed skill, or employer-provided access. Capability, as distinct from these, is not directly measured anywhere in the public literature.
- No dataset reviewed here isolates age from occupation, education, income, firm size and job level simultaneously. The New York Fed data strongly suggests that the age coefficient would shrink under those controls. That regression is not published in any source reviewed.
- The direction of causation between access and use is not established. Employers may withhold tools from cohorts they judge less likely to use them; cohorts may use tools less because the tools are withheld; both may operate at once. No study reviewed here disentangles them.
- Judgment, restraint and accountability, the domains where older cohorts appear to perform relatively better, are the least measured constructs in this entire literature. The Deloitte and Nature findings are the two substantive data points. Both are single studies, one national and one self-reported. This is the largest open space in the evidence base.
- Longitudinal evidence is almost entirely absent. Nearly every source is a repeated cross-section. Whether the observed gaps are cohort effects that persist or period effects that close as tools become ambient is an open empirical question.

10. Implications for instrument design

Four conclusions follow for anyone building a measurement instrument in this space.

Report bands, never binaries. The 50-plus population is not a cohort. AARP's twenty-two point internal spread between the fifties and the seventies is larger than several of the gaps routinely reported as generational. Collecting continuous age and reporting in bands preserves this; collecting in bands destroys it permanently.

Control variables are the credibility test, not a methodological nicety. The single strongest objection any sophisticated reader will raise to a generational finding is that it has measured occupation and called it age. The New York Fed gradients make that objection credible in advance. An instrument that collects function, tenure and level, and reports them alongside every cohort delta, answers the objection before it is made. An instrument that does not will be dismissed by exactly the audience it needs.

The judgment finding is thinly evidenced, which is what makes it valuable. Two published studies support the proposition that experience-linked judgment holds up better than tool recency: the Deloitte literacy work and the Nature seniority and task-type findings. Both are suggestive rather than definitive. An instrument that measures judgment and accountability directly, at scale, across organizations, enters a field with almost nothing in it. That is a genuine research position, not a rhetorical one.

Access and permission are the domains with the most independent corroboration. Randstad, AARP, Generation, ManpowerGroup, Gallup and the New York Fed all document distribution failures separately and by different methods. These are the domains where an organization's own data is most likely to contradict its own narrative, and where the finding will be hardest to argue with.

One caution. The corroborating evidence for a distribution-failure thesis is strong on access and training and weak on outcomes. The productivity literature currently points the other way, toward larger gains for less experienced workers. A thesis that rests on experience as advantage should be argued on judgment, restraint and defensible decision-making rather than on throughput, and should acknowledge the productivity finding rather than route around it. Acknowledging it is also the stronger position: if AI compresses the value of routine experience, the experience that remains valuable is precisely the judgment-linked kind.

Source register

Source	Population and n	Field date
Pew Research Center, Americans and AI 2026	U.S. adults, n = 5,119	Feb 2026
Bick, Blandin and Deming, Real-Time Population Survey / GenAI Adoption Tracker	U.S. adults 18-64, 8 waves, ~40,000	May 2026
Gallup Workforce Study, Q2 2026	Employed U.S. adults, n = 22,573	May 2026
Gallup Workforce Study, Q1 2026 and prior waves	Employed U.S. adults, n = 23,717	Feb 2026
U.S. Census Bureau, Business Trends and Outlook Survey	U.S. businesses, biweekly panel	Dec 2025 to May 2026
U.S. Census Bureau, CES Working Paper 26-25 (AI supplement)	U.S. firms, nationally representative	Nov 2025 to Jan 2026
Federal Reserve Bank of New York, Survey of Consumer Expectations	U.S. population, representative	Nov 2025
Stanford HAI, AI Index Report 2026	Global, multi-source compilation	2026
Anthropic Economic Index (Feb 2025, Sep 2025, Jan 2026, Mar 2026 reports)	Claude.ai and API conversations, 150+ countries	through Feb 2026
Chatterji, Cunningham, Deming et al., How People Use ChatGPT (NBER w34255)	~1.1M classified conversations	May 2024 to Jun 2025
Anthropic, Coding agents in the social sciences	Quantitative social scientists, n = 1,260	Feb to Mar 2026
AMA, 2026 Physician Survey on Augmented Intelligence	U.S. physicians, n = 1,342	2026
Doximity, 2026 State of AI in Medicine Report	U.S. physicians, n = 3,151, 15 specialties	Nov 2025 to Jan 2026
Doximity, 2026 Physician Compensation Report	U.S. physicians	2026
Thomson Reuters Institute, 2026 AI in Professional Services Report	Legal, tax, risk, government professionals	2026
American Bar Association, Legal Technology Survey Report	U.S. lawyers	2025
Gallup and Walton Family Foundation, teacher survey	U.S. K-12 teachers, n = 2,232	2024-25 school year
RAND Corporation, teacher and student AI surveys	U.S. K-12 teachers and students	2024-25
Nature Research Intelligence and Fudan University, AI for Science 2026	Researchers worldwide, n = 10,480	Mar 2026
Deloitte Access Economics for RMIT Online, Beyond Prompting	Australian workforce	Mar 2026
ManpowerGroup, 2026 Global Talent Barometer	Workers, 19 countries, ~14,000	Jan 2026
AARP Research, Tech Trends and Older Worker	U.S. adults and workers 50-plus	2025 to 2026

Source	Population and n	Field date
surveys		
Randstad, generational AI adoption research	Global workforce	2025
Generation, Age-Proofing AI	U.S. workers 45+ (n = 2,610) and employers (n = 1,488)	2024
Brynjolfsson, Li and Raymond, Generative AI at Work (QJE)	Customer support agents, field experiment	2025
MIT Project NANDA, The GenAI Divide: State of AI in Business	52 interviews, 153 leaders, ~350 employees, 300 deployments	2025 (preliminary)
Wharton, generative AI ROI study	U.S. businesses measuring ROI	Jun to Jul 2025
NBER Working Paper 34836, executive survey	Executives, US/UK/Germany/Australia, n = ~6,000	2026
McKinsey, State of AI	Global executives	2025 to 2026
PwC, 2026 CEO survey	Global CEOs	2026
IDC/Microsoft, IBM, Accenture ROI studies	Enterprises; vendor-sponsored	2025 to 2026
Vendor pricing pages and licensing-consultancy analyses	Published list prices and reported negotiated rates	2026

Verification note. Every figure in this review carries a source grade recorded in the accompanying verification audit. Grade A figures were read from the primary publication: Pew Research Center, Eurostat, the Federal Reserve Bank of New York, Nature Research Intelligence and Fudan University, the Walton Family Foundation, NBER Working Paper 34836, and the Quarterly Journal of Economics. Grade B figures come from official domains read at excerpt rather than in full, including Gallup, the U.S. Census Bureau, Stanford HAI, the American Medical Association, Doximity, Thomson Reuters and the Anthropic Economic Index. Grade C figures, traceable only to secondary or commercially interested reporting, are flagged in place throughout the text and should not be cited without independent confirmation. Four items remain unresolved and are identified as such where they appear: the American Bar Association technology survey figures, which sit behind a paywall; the RAND student-instruction figure; the Generation employer-consideration figures on candidates over 60; and all enterprise seat pricing, none of which is published by the vendors.