



INTERNATIONAL RESEARCH INSTITUTE — MARKET STUDY

Education Technology and the Future of Workforce Skills

Education technology · 2025-10-14 · Licence CC BY 4.0

The state of the market

Education technology has moved from a peripheral procurement line to a structural component of how economies build and renew human capital. The proximate cause is not the appeal of digital classrooms but the changing shape of work. As task composition shifts across a widening set of occupations, the interval over which a given bundle of technical skills remains current has shortened, and the burden of keeping workforces productive has migrated from a one-time investment in initial education toward continuous, in-career learning. Education technology is the delivery mechanism most capable of meeting that demand at scale, which is why its centre of commercial gravity is moving from schools toward employers and working adults.

The market itself is large but poorly measured. Published estimates span a wide band, and much of the divergence reflects definitional choices rather than genuine disagreement about activity — whether one counts hardware, connectivity, publisher revenue, corporate learning spend and consumer subscriptions within a single figure. On our reading of the available expenditure data, annual revenue plausibly sits in the range of USD 250–340 billion for 2024, with corporate and workforce-oriented learning the most defensible and highest-margin pool. We present all such figures as transparent estimates rather than measured totals.

By the numbers

| INDICATOR | VALUE |
|--|---------------|
| Est. global edtech revenue, 2024 — Indicative range; estimate | USD 250–340bn |
| Corporate & workforce learning — Largest, highest-margin pool | USD 95–130bn |
| Distinct market segments — Different buyers and economics | 6 |
| Assessment & credentialing — Smallest today, the strategic prize | USD 15–25bn |

Two forces now dominate the outlook. The first is the diffusion of generative artificial intelligence, at once the strongest demand tailwind the sector has seen and its most acute competitive threat: cheap, fluent content generation erodes the defensibility of businesses built on delivering explanatory material or answering student questions, while raising the value of the functions AI cannot easily counterfeit — reliable assessment, credible credentials and verified evidence that a person can do the work. The second is a persistent efficacy gap. Adoption has repeatedly run ahead of rigorous proof that specific tools improve learning or employment outcomes, and closing that gap is the binding constraint on public procurement, durable pricing and the sector's social legitimacy.

For decision-makers the strategic questions are therefore less about whether to adopt digital learning and more about how to govern it. Governments face choices over credential recognition, data protection, procurement standards and the equity consequences of uneven access. Employers must decide how much skill formation to internalise versus buy. Investors confront a sector where the addressable need is genuine and durable but where business-model quality varies enormously and where several prominent failures have already tempered earlier enthusiasm. This report maps the structure, the drivers, the regional divergence and the plausible paths to 2030.


Headline judgements

- **The demand is structural, not cyclical.** Continuous reskilling is driven by lasting changes in task composition across occupations, not a temporary technology cycle — which favours workforce and adult-learning segments over traditional classroom software.
- **Headline market figures conceal more than they reveal.** Commonly cited totals bundle segments with different unit economics, buyers and defensibility; a credible view requires treating consumer subscriptions, K–12 software, higher-education provision and corporate learning as distinct markets.
- **Corporate and workforce learning is the commercial core.** Employers are the buyers with both budget and a direct financial return from skill formation, carrying higher willingness to pay and clearer outcome accountability than publicly funded classroom technology.

- **Generative AI reprices the value chain.** Content delivery and homework-assistance models lose defensibility; assessment, credentialing, verification and workflow-embedded learning gain it. Several first-wave listed edtech businesses have already seen valuations fall sharply.
- **The efficacy evidence base remains thin.** Independent evidence that specific products produce learning or wage gains is scarce relative to the scale of deployment — the sector's central credibility problem and most important research priority.
- **Outcomes are diverging by geography and income.** Access to devices, connectivity and effective instruction is highly uneven; absent deliberate policy, the technology amplifies existing human-capital advantages rather than compensating for them.
- **The credentialing layer is the strategic prize.** Whoever provides the trusted, portable record of what a person can do captures a durable position between learners, educators and employers — more defensible than content provision.

1. Context and why it matters

For most of the twentieth century, the dominant model of human-capital formation was front-loaded: a concentrated period of initial education, followed by a working life in which accumulated skills depreciated slowly enough that periodic on-the-job updating sufficed. That model rested on an assumption of relative stability in the task content of jobs. Over the past two decades that assumption has weakened. Automation of routine cognitive and manual tasks, the spread of software into nearly every occupation, and now the arrival of capable generative models have accelerated the churn in what workers are actually asked to do.



The model of human-capital formation is shifting from a front-loaded period of initial education toward continuous, lifelong learning delivered to working adults. — IRI

The consequences show up in a consistent, if imprecisely measured, pattern across labour-market research. International bodies that survey employers report that a substantial share of the workforce expects material change to the skills their roles require within a five-year horizon, and that a majority of workers will need some retraining over comparable periods. Adult-skills assessments across advanced economies reveal large populations of working-age adults with literacy, numeracy and digital-problem-solving proficiency below the level modern work increasingly demands. These findings are not precise forecasts, and methodologies differ, but the

direction is robust and widely corroborated: the stock of skills in existing workforces is a poor match for the flow of skills employers say they need, and the mismatch is growing rather than resolving.

This is why education technology matters as an economic question rather than a consumer-software one. If skill formation must become continuous and lifelong, traditional delivery institutions — schools and universities operating on multi-year cycles with fixed physical capacity — cannot meet the whole of the demand. The marginal unit of learning increasingly has to be delivered to a working adult, at low cost, at a time and place that fits around employment, and in a form modular enough to target a specific capability gap. Those are precisely the properties digital delivery can provide. Whether it provides them well, cheaply and equitably is the substance of this report.

Two framing points recur throughout. "Education technology" is not one market but a label for loosely related activities spanning consumer apps, institutional procurement, publishing, corporate services and public infrastructure, so analytical care requires disaggregation. And technology is a delivery and measurement mechanism, not a substitute for the determinants of learning — motivation, instruction quality, time on task and coherent curriculum; the sector's most consequential failures have come from treating the tool as the intervention.

2. Market structure and scale

Any single number for the global education-technology market should be treated with caution, because the total depends heavily on where the boundary is drawn. Broad definitions that include devices, connectivity, publisher digital revenue, institutional software, online higher education, corporate learning and consumer subscriptions produce figures well above narrow definitions counting only pure-play digital-learning software and services. Rather than defend a single headline, we disaggregate the sector into segments with distinct buyers and economics and give a transparent estimated range for each.

The estimates below are built by triangulation: education-expenditure data provide the outer envelope of total spending (public plus private, of the order of several trillion US dollars globally per year); corporate learning benchmarks — per-employee annual training spend applied to formal-sector employment — anchor the workforce segment; venture-funding records and listed-company disclosures calibrate the pure-play digital pools. Each segment figure is an estimate, not a measured total.

| SEGMENT | EST. 2024 REVENUE (USD BN) | INDICATIVE GROWTH TO 2030 | PRIMARY BUYER | BASIS AND CONFIDENCE |
|--|----------------------------|---------------------------|--------------------------|--|
| Corporate & workforce learning (incl. L&D services, upskilling platforms) | 95–130 | Moderate–high | Employers | Per-employee training benchmarks × formal employment; corporate disclosures. Medium confidence |
| K–12 institutional software & digital content | 45–65 | Moderate | Schools / ministries | Public procurement and publisher digital revenue. Medium confidence |
| Higher-education digital (online programmes, OPM, courseware) | 40–60 | Low–moderate | Universities / students | Listed providers, tuition-share models. Medium–low confidence |
| Consumer & lifelong learning (language, skills, test prep, MOOCs) | 30–45 | Moderate | Individuals | App-store and subscription disclosures. Medium confidence |
| Assessment, credentialing & verification | 15–25 | High | Employers / institutions | Testing incumbents + emerging credential providers. Low–medium confidence |
| Enabling infrastructure attributable to learning (LMS, analytics, tutoring AI) | 20–30 | High | Institutions / employers | Software vendor revenue apportioned. Low confidence |
| Indicative total | 250–340 | Moderate overall | Mixed | Sum of ranges; not additive-precise. Estimate |

Estimated 2024 revenue by segment — stated ranges

| CATEGORY | LOW ESTIMATE (USD BN) | HIGH ESTIMATE (USD BN) |
|----------------------------|-----------------------|------------------------|
| Corporate & workforce | 95 | 130 |
| K–12 institutional | 45 | 65 |
| Higher-ed digital | 40 | 60 |
| Consumer & lifelong | 30 | 45 |
| Assessment & credentialing | 15 | 25 |
| Enabling infrastructure | 20 | 30 |

Endpoints of the ranges tabulated above. Estimates; segment ranges are not precisely additive.

Several structural features follow. The corporate and workforce pool is the most commercially attractive: the buyer has budget, a direct financial return from a more productive workforce, clearer accountability for outcomes than a publicly funded school system, and higher willingness to pay per learner. The K–12 institutional segment is large but constrained by procurement cycles, tight budgets and cautious efficacy requirements; margins and pricing power are correspondingly weaker. Higher-education digital provision has been the site of the sector's most visible business-model stress, as tuition-share arrangements between universities and online-programme managers came under regulatory and reputational pressure. Consumer learning has real subscription revenue but high churn and heavy dependence on a few scaled brands.

Indicative composition of the sector, 2024

| SEGMENT | SHARE |
|----------------------------|-------|
| Corporate & workforce | 38% |
| K–12 institutional | 18% |
| Higher-ed digital | 17% |
| Consumer & lifelong | 13% |
| Enabling infrastructure | 8% |
| Assessment & credentialing | 7% |

Shares derived from the midpoints of the segment ranges above; illustrative estimates that are not precisely additive. Corporate and workforce learning is the largest and most defensible pool.

The segment we flag as strategically most important relative to its current size is assessment, credentialing and verification. It answers the question employers actually care about — can this person do the work? — and it is the layer least exposed to commoditisation by generative AI, a point developed in Sections 4 and 6.

The sector's funding history is instructive. Venture investment rose steeply during pandemic-driven remote learning around 2020–2021, reaching a level widely estimated in the tens of billions of dollars globally at its peak, before contracting sharply in 2022–2023 as demand normalised and capital costs rose. The correction was informative: it separated businesses with durable demand from those whose growth had been a temporary artefact of school and campus closures. The most prominent cautionary episode was the collapse in valuation of a major emerging-market tutoring company that had raised enormous sums — a reminder that scale of funding is not evidence of underlying unit economics.

The venture-funding cycle, 2020–2023

| WHEN | MILESTONE | DETAIL |
|--------------------|-------------------------|--|
| 2020–2021 | Pandemic funding surge | Venture investment rose steeply during remote learning, reaching a level widely estimated in the tens of billions of dollars globally at its peak. |
| 2022–2023 | Sharp correction | Funding contracted as demand normalised and capital costs rose, separating durable demand from growth that was a temporary artefact of school and campus closures. |
| Cautionary episode | A high-profile collapse | A major emerging-market tutoring company that had raised enormous sums saw its valuation collapse — scale of funding is not evidence of underlying unit economics. |

3. Demand-side drivers: the shortening useful life of skills

The demand case rests on a single, durable proposition: the useful life of specific technical skills is falling, while the working life over which those skills must be maintained is lengthening as populations age. The gap between those two trends is the structural market.

The shortening of skill life has several sources. The most direct is software and automation, which continuously reshape the task content of occupations. The fastest-moving is generative AI, which is altering the value of particular skills within white-collar work — reducing the premium on some forms of routine drafting, coding and analysis while raising it on judgement, oversight and the ability to direct and verify machine output. Whichever technology is proximate, the effect on the learning market is the same: workers must acquire new capabilities repeatedly across a career, in increments small enough to fit around employment.

The lengthening of working life compounds this. As demographic ageing raises the average age of workforces in advanced and many middle-income economies, a growing share of the labour force is decades removed from initial education and therefore dependent on in-career learning to remain productive. For such workers the alternative to continuous learning is premature exit from the workforce, which economies with shrinking working-age populations can ill afford.

There is an important asymmetry in who currently meets this demand. Formal-sector employers in higher-productivity industries invest substantially in structured learning, and their workers benefit disproportionately, while workers in small firms, informal employment and lower-productivity sectors receive far less despite often facing greater displacement risk. Left to itself, the market directs continuous learning toward those who already have the most human capital, not those who need it most — the equity problem developed in Section 5.

A final demand-side point concerns completion. The economics of digital learning depend not on enrolment but on completion and application. Open online courses demonstrated early that very large enrolment can coexist with very low completion, and that those who complete are disproportionately the already well-educated. The products most likely to succeed engineer completion and on-the-job application — by embedding learning in work and tying it to a recognised credential — rather than merely maximising sign-ups.

4. Supply-side dynamics: how generative AI reprices the value chain

Generative AI is the defining supply-side variable for the sector this decade. Its effect is best understood not as uniform uplift but as a repricing of different functions within the learning value chain: some become cheaper and less defensible, others more valuable.

The functions most exposed to commoditisation are content generation and explanation. When a general-purpose model can explain almost any concept, generate practice problems and answer a learner's questions at near-zero marginal cost, businesses whose value rested on assembling and delivering explanatory content lose pricing power. The early warning is concrete: at least one prominent listed company built on paid homework assistance saw its business deteriorate rapidly once free, capable models became widely available. The lesson generalises to any model whose moat was proprietary content or human-answered questions.

The functions that gain value are those AI cannot easily counterfeit. Reliable assessment — establishing what a learner actually knows and can do under conditions that resist gaming — becomes more valuable precisely because AI makes it easier to produce work one did not author. Credentialing and verification gain for the same reason: employers need trustworthy signals in a world where output is cheap to fabricate. And learning embedded directly in the workflow — capability delivered at the moment of need inside the software a person already uses — becomes both more feasible and more defensible, anchored to a specific work context rather than generic content.

The net effect is a shift in where defensibility lives. In the previous generation, scale in content and distribution conferred advantage. In the emerging configuration, advantage accrues to whoever controls trusted assessment and credentialing, owns proprietary data on real learning and employment outcomes, and is embedded in the systems where work actually happens — implications developed in Section 6. A caveat belongs here: generative AI also introduces new failure modes, including the risk that learners offload the very cognitive work that produces skill, so the tools that improve outcomes will be those designed around how people actually learn rather than those that merely make answers frictionless.

5. Regional and comparative lens

The sector's trajectory looks very different across regions, and the divergence is central to the story. Three broad patterns are worth distinguishing.

In high-income economies with ageing workforces, the dominant question is renewal: how to keep large, experienced but slowly depreciating labour forces productive as task content shifts. Here the corporate learning market is most developed and digital infrastructure near-universal; the binding constraints are efficacy evidence, worker motivation and the integration of learning into work. The commercial opportunity is real but buyers are sophisticated and increasingly demand proof of return.

In large middle-income economies with young, growing and digitally connected populations, enormous latent demand meets immature market structures. Countries with hundreds of millions of young people and expanding smartphone access represent the largest potential learner pools in the world. Yet this is also where the sector's most severe governance and business-model failures have occurred, where aggressive growth outran unit economics and consumer protection. The opportunity is genuine but the execution risk — and the case for public oversight of consumer marketing and financing — is high.

In lower-income economies, particularly across parts of Sub-Saharan Africa and South Asia with fast-growing working-age populations, the constraint is more fundamental: reliable electricity, affordable connectivity, device access and, above all, effective foundational instruction. Where children and adults lack secure literacy and numeracy, sophisticated digital tools cannot substitute for the missing foundation, and the most cost-effective interventions are often simple, low-bandwidth and teacher-supporting rather than teacher-replacing. This is where the mismatch between where investment flows and where human-capital returns are highest is starkest.

Three regional patterns compared

| REGION TYPE | DEMOGRAPHIC PROFILE | BINDING CONSTRAINT | COMMERCIAL POSTURE |
|--|---|---|--|
| High-income, ageing | Large, experienced, slowly depreciating workforce | Efficacy evidence, worker motivation, integration into work | Most developed corporate market; sophisticated buyers demand proof of return |
| Large middle-income | Young, growing, digitally connected — the largest potential learner pools | Governance and business-model failures; consumer protection | Genuine opportunity but high execution risk |
| Lower-income (parts of Sub-Saharan Africa, South Asia) | Fast-growing working-age populations | Electricity, connectivity, devices and foundational instruction | Low-bandwidth, teacher-supporting tools; investment–need mismatch starkest |

Synthesised from Section 5; qualitative comparison.

The comparative lens yields a clear cross-cutting conclusion: education technology is not inherently equalising. Its benefits accrue most easily to those with the connectivity, devices, prior education and institutional support to use it well. Absent deliberate policy — subsidised access, foundational-skills investment, public-interest content and standards that protect learners — the technology tends to widen human-capital gaps between and within countries rather than narrow them. Whether the coming decade produces convergence or divergence is a policy choice, not a technological inevitability.

6. Competitive landscape

The competitive structure of the sector reflects its segmentation. There is no single dominant firm across education technology, because the segments are genuinely different markets with distinct competitive dynamics.

In consumer learning, a small number of scaled brands with strong distribution hold durable positions, particularly in language learning, while a long tail of applications competes on acquisition cost and struggles with retention. In higher-education digital provision, the model of universities partnering with online-programme managers on tuition-share terms has come under sustained pressure, and several high-profile players have retrenched. In corporate learning, the market remains fragmented across content libraries, platform providers and services firms, with consolidation underway as buyers seek fewer, more integrated relationships. In

assessment and credentialing, established testing and certification incumbents hold strong positions that newer entrants contest with digital, skills-based alternatives.

The strategic contest that will define the next several years is over the credentialing and verification layer. Historically the trusted record of what a person can do has been held by universities (degrees) and professional bodies (certifications); digital credentials, skills-based hiring and employer-recognised micro-credentials are attempts to build a more granular, portable alternative. Whoever establishes the trusted standard captures a position far more defensible than content provision, because it benefits from network effects — a credential is valuable in proportion to how many employers recognise it — and from switching costs once embedded in hiring systems. This is why we regard the credentialing layer as the sector's strategic prize despite its modest current revenue.

Two structural risks temper the competitive outlook. First, the largest general-purpose technology platforms can enter learning adjacencies cheaply, bundling AI-driven tutoring into products people already use, which caps the pricing power of standalone tools. Second, public institutions remain the dominant providers of education globally, and in many segments the relevant competition is not another vendor but the option of in-house or public provision. The addressable private market is therefore smaller than the total education-spending envelope suggests, and firms that mistake the latter for the former routinely overestimate their opportunity.

7. Risks and open questions

Several risks bear materially on the outlook.

The efficacy gap. The most important uncertainty is whether widely deployed tools actually improve learning and employment outcomes. Independent evaluation remains scarce relative to the scale of adoption, and much of the evidence that exists is vendor-sponsored or measures engagement rather than durable capability. Until the evidence base thickens, public procurement will stay cautious, pricing power will be limited, and the sector will remain vulnerable to the charge that adoption is driven by novelty rather than proof. Closing this gap is the single most valuable investment the field could make.

***The binding constraint** — Independent evidence that specific products produce learning or wage gains remains scarce relative to the scale of deployment. It is the sector's central credibility problem — and until it thickens, public procurement stays cautious and durable pricing power stays out of reach.*

Data protection and learner analytics. Learning systems generate detailed data about individuals, including minors. The collection, use and security of that data raise significant privacy questions, and regulatory regimes are tightening. Products built on extensive data capture face both compliance cost and reputational exposure; mishandling of children's data in particular is a source of acute legal and political risk.

Consumer protection. Aggressive marketing, opaque course financing and inflated claims about employment outcomes have already produced regulatory and reputational problems, most visibly in the collapse of a heavily funded consumer-tutoring business. As skills-based training becomes financially consequential for individuals, oversight of claims and financing is likely to increase.

Over-reliance and skill hollowing. If learners and workers offload core cognitive tasks to always-available AI, measured productivity may rise while underlying capability falls, storing up problems that surface only when the tools are unavailable or wrong. The effect is real and not yet well understood.

Macroeconomic and funding sensitivity. Corporate learning budgets and venture funding are both cyclical. A downturn compresses the discretionary training spend that anchors the most attractive segment, and the sector has already shown how quickly capital can retreat. Models dependent on continuous external funding rather than customer revenue are especially exposed.

Where this goes: three scenarios to 2030

We do not offer a point forecast. Instead we describe three internally consistent scenarios, distinguished chiefly by how two uncertainties resolve: whether the efficacy evidence base strengthens, and whether governance and access policy is deliberate or laissez-faire.

Three scenarios to 2030

A — Evidence-led integration — Constructive

Independent evaluation matures, buyers reward demonstrated outcomes, and AI is deployed around sound learning design. Credentialing standards consolidate and become portable while policy subsidises access.

| METRIC | VALUE |
|-----------------|---------------------------|
| Revenue path | Steady growth |
| Concentrated in | Workforce & credentialing |
| Equity | Gaps narrow |

B — Fragmented muddling-through — Base case

Adoption continues and AI reprices the value chain, but the efficacy gap closes only partially and governance stays patchy. Benefits accrue unevenly, punctuated by periodic corrections as over-funded models fail.

| METRIC | VALUE |
|--------------|------------------------------|
| Revenue path | Below optimistic projections |
| Benefits | Uneven |
| Probability | Most probable |

C — Hype-and-retrench — Adverse

Weak evidence, efficacy and consumer-protection failures and a funding downturn sour sentiment. AI commoditises content faster than incumbents adapt, credentialing fails to standardise, and access gaps widen.

| METRIC | VALUE |
|------------------|------------------------|
| Revenue path | Stagnates or contracts |
| Exposed segments | Content delivery |
| Equity | Gaps widen |

Scenario A — Evidence-led integration (constructive). Independent evaluation matures, buyers reward demonstrated outcomes, and generative AI is deployed around sound learning design rather than frictionless answers. Credentialing standards consolidate and become portable across employers, while public policy subsidises access and invests in foundational skills, narrowing rather than widening gaps. The sector's revenue grows steadily, concentrated in workforce learning and credentialing, and its social contribution is broadly positive. This is achievable but requires deliberate choices, particularly on evaluation and access.

Scenario B — Fragmented muddling-through (base case). Adoption continues, AI reprices the value chain roughly as described, and some segments consolidate, but the efficacy gap closes only partially and governance remains patchy. Benefits accrue unevenly, concentrated among already-advantaged workers and regions, and the market grows but below the more optimistic projections, punctuated by periodic corrections as over-funded models fail. This is, in our judgement, the most probable path absent decisive policy intervention.

Scenario C — Hype-and-retrench with widening gaps (adverse). Weak evidence, high-profile efficacy and consumer-protection failures, and a funding downturn combine to sour buyer sentiment. AI commoditises large parts of the content-delivery business faster than incumbents adapt, credentialing fails to standardise, and access gaps widen as public investment lags. Revenue stagnates or contracts in real terms in the exposed segments, and the sector's net contribution to human capital disappoints.

Across all three, two directional judgements hold: workforce and credentialing layers outperform classroom-software and content-delivery layers, and the distribution of benefits is determined more by policy than by technology.

Practical implications

For governments and policymakers. Treat efficacy evidence as a public good and fund it: independent evaluation of widely used tools is undersupplied by the market and is the precondition for sound procurement. Establish clear standards for the recognition and portability of digital credentials, so skills become legible across employers and borders. Protect learners through data-protection enforcement and oversight of outcome claims and course financing. Above all, treat access and foundational skills as the equity levers they are: subsidised connectivity and devices, and sustained investment in basic literacy and numeracy, determine whether technology narrows or widens human-capital gaps.

For business and investors. Distinguish sharply between segments; do not price a workforce-learning or credentialing business as if it shared the economics of consumer content. Favour models whose defensibility survives cheap generative content — assessment, verification, proprietary outcome data and workflow-embedded learning — over those whose moat was content or answers. Demand evidence of learning and employment outcomes, not engagement metrics, and discount claims resting on funding scale rather than unit economics. For employers as buyers, the central choice is how much skill formation to internalise versus purchase, and how to make internal skills portable enough to retain and redeploy talent.

For education providers and universities. The credentialing franchise institutions have long held is contestable and should be actively defended and modernised through granular, stackable and verifiable credentials that reflect current capability. Partnerships with technology providers should be structured around demonstrated outcomes and transparent economics, learning from the stress in tuition-share arrangements. The durable institutional advantage is trust — in assessment and in the credential — and should be protected accordingly.

A note on method and data

This report is an analytical synthesis rather than a primary statistical study, and we are explicit about what is established and what is estimated. Our qualitative claims about labour-market restructuring, the shortening useful life of skills, adult-skills deficits and demographic divergence rest on the broad and consistent body of evidence produced by international economic and labour organisations, adult-skills assessment programmes and national statistical agencies; these directional findings are well corroborated even where specific percentages vary by methodology.

Our market-sizing figures are estimates produced by triangulation, not measured totals. We bound the total by reference to global education-expenditure data (public and private), then build segment ranges from three independent anchors: corporate learning benchmarks (per-employee annual training spend applied to formal-sector employment), venture-funding records, and the disclosures of listed education-technology companies. Where sources conflict — as they frequently do, largely because of differing definitions of "edtech" — we report a range and state the basis and our confidence. The segment ranges in Section 2 are not precisely additive, and the indicative total should be read as an order-of-magnitude estimate. All forward-looking figures are conditional on stated assumptions and should not be read as forecasts.

We have deliberately avoided attributing invented precise statistics to real institutions. Where we cite a pattern — the surge and contraction in venture funding around 2021–2023, or the deterioration of a homework-assistance business following the diffusion of generative AI — we describe the well-documented direction without fabricating exact numbers; readers seeking precise current figures should consult the primary sources directly. The principal limitations of this analysis are the thinness of independent efficacy evidence, the definitional instability of the market, and the speed at which generative AI is changing the competitive landscape, any of which could move the estimates materially.

Sources and evidence

The labour-market and adult-skills evidence rests on the recurring international programmes below, read directly; venture-funding and listed-company figures are used to calibrate the segment ranges and are referenced by category where proprietary.

- World Economic Forum (2023). *The Future of Jobs Report 2023*. WEF, Geneva.
- OECD (2024). *OECD Skills Outlook and Survey of Adult Skills (PIAAC): results and technical report*. OECD Publishing, Paris.
- OECD (2023). *Education at a Glance 2023: OECD Indicators*. OECD Publishing, Paris.

- UNESCO (2023). *Global Education Monitoring Report 2023: Technology in Education — a Tool on Whose Terms?* UNESCO, Paris.
- World Bank (2023). *Education technology and learning outcomes: an evidence review*. World Bank, Washington, DC.
- International Labour Organization (2023). *World Employment and Social Outlook: Trends 2023*. ILO, Geneva.
- International Telecommunication Union (2024). *Measuring Digital Development: Facts and Figures 2024*. ITU, Geneva.
- HolonIQ and comparable venture databases (2022–2024). Education-technology funding and market-sizing trackers, referenced by category for direction and magnitude.
- Listed education-technology company disclosures (2021–2025), used to calibrate segment revenue pools.

Suggested citation

Ellison, N., Devlin, M., and Raghunathan, P. (2025). *Education Technology and the Future of Workforce Skills*. Flagship Report. International Research Institute, Education & Human Capital Programme. Licensed under CC BY 4.0.