GSV Ventures

ENVISIONING PRE-K TO GRAY:

A Forecast for Learning & Earning in 2025-2026

AUTHORS

Jennifer Lee & Claire Zau

Foreword

As the sun sets on summer vacations and we head back to school and back to the office, GSV Ventures has an eye on the future: **What will learning and work look like this academic year, this fiscal year, and beyond?**

It's no secret that AI has unlocked a new axis of possibility in the way we learn, work, and live—with potential for exponential acceleration—yet the world is largely still grappling with what that promise looks like in practice. While the universe of possibility remains vast, there is a lens of likelihood we see bringing the future into focus.

We are entering an era of AI ubiquity: no longer an isolated tool, but an omnipresent force multiplier and the connective tissue embedded throughout our daily tasks and touchpoints. The winners in this new era will balance innovation with responsibility, automated convenience with cognitive function, and the delicate relationship between human and machine.

In this look ahead at the future of learning and work, we aim to define:

- Major thematic observations
- What makes this moment an inflection point
- Key predictions (and their precedent)
- Short- and long-term projected impacts

In addition to this report, we'll deep-dive into each theme in weekly live sessions throughout the fall, joined by relevant experts to address your most pressing questions. Sign up **here** to stay updated about session registration and recordings.

REPORT AUTHORED BY

Jennifer Lee
Partner,
GSV Ventures

Claire Zau
AI Lead & Partner,
GSV Ventures

Table of Contents

THEME 1:

Fostering Trust in the Age of AI

K-12 | HIGHER ED | CORPORATE & LIFELONG

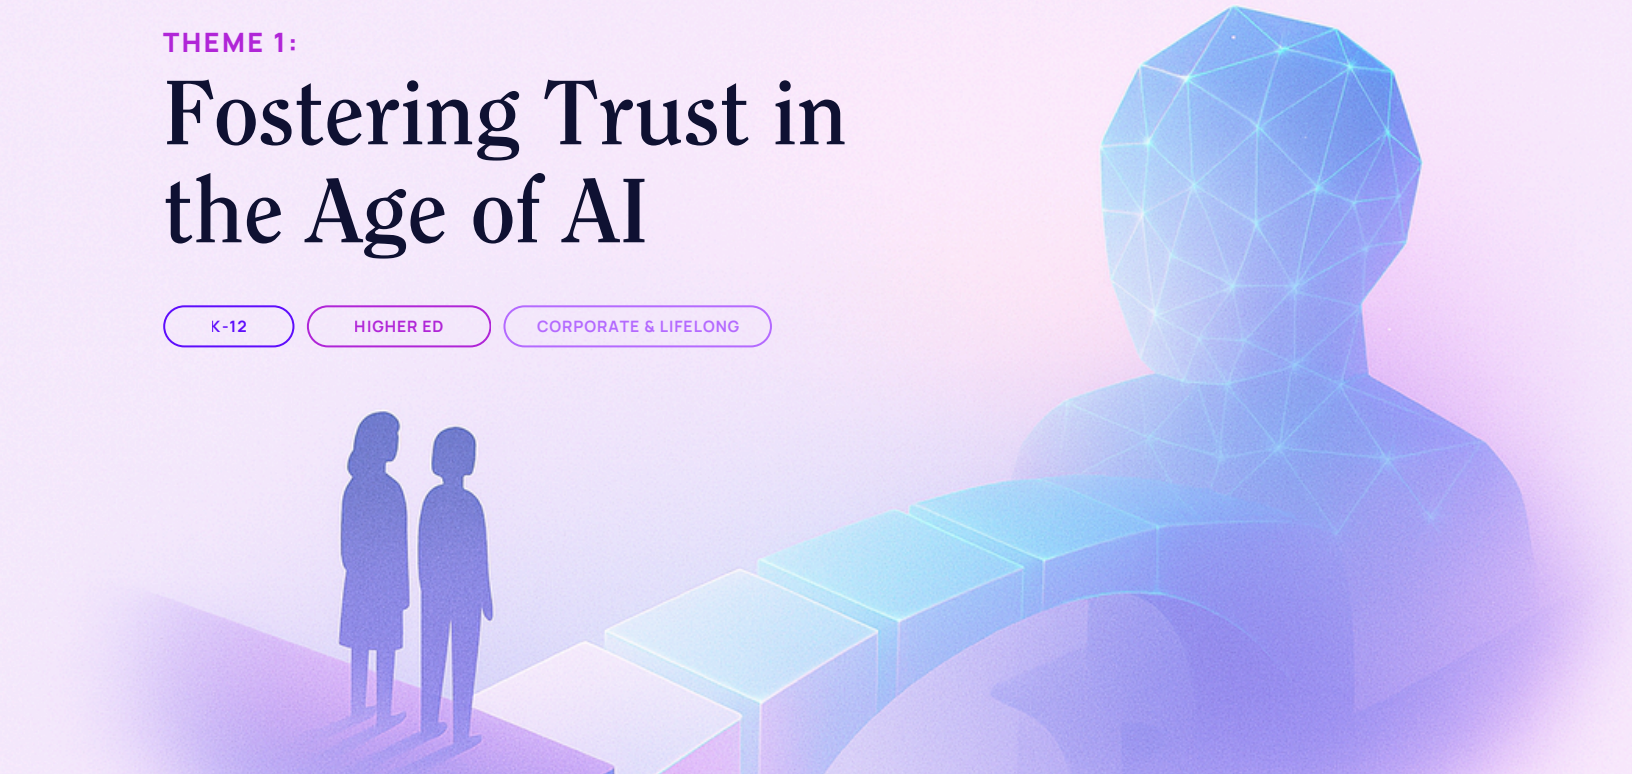

Thesis

Building trust will be the cornerstone of success for any AI company or product. Whether leveraging AI agents for employee upskilling or utilizing LLMs for tailored lesson plans, the widespread adoption of AI will hinge on the trust that learners and stakeholders place in these systems.

Historical Context

New technologies historically face initial distrust as people adapt to unfamiliar innovations. We've seen this pattern with electricity (safety and reliability concerns), elevators (fear of falling), and cloud computing (security and data privacy anxieties). Trust curves typically progress through phases—and AI is no exception.

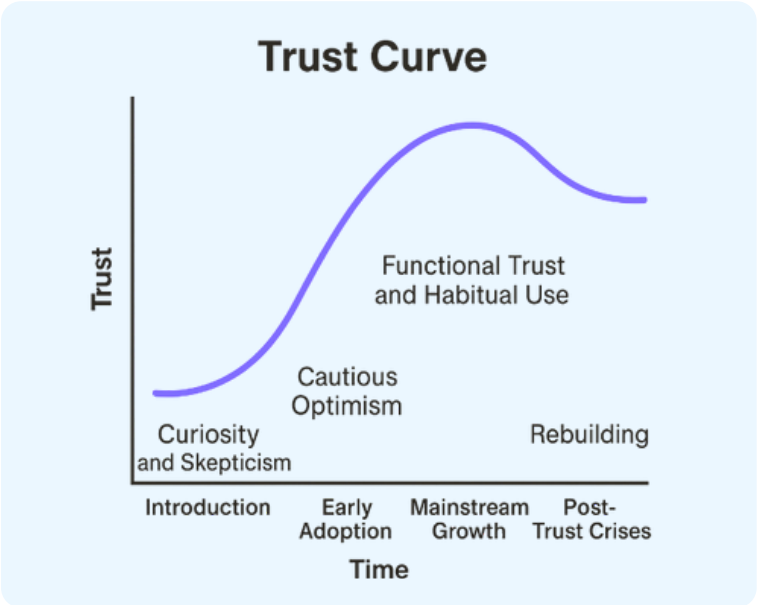

Why Now?

Early in the Adoption Curve, Particularly for AI Agents

For AI, we are arguably between the Introduction and Early Adoption phases, where trust is still a critical barrier to widespread adoption. Despite AI's improved performance in many exams and benchmarks, a narrative of distrust persists and in some cases is growing, often linked to frequent hallucinations, biases in training data, and the perceived ability to manipulate models.

Recent research highlights this: the decline in trust when financial recommendations incorporate AI and a survey of IT leaders by Capgemini show a decline in trust in fully autonomous AI agents from 43% to 27%. Recent research by Google DeepMind also highlighted LLMs' overconfidence in initial choices and underconfidence when challenged. Thus, there is significant opportunity for companies to address trust directly, and we believe this will be a critical factor for future success.

Explosion of Content Across Modalities

With each model update, GenAI's capability to produce increasingly realistic content across all output modalities (text, image, voice, video) faster and cheaper is well discussed. This explosion of AI content has created concerns about people's ability to determine what is "real"; our favorite benign examples this summer were the backyard videos supposedly catching animals jumping on trampolines. However, just as important, growing infobesity has also raised concerns of authenticity, reliability and factuality—as well as both the discoverability and identification of high-quality, relevant content.

Predicted Impacts

Short Term

- **The speed of business is the speed of trust.** While users will naturally progress along the trust curve over time and with AI readiness policies, companies prioritizing integrity and trust via the key factors on the following page will supercharge their growth.
- We are already seeing early examples of **companies actively looking to build trust in high-stakes industries** (e.g. law, finance, healthcare) and for more vulnerable populations (e.g. younger students). This has mostly centered around optimizing models for retrieval of more relevant or accurate information; human-in-the-loop approaches, like educator override capabilities; and use-case-specific or goal-specific evals.

FACTOR	DEFINITION	HOW WINNERS WILL FOSTER TRUST	EXAMPLES IN LEARNING
Reliability	The extent to which a system consistently performs as intended without failure or errors	Improvements in reliability have been ongoing and will continue, but with diminishing returns on building trust.	Less than 1 year ago, ChatGPT couldn't count the number of "r"s in "strawberry." Today, Google DeepMind and OpenAI both achieved gold medal-level performance at this year's IMO.
Predictability	The degree to which behavior matches users' expectations in similar situations	Create and maintain functional memory systems and dynamic contextual relationships.	In the workforce, AI agents will recommend key training based on consistent evaluation logic and with the context of current and past performance.
Benevolence	The belief that the system and its designers intend to act ethically and in the user's best interests	Differentiate by creating evals and benchmarks that extend beyond the standard accuracy and performance metrics to include community values and goals.	In K-12, winners will proactively show evidenced-based positive impact. Examples are: • Meeting ESSA tiers like CourseMojo • LearnLM's pedagogical benchmarking • Latimer AI's BiQ benchmarking
Control	The extent to which users can influence the system's actions or override its decision	Enable transparency that empowers clients to influence model outputs and adapt AI agents to evolving priorities.	Ability to apply guardrails for given environments (e.g. age-appropriate experiences) and desired data privacy and handling rules

Long Term

- As AI blurs lines around reality and humanity, **verification layers will become essential**. Technologies such as blockchain-based provenance, NFT-backed digital authenticity, and biometric identity verification (think Clear's LinkedIn verification for recruiting) could help re-establish trust in digital content, as well as both digital and physical interactions.
- We anticipate the **creation of a trust-based AI industry**—where reliability, predictability, benevolence, and control are central pillars. This will include the emergence of entirely new professional roles designed to ensure AI systems remain reliable, explainable, and aligned with human needs. Early examples of this are already starting to surface, typically as part of SWE roles. For example, Cisco posted an AI Software Engineering Tech Lead (AI Validation / AI Artifacts) role, in which one responsibility is to "Collaborate across engineering, security research, and product teams to establish robust frameworks for evaluating AI systems' trustworthiness and resilience."

Eventually, this trust-building need will extend beyond existing engineering roles for AI evals, benchmarking, and red-teaming. **Potential new professions may include:**

- **AI Curation Manager** responsible for reviewing, selecting, and validating AI-generated outputs to meet factual and ethical standards before they reach end users
- **Control & Override Designer** to build and maintain user interfaces and mechanisms for human oversight, intervention, and correction of AI actions that would also help fine tune and retrain the model
- **Predictability Architect** can ensure AI systems behave consistently across similar inputs, sessions, and contexts by auditing cross-session memory and contextual persistence
- **Academic Feed/Experience Strategist** might design, monitor, and continuously refine the data pipelines that either 1) trains educational AI systems, or 2) uses AI tools directly to guide student journeys across academic courses, relevant clubs, campus resources, work experiences, etc.

- **Trusted, human curation and validation will become the ultimate premium good.** This will lead to the return of premium versions of service roles that have been or are being displaced by technology, such as Premium Travel Agents, Personal Taste Shoppers, etc.

THEME 2:

AI is Air: Ambient AI in Every Breath, Step, and Swipe

K-12 HIGHER ED CORPORATE & LIFELONG

Thesis

We are entering the era of ambient AI, where it is embedded in every step, swipe, and click. It will serve as a continuous layer: capturing data passively, responding proactively, and eventually predicting needs before we can articulate them. This shift expands the surface area of data collection, raises profound questions about agency and privacy, and pushes us to rethink what it means to teach, learn, and live in a world where AI is always on. This layer doesn't just reshape how data is captured; it also rewrites what can be measured—something we'll explore in Theme 3: From Testing to Continued Intelligence.

Historical Context

For most of modern education, technology has been reactive, not ambient. Data collection happens in isolated, intentional moments: a quiz at the end of a chapter, a standardized test, a project submission at term's end. These are formal checkpoints designed to measure progress at a fixed point in time. Yet, these snapshots miss almost everything in between—moments of confusion, sparks of ideation, the steps in a problem-solving process, or the time spent wrestling with a concept.

As a result, most tools from the first wave of generative AI have inherited these same limitations. Grading assistants, quiz generators, lesson planners, and chatbot tutors mostly serve to make existing workflows faster, cheaper, and more scalable. However, they remain fundamentally reactive and discrete, designed around the same fixed-point, after-the-fact data rather than the continuous reality of how learning actually happens.

Why Now?

Explosion of Input Modalities

AI can now potentially capture the entire spectrum of human activity—audio, video, biometric, behavioral, and spatial data—weaving them into continuous context. Microphones, cameras, and wearables are on the rise. Meeting tools like Otter, Fireflies, and Granola already transcribe, tag, and summarize conversations automatically. Even OpenAI is reportedly developing a multimodal IO device designed to see, hear, and process context continuously, potentially integrating camera, microphone, and wearable sensors into a single AI companion.

As these modalities converge, they create a "ubiquitous data profile": a continuously updating composite of a learner's context and activity. That's both an unprecedented opportunity for personalization and a profound challenge for privacy, consent, and governance.

Agentic Capabilities

The ambient data layer allows AI agents to become even more active participants in the workflow. In an education setting, a student's verbal reflection might trigger an automatically generated study plan or funnel them into a quiz experience; a teacher's voice memo could notify relevant support staff; a learner's pacing on a coding exercise might prompt a real-time hint.

Persistent Memory

Memory is currently one of the biggest focuses for AI labs. It's drawing significant research investment, and many argue it's the new moat—a core product primitive that will define the next generation of AI system. With memory-aware systems, AI can build longitudinal context, recall past interactions without the user re-explaining, and adapt responses based on long-term patterns. The fact that ChatGPT, Gemini, and Anthropic are all building persistent memory into their products signals that it's becoming a baseline expectation. That memory layer is powerful because it will stitch together currently fragmented learning experiences into a coherent, continuous profile.

Predicted Impacts

Short Term

- **"Shadow profiles" will become the norm**: Systems will begin passively building learning profiles from multimodal data (clicks, pacing, tone, writing process), often without the user explicitly "submitting" anything. While this can unlock hyper-personalized support, it risks creating opaque "shadow records" of student behavior.

  - Without strong transparency and user controls, learners and educators might not know exactly what is tracked, how it's interpreted, or how it might follow them over time. The tradeoff is delicate: people will happily share data when the benefit is clear. We've seen this play out before with fitness wearables that collect health data or social apps that track location—people opt in when the value is tangible, but push back hard when the exchange feels one-sided. The moment it feels extractive or intrusive, trust collapses.

- **Teacher–student copilot integration**: Today, AI tools for teachers and for students tend to live in separate silos, each with their own data and workflows. Ambient AI will further enable bi-directional intelligence: the student's learning journey automatically informs the teacher's dashboard, while the teacher's plans and feedback flow back into the student's context in real time. For example, if a student's copilot notices repeated difficulty in a concept, it could alert the teacher's copilot, which then adjusts tomorrow's lesson plan or suggests targeted peer workgroups.
- **Data privacy issues**: The volume and sensitivity of multimodal, behavioral data will force urgent conversations about interoperability standards, FERPA compliance, vendor contracts, and informed consent. Districts will need policies for how continuous data streams are stored, who can access them, and how long they're retained. Vendors will be pressured to provide clear opt-in/out controls, explainability features, and portability of a learner's data profile.

Long Term

- **Blurring boundaries between "school" and "life":** When persistent copilots follow students across devices and contexts, learning moments become continuous and location-agnostic. A part-time job could generate performance feedback relevant to career skills. A VR game session at home could feed into teamwork assessments. The notion of "class time" vs. "personal time" will grow more fluid, raising both exciting opportunities for contextualized learning and concerns about overreach.
- **Rise of the Learning Operating System**: A few dominant platforms are likely to emerge as the backbone for ambient AI in education—controlling the agents, the persistent memory layer, and the syncing of profiles across devices. These Learning OS providers will compete on depth of personalization, privacy protections, and ecosystem reach. Whoever owns this layer will hold a powerful position.
- **New behavioral norms in an "always on" environment**: As people grow accustomed to AI systems observing, remembering, and sometimes anticipating their needs, behaviors will subtly shift. Students may self-censor or "perform" for the AI, knowing it's recording their work. Educators might feel like they have more time freed up, knowing that AI can take on some of the administrative load. Conversely, they may feel more constrained under continuous observation, aware that every interaction is being recorded and analyzed. Over time, the presence of ambient AI could shift how people express themselves and redefine what autonomy feels like in the classroom.

THEME 3:

From Testing to Continued Intelligence

K-12 | HIGHER ED | CORPORATE & LIFELONG

Thesis

Theme 2's "Ambient AI" will also turn learning and work into rich data streams. AI's ability to harvest multi-modal data (audio, video, biometrics, spatial context) and convert it into meaningful signals will revolutionize assessment. The winners won't just measure more, they'll decide better. They'll sift through the inundation of information and convert relevant ambient signals into trustworthy, equity-minded actions that save time, personalize support, and raise real-world performance without turning school or work into a surveillance state.

Historical Context

Peter Drucker's famous maxim "what gets measured gets managed" has guided many an organization. The catch: we mostly measured what was easy and not what truly mattered.

As discussed in Theme 2, we have focused on reactive, discrete proxies (e.g. isolated summative assessments like standardized tests) while failing to provide a continuous, linked narrative of student learning—overlooking vital competencies like collaboration, reasoning under stress, or problem-solving grit.

These metrics exacerbated stress, especially in under-resourced communities; failed to predict real-world performance; and raised concerns of bias. Education, hiring, and performance management systems bent to the metrics—and thus often missed the point.

Why Now?

Ambient AI

What's changed is both technological and cultural. Advances in AI will be able to decode ambient signals: labs, presentations, peer interactions, debates, and task execution can all be transcribed, tagged, and analyzed in real time. Tools like Otter, Fireflies, and Granola already tag and summarize meetings. In K-12 classrooms, Socrait, a voice-enabled AI teaching assistant, converts teacher speech into attendance logs, behavior notes, summaries, and task reminders. This turns ambient speech into not only micro-assessment data but also interventions without extra work. Outside the classroom, Trophi.ai tracks telemetry, video, and voice in sim-racing to benchmark performance and deliver real-time coaching. In the workforce, Gong and others ingest call recordings, screen-share data, and CRM activity to assess sales and service performance, identify skill gaps, and trigger targeted coaching interventions.

Predicted Impacts

Short Term

- **Wearable EdTech arrives:** Classroom headsets, lapel mics, and computer-vision lab tools will become the norm and feed daily micro-assessments. Policies for when and how these devices should be used will also develop, similar to the cell phone usage debate in schools.
- **Beyond the grade-book:** The next evolution of "LMS" will act more like a learning operating system: ingesting multi-modal streams, normalizing events, and outputting composite progress scores tied to competencies (communication, problem-solving, procedural accuracy) on a real-time basis. This will happen first in innovative schools with integrated systems—a "School OS."
    - While these new formative assessments will eventually dominate, **summative, standardized assessments** will continue to exist for comparability.
- **Stronger MTSS:** Whole-learner views improve Tier 1–3 interventions with earlier, more specific nudges.
- **Corporate specialization:** Vertical-specific, conversational, and project-based assessments (code reviews, safety simulations, SOP adherence) replace generic interviews and hiring tests; AI copilots get their own dashboard for agent performance.
- **This assessment renaissance will have growing pains**, including:
    - **Privacy & governance**: New consent flows; data minimization by default; needed definition of data retention windows.
    - **Cognitive overload**: Without curation, teachers and managers will drown in data paralysis. Winners will auto-summarize into 1-3 next actions, not 30 charts.

Long Term

- **Skills become the substrate:** Continuous, in-context formative signals build durable skill graphs for learners and workers that are portable, verifiable, and more predictive than course titles or GPAs.
- **Beyond the grade-book—everywhere:** True Learning Operating Systems become the backbone for Ambient AI and utilize all the new forms of micro- and invisible assessment. Even schools running mixed tool stacks will implement these mixed-data assessments and robust learning profiles as interoperability improves. Eventually, competency profiles will travel across classes, after-school programs, and CTE, as mentioned earlier.
- **Work as measurable evaluation:** Everyday artifacts (e.g. tickets closed, customer escalations defused, machine checks completed) feed performance reviews and personalized growth plans. "Just-in-time" micro-training is identified and assigned agentically, while humans audit and coach.
- **Licensure reimagined:** In regulated fields (CNA, HVAC, advanced manufacturing), AI-native assessment platforms could unify:
  1) Practice and prep,
  2) Adaptive, simulation-based testing, and
  3) Direct pipelines to state licensure systems.

  Real-time evidence from training can satisfy parts of certification, compressing time-to-license in fields with labor shortages.
- **Persistence of summative assessments:** National and international yardsticks like NAEP and PISA endure for comparability, but will evolve to include adaptive tasks and simulations.

THEME 4:

Fallout from Automation of Entry-Level Jobs

CORPORATE & LIFELONG

Thesis

The rapid acceleration of AI agent capabilities threatens to automate away foundational entry-level jobs—roles that have long served as critical rungs on the white-collar career ladder. As AI absorbs routine work, the result could be disruptions to talent pipelines, workforce training, and the role of human judgment. Education systems, employers, and policymakers will need to rethink how people acquire experience, how careers start, and what skills remain uniquely human.

Historical Context

For decades, entry-level white-collar jobs—analyst roles in finance, junior reporters in media, research assistants in marketing—served as training grounds for the next generation of professionals. These early roles allowed workers to "learn by doing," honing skills through repetitive tasks, close supervision, and gradual exposure to higher-responsibility work. The lower rungs of this ladder have always been more task-heavy, repetitive, and, by extension, more automatable than the higher rungs. Today, AI can increasingly perform these tasks faster, cheaper, and often with fewer errors.

Anthropic CEO Dario Amodei has issued perhaps the bluntest version of this warning: AI could wipe out half of all US entry-level white-collar jobs and spike unemployment to 10–20% within the next one to five years. He's urged both industry and government to stop "sugar-coating" what's coming and to prepare for a possible mass elimination of junior roles across technology, finance, law, consulting, and other white-collar professions.

Why Now?

Acceleration in AI Agent Capabilities

Frontier models' ability to autonomously execute complex tasks has been doubling roughly every seven months. At ~50% reliability, AI systems can now handle tasks taking a human ~4 minutes end-to-end. Within 2–4 years, they may reliably complete multi-day projects, encompassing exactly the kind of work traditionally assigned to junior employees.

Cultural Signals of Anxiety

The degree, once a reliable entry pass, is now viewed by many as insufficient on its own. When trust in the degree-to-career pathway is already weakened, a sharp pullback in entry-level hiring risks triggering a broader crisis.

Automation Risk is Uneven

Clerical, administrative, and certain analytical roles face steep decline, while fields like green energy, care work, and specialized technical services are projected to grow. The jobs being created are not mapping cleanly to the jobs being lost and often emerge in different regions, industries, and skill bands.

Demographic Cliff

The US is seeing a historic decline in its working-age population, with fewer young people entering college and the workforce each year. Traditionally, this would have meant labor shortages and wage pressure. However, in an AI economy, these demographic headwinds could actually offset some of the disruption from automation.

49%

of US Gen Z job seekers believe AI has reduced the value of their college education (Fortune)

~78M

Jobs projected as a net gain by 2030: 170M created vs. 92M displaced (World Economic Forum)

6M

Projected worker deficit by 2032 due to baby boomer retirements & shrinking demographics (Lightcast)

Predicted Impacts

Short Term

- **Augmented vs. Replaced:** Productivity tools will target the most automatable, repetitive entry-level tasks. Humans remain accountable for ethical decisions, client interactions, and nuanced problem-solving, while AI executes much of the operational load. Rather than performing the final deliverable, junior knowledge workers may shift toward roles that feed, monitor, and refine AI systems: preparing training data, supervising AI outputs, curating domain-specific knowledge, or providing the human judgment that bridges gaps in AI reasoning. In many ways, these roles look more like managing a team—except the "team" is made up of AI agents.
- **Short-cycle upskilling:** As degree value declines, employers will demand short-cycle upskilling, micro-credentials, and targeted training. "Job-ready" signals will matter more than broad academic qualifications. The risk is that, without coordinated standards, the credential market will become noisy and confusing, making it harder for both employers and workers to navigate. Even in today's market, only 1 in 8 credentials deliver material wage gains for workers, according to research by Burning Glass Institute.
- **Impacts on education professions:**
  - **In K-12**, classroom teaching may see a resurgence in recognition and value as one of the few "human-first" roles resistant to automation. Skills like building trust, motivating students, and reading emotional cues remain difficult for AI to replicate at scale. Isabelle Hau, in her book Love to Learn, argues that "relational intelligence" (RQ) deserves equal footing alongside IQ and EQ. In this sense, teaching is not just a role AI can't replace; it becomes a profession where relational intelligence is celebrated. Historically, economic downturns have driven talent back into the classroom. During the 2008 recession, for example, many professionals from other fields entered teaching as a stable, mission-driven career path. A similar pattern could emerge if workers displaced by AI seek the stability and social value of teaching.
  - **In higher education**, while teaching and research retain a strong human element, administrative functions are particularly exposed. The sector has long been criticized for administrative "bloat," with large layers of staff dedicated to enrollment, advising, compliance, and campus services. AI agents may be introduced to reduce headcount in these areas by automating student records management, financial aid processing, and even parts of academic advising.

- **In the corporate learning world**, AI promises to streamline L&D functions, particularly in automating content creation and assessment design. This will shift the role of L&D professionals away from manual program building and toward higher-order tasks: defining capability frameworks, ensuring cultural alignment, and curating AI-generated materials for accuracy and relevance.

Long Term

- **Corporate ladder bifurcation:** The traditional, linear climb from junior analyst to senior leader may fracture into two distinct tracks:
  - **The AI-augmented, nonlinear career**, where workers bypass traditional entry-level work and jump directly into mid-level, project-based roles alongside AI collaborators. These individuals will need to master orchestration—knowing what to delegate to AI, when to override it, and how to integrate its outputs into complex deliverables.
  - **The human-first profession**, concentrated in fields where trust, empathy, and embodied presence remain central—teaching, healthcare, therapy, social work.
- **AI-first apprenticeships:** As degrees lose their signaling power and micro-credentials flood the market, employers face growing uncertainty over how to gauge true skill levels. We could see the rise of AI-first apprenticeships, where new hires learn by overseeing and refining machine-generated work in real time. Workers build capabilities while already contributing value, guided by AI systems and human mentors. This model reintroduces the old idea of "earning while learning."
- **Potential skills erosion:** If entry-level workers only ever see polished AI outputs, there is concern that they may miss critical learning moments: the trial-and-error, debugging, and iterative problem-solving that build deep expertise. Over time, this could produce a generation of professionals adept at managing AI but less capable of doing the foundational work themselves if the AI fails. This could lead to a hollowing out of core competencies across industries.
  - To counteract this, institutions and employers may try to design intentional friction back into the system, creating spaces where humans still do the messy parts of the work. Even as AI becomes capable of executing end-to-end processes, embedding intentional learning loops (debugging exercises, manual checks, or shadow work alongside AI) may become essential to ensure future workers understand not just the outputs, but the flow of reasoning and problem-solving beneath them.

THEME 5:

Growing Privatization of Education

K-12

Thesis

EdTech innovation, combined with supportive policy changes, will accelerate the adoption and expansion of school choice options. This will be particularly evident in the form of alternative private education models.

Historical Context

Over the past five years, private school enrollment has grown both in the US and globally, supported by rising demand from families seeking personalized education and better-fit learning environments.

- According to the Cato Institute, the 2024-2025 academic year saw private school growth extend for a fourth consecutive year.
- NCES released its 2024 edition of the Report on the Condition of Education, which reflects 2021 enrollment data. Private US K-12 enrollment was 5% higher in Fall 2021 than in Fall 2011. Enrollment was relatively stable from 2019 to 2021, while public school enrollment declined by ~2% over the same period.

46%

of US private schools reported enrollment increases for '23-'24 (Cato Institute)

13%

Global growth in international private school student enrollment from 2020 to 2025 (ISC Research)

1M+

Participate in private school choice programs in the US as of 2024 (EdWeek)

Why Now?

Demand-side factors

In the US, the Education Savings Accounts (ESA) movement and other funding mechanisms have gained significant traction, enabling families to direct public funds toward tuition at private schools, microschools, or other alternative education providers.

- Over 1 million US students now participate in private school choice programs (ESAs, vouchers, tax-credit scholarships, etc.) as of 2024, up from ~500,000 in 2019. Twelve states and D.C. now offer universal or near-universal eligibility for private school choice.
- Iowa launched its ESA program in 2023, which drove an estimated 4,500–5,500 additional students to enroll in private schools—a 60% increase in private school enrollment statewide.

Globally, macroeconomic growth—especially the rise of the middle class—is increasing the number of families able to afford private education and creating sustained demand. Finally, the post-COVID normalization of nontraditional models such as homeschooling, microschools, and hybrid models has lowered the barrier to entry for new school formats and increased family choices.

Supply-side factors

On the supply side, technological advancements have lowered the operational cost of running schools, enabling the emergence of AI-native and tech-enabled private schools that offer more affordable options than traditional brick-and-mortar institutions.

For example, platforms like Outschool have made it easier for families to set up and manage homeschooling: identifying alternative educational resources, handling payments, and facilitating access to high-quality instructors online. In India, LEAD School equips affordable private schools with an integrated system for curriculum, teacher training, and assessments, dramatically reducing costs while raising instructional quality. These are just two among many products and providers that show how technology makes it easier to deliver high-quality instruction at scale.

Predicted Impacts

Short Term

- **Greater experimentation in school formats**, with a rapid turnover of models as communities test and refine new approaches to meet local needs. This will include schools that are very AI-forward, schools that will anchor in traditional teaching methods, and everything in between. The variety will also apply to pricing, from premium to ultra-affordable versions, particularly after factoring in ESA or student voucher related funding sources.
- **For parents, taxpayers, and policymakers**, the confusion and lack of transparency around the academic benefits of these new alternative schooling models will drive need for increased regulation, with states introducing mandated testing, accreditation requirements, or accountability measures to ensure quality and transparency.

Long Term

- In the US, the trend towards privatization will reshape the public education landscape by **drawing more students into affordable alternative options**.
  - Arizona is a key leader in the school choice movement, having supported both a robust charter school system and money for home schooling/private school vouchers. In 2024, 89,000 students received Empowerment Scholarship Accounts (ESAs), with an additional voucher program awarding 62,000 scholarships. Also, over 230,000 students attended charter schools. Only 75% of Arizona's students now attend traditional public schools, which is one of the lowest rates nationally.
  - In just one year, 20 public schools in the Phoenix area alone closed due to enrollment drops. (We note that declining birth rates, plus the increase in alternative schooling discussed here, will have additional impacts on public school enrollment.)
  - In one district in particular, Roosevelt Elementary School District, enrollment fell from 10,500 in 2010 to 7,000 in 2023. With school funding tied to student enrollment, as students left for alternative schools, Roosevelt faces a $5M annual budget deficit. It is considering closing between 7–9 of its remaining 25 schools.
  - The long-term impacts on the displaced students, impacted teachers, and staff—as well as the communities surrounding these closed schools—will remain to be seen.
- **The market will evolve into a dynamic "landscape for experimentation"** characterized by a wide variety of hyper-local, specialized microschools and hybrid programs that emerge and dissolve based on changing community priorities, parent demand, and policy shifts.
- **Alternative school models will increasingly position themselves as "future readiness,"** emphasizing skills such as adaptability, problem-solving, and digital literacy.
- **Blended learning environments will become more common**, integrating in-person instruction with online and AI-powered learning experiences to create flexible, personalized education pathways.

THEME 6:

Volatility as Education's New Normal

K-12

Thesis

Amidst growing budget and policy uncertainty at both the federal and state levels, school districts are pulling back on EdTech spending and procurement decisions—forcing vendors to demonstrate clear student outcomes, teacher value, and fiscal efficiency to remain in budgets. The effects of this volatility are cascading into classrooms, where both students and teachers are confronting real consequences.

Historical Context

The pandemic triggered a wave of rapid EdTech adoption across the country, as districts leveraged emergency ESSER funding to enable remote learning. Many tools were purchased reactively, now leading to bloated tech stacks. In the past two years, the AI boom has added further saturation, prompting the launch of new point solutions often with limited efficacy data.

2,700+

Tools used annually by the average district, with 1,400 accessed monthly (LearnPlatform)

Why Now?

District and school leaders are facing volatility from both a budget and policy standpoint, the effects of which cascade down into the classroom.

Budget Volatility
In the past few months alone, the administration froze over $6B in K-12 funding, covering after-school programs (Title IV-B), teacher training (Title II-A), student support (Title IV-A), ELL (Title III-A), and more, just before the July 1 anticipated disbursement. A few weeks later on July 18, the administration released $1.3B for after-school and summer programming. On July 25, they released an additional $5.5B covering the remaining delayed grant funds. Nevertheless, the initial delay caused widespread disruption and panic, postponing hiring, halting programming, and threatening staff layoffs.

Policy Ambiguity
From a policy perspective, debates over school choice, DEI mandates, student immigration status, and curriculum and content are adding additional ambiguity in district decision-making. District leaders are now forced to make tough tradeoffs—consolidating tech tools, delaying hiring, and scaling back non-core services—while still being held accountable for student outcomes. For many, this is also happening in the face of declining enrollments and, thus, budgets.

Teachers are also being asked to do more with less: address learning loss, re-engage disengaged students, and integrate new AI tools, all with fewer support services and minimal clarity on long-term funding. Students—especially ESL learners, undocumented students, and those from low-income communities—are at risk of losing access to services like counseling, SEL support, and after-school programs that were expanded during ESSER.

Predicted Impacts

Short Term

- **Core systems remain protected while tool consolidation accelerates:** Districts are maintaining spend on essentials like LMS, SIS, and cybersecurity, while cutting platforms that are redundant or lack evidence of impact to lower costs and reduce technical maintenance overhead.
- **New AI investments face higher scrutiny**: AI-powered grading or tutoring tools may see initial uptake, but reliable scale will require clear reductions in teacher burden and/or improvements in student outcomes.
    - Aiding school leaders are continued efforts to identify high-quality EdTech tools faster and cheaper, such as ISTE's procurement guidelines and EdTech Quality Indicators.

From 1,000 to 346
Denver Public Schools streamlined its digital tools, reducing complexity for teachers and students—saving $1M in the process. (EdWeek)

- **Procurement cycles slow dramatically:** According to Whiteboard Advisors, in late July, 53% of EdTech vendors reported districts were "in a holding pattern" and 28% had already seen cuts. About half expected deals to be delayed into August, with 22% facing 60+ day delays.
- **Teachers face growing pressure**: With uncertain budgets, fewer aides, and growing expectations, teachers are being stretched even thinner—all while being inundated with a plethora of AI tools advertising time-savings on key instructional tasks like lesson planning and grading or administrative tasks like attendance tracking.
    - It's worth noting that workers across many industries are also facing high degrees of uncertainty (mainly via the risk of AI job displacement discussed in Theme 4) while simultaneously tasked with achieving greater output and efficiency through AI tools. Simply put, the stress for many employees, including teachers, feels like it is reaching a boiling point.
- **Students lose access to auxiliary services:** SEL tools, tutoring programs, and after-school supports are at risk of being cut as districts refocus on core operations.

Long Term

- **Wave of EdTech consolidation**: Larger vendors will continue to acquire point solutions to offer integrated platforms.
- **Centralized state procurement:** Similar to textbook adoption cycles, states may implement "AI tool approved lists," requiring EdTech products to meet specific outcome, privacy, and interoperability standards.
- **Equity concerns rise**: Without protected, targeted funding, ESL students and undocumented learners may see support services eliminated, widening existing achievement gaps.
- **Teacher burnout and attrition may worsen**: While there was a drop in reported teacher burnout from 60% in 2024 to 53% in 2025, teachers still report almost 2x the levels of job-related stress to similar working adults. The lack of sustained funding and planning clarity may erode morale. In addition, the promise of AI-driven efficiency may not be fully felt, as teachers are asked to do more and more with these new tools, despite limited training and resources.
    - On the optimistic side, as mentioned in Theme 4, there is potential for a new wave of educators displaced by AI job automation, as teaching becomes not only recognized but appreciated as a profession where relational intelligence is critical.

THEME 7:

The New SEL: Navigating Human-AI Relationships

K-12 | HIGHER ED | CORPORATE & LIFELONG

Thesis

The rise of AI companionship is reframing not just how humans interact with technology, but how humans interact with each other. Generative AI can now remember, adapt, and respond over years, creating relationships that can feel deeply personal. As these bonds grow more sophisticated, they raise new questions about emotional health and the role schools will or should play in preparing kids for this new reality, such as "Should SEL evolve to include healthy human-AI relationships alongside human-human ones?"

Historical Context

In the digital era, technology has evolved from connecting people to one another to providing companionship itself. The internet allowed individuals to find their "tribe" online, especially those who felt isolated in their physical communities. In the 1990s, Tamagotchis popularized the idea of emotionally and "physically" caring for a virtual pet—but human bonding with machines predates all this. In 1966, MIT professor Joseph Weizenbaum created ELIZA, a chatbot simulating a Rogerian psychotherapist. He wrote that "some subjects have been very hard to convince that Eliza (with its present script) is not human."

Why Now?

Normalization of Chatbots

GenAI has transformed both human-machine and human-human interactions. Research by Common Sense Media shows that 33% of teens use AI companions for social interaction and relationships—with more concerning percentages highlighted below. Furthermore, an HBR article showed that therapy/companionship became the top GenAI use case in 2025.

72%

of teens surveyed have used AI chatbots

Data from Common Sense Media

31%

Find conversations with AI companions as satisfying or more satisfying than those with real-life friends

33%

of AI companion users have chosen to discuss important or serious matters with AI instead of real people

Improving Memory Systems, Agentic Capabilities & Customization Driving the Anthropomorphism of AI

AI model attachment is not only real but emotionally charged, as evident by some users grieving the loss of ChatGPT-4o when replaced by ChatGPT-5.

AI companions can now maintain lengthy conversations spanning months and potentially years, building detailed models of a user's personality, habits, and emotional state. In addition, users can customize ChatGPT and other models with "personality styles," giving them distinct, consistent tones and even personas. The Wall Street Journal reported that in June, 2.9% of conversations with Anthropic's Claude were "affective," meaning driven by emotional or psychological needs such as role-playing. These exchanges can be validating, with some users feeling "seen, heard, and remembered"—but also risk leading to delusion or over-reliance.

Predicted Impacts

Short Term

- **Normalization of AI companionship as the market floods with products targeting all age groups:** As of July 2025, AI companion apps have been downloaded 220 million times globally. This includes the emergence of AI chatbots for young children, marketed as "safe chat" friends. In 2025 to date, GSV has encountered over a dozen AI companions targeted at children as young as 4 years old. In combination with voice-based smart home devices, these multi-modal products are shaping a generation that sees AI conversation as normal.

- **AI literacy will expand to include safety and social-emotional dimensions:** Students should not only learn how AI works, but also how to engage with it responsibly, ethically, and with self-awareness. This means embedding AI safety concepts alongside relationship skills and even within SEL curricula. A practical model could mirror SexEd/Personal Safety education, where age-appropriate lessons—potentially delivered by outside experts—equip students with the judgement and habits to navigate AI in healthy, responsible ways.

Long Term

- **AI companions will become agentic:** Imagine an AI avatar that is self-acting and can adventure, develop, and play autonomously in AI-generated worlds, even while you are not actively engaging with it. For many, this could increase the risk of anthropomorphizing AI, creating further confusion and emotional FOMO when users are absent.
- **Societal norms may establish a hierarchy of relationships:** Human-human in real life, human-human digital, human-AI, and even AI-AI. There will also likely be continued debate and evolution of how these relationship categories are prioritized and even ranked.
- **Relational intelligence** could expand to cover healthy human-AI relationships alongside human-human ones.