




Roger
CAMRASS

FROM KNOWING TO BECOMING

REAUTHORING EDUCATION
FOR THE INTELLIGENT ERA

A summary of the discussion · 24 June 2026

When machines supply the answers, learning must change its purpose

The June briefing in the *AI and the Human Experience* series asked how education must change as artificial intelligence takes over cognitive work. Pairing a conceptual pre-read with a live panel of four, it traced a single shift: from an education built to transmit knowledge toward one built to form judgement, character and purpose — from *knowing* to *becoming*.

Few institutions carry the past as heavily as education does. The school and the university were among the great engineering achievements of the Industrial and Digital eras — systems designed, with remarkable success, to take in raw human potential and output disciplined, specialised, broadly interchangeable professionals fit for factories, bureaucracies and predictable corporate hierarchies. For two centuries that blueprint prioritised the transmission of information, the accumulation of knowledge and the measurement of both through standardised tests. It worked because the world it served rewarded exactly those things.

That settlement is now dissolving. As artificial intelligence moves beyond automating manual labour to performing advanced cognitive work — writing, coding, analysis, structured problem-solving — at superhuman speed and negligible cost, the foundational value of an education built around knowledge transfer is quietly eroding. When a machine can supply a competent answer to almost any well-posed question in seconds, the central purpose of learning can no longer be the acquisition of answers. It must pivot from the economic logic of “doing” towards the harder, more human work of “being”: the formation of judgement, character, purpose and the capacity to author one’s own life.

This was the territory our June briefing set out to explore. The pre-read circulated beforehand, *Transforming Education for the Intelligent Era*, framed the challenge as five structural shifts. The live discussion then tested those propositions against experience, with four panellists offering deliberately different vantage points.

The panel



Joshua Taylor

Recent A-level student
Weighing the real value of a university degree.



Mark Bridle

Incoming digital lead, Bootham School
Adopting technology without losing a school’s soul.



Tracey Jessup

Chief Transformation & Resource Officer, De Montfort University
The institutional machinery of higher education.



Prof. Alan Brown

University of Exeter
A candid look at the business model beneath it all.

A wider community of contributors — among them John Neill, formerly Chair and Founder of Unipart; Bootham headmaster Deneal Smith; Craig Walker; Andy Singleton; Asheesh Malhotra; Ric Merrifield and Trish Wheeler — enriched the conversation. What follows weaves the pre-read and the discussion into a single account, amplified where useful by recent evidence.

TABLE 1 — THE PRE-READ'S FIVE STRUCTURAL SHIFTS

1	Industrial infrastructure → personalised pathways	From age-based cohorts and fixed semesters to competency-based progression by demonstrated mastery.
2	Assessment → the “AI Mirror”	From examinations and grade points to reflective portfolios, narratives and developmental maps.
3	Re-authoring the institution	Schools become studios; universities become communities of inquiry rather than lecture halls.
4	The hybrid human-AI mentor	AI carries routine instruction; teachers are freed to cultivate the human capability stack.
5	The safeguard against de-agency	A single disciplined rule — systems suggest, but humans decide.

1.0 The dissolving blueprint

The first and most structural challenge is the one that least announces itself: the standardised, one-size-fits-all architecture of modern schooling. Industrial education sorts diverse young people into synchronised, age-based cohorts and moves them through fixed, time-based semesters, largely indifferent to the wide variation in how individuals develop. The pre-read argued that this systemic conformity routinely produces disengagement, anxiety and a thinning of personal identity, and that the alternative — genuinely personalised and competency-based pathways — requires abandoning cherished fictions such as “seat time” and fixed graduation dates in favour of progression by demonstrated mastery.

Joshua’s testimony made the abstraction concrete. He described using AI, in his case to revise for A-levels, not as a shortcut but as an accelerant: generating flashcards in moments, and producing clearer explanations of difficult material than his textbooks managed, including notoriously dense philosophical texts. Yet he was emphatic that the tool is only as good as the discipline around it. Used passively, it hollows out learning; used actively, it deepens engagement. That distinction — between dependency and partnership — turned out to be the hidden spine of the entire evening, recurring in almost every contribution that followed.

2.0 The vanishing first rung

If the purpose of education is shifting, so too is the destination it has long pointed towards: the graduate job. Joshua was open about the uncertainty facing his cohort. The traditional reasoning — work hard, earn the degree, secure the entry-level role and learn the trade on the job — is fraying precisely because the tasks that defined entry-level work are the ones AI now does most cheaply. Routine research, first-draft analysis, basic coding and administrative grind were the rungs by which juniors used to climb; they are also exactly what the technology automates first.

The recent evidence bears this out, while cautioning against fatalism. In the United Kingdom, the Institute of Student Employers reported that graduate recruitment fell by around eight per cent in 2024/25, with a further

seven per cent decline forecast, and that the contraction is concentrated among a small number of very large employers. Tech graduate roles have been hit hardest, falling by almost half against the previous year. A McKinsey analysis found that job advertisements for the kinds of occupations typically filled by graduates have dropped by roughly a third since mid-2022, and the World Economic Forum's Future of Jobs research records that around forty per cent of employers expect to reduce headcount where AI can absorb the work. Stanford's digital economy researchers have measured a meaningful fall in employment, specifically among 22- to 25-year-olds, in the most exposed occupations.

Yet the panel resisted a counsel of despair. Andy Singleton observed that some of the decline is cyclical and temporary, and that firms are reorganising around “AI-native” structures in which the technology handles the bulk of routine work and humans are redeployed to higher-value judgement — a shift that ultimately creates new openings even as it closes old ones. The consensus was that a broad degree retains real value, but that universities must stop training students for the very tasks machines have absorbed. The premium has moved decisively towards the human capabilities AI does not possess: communication, presentation, collaboration, the ability to frame a problem rather than merely solve it. Asheesh Malhotra pressed the corollary that a stubborn gap remains between graduating and being genuinely ready for work, and that curricula must be redesigned to close it.

**“Used passively, it hollows out learning;
used actively, it deepens engagement.”**

3.0 From policing to pedagogy: assessment and the AI Mirror

Nowhere is the old model more visibly under strain than in assessment. For centuries, academic success has been gatekept by examinations, grade-point metrics and timed recall — instruments that, the pre-read argued, measure conformity and memory far more reliably than they measure understanding. When a chatbot can produce a polished essay on demand, those instruments do not merely become unfair; they become meaningless, and institutions are tempted to pour their energy into detection rather than into rethinking what is worth assessing at all.

The panel was strikingly united in finding this the wrong priority. Tracey Jessup reframed AI-enabled cheating not as a unique civilisational crisis but as the latest entry in a long history of academic shortcuts, and described how De Montfort has responded constructively rather than punitively — deploying an AI-integrated tool that gives students immediate, formative feedback on drafts, effectively scaling the personalised support that human tutors cannot provide at volume. Alan Brown was blunter still, noting that too many institutions invest in surveillance software while leaving impersonal lectures and weak teaching untouched. Joshua, from the student side, argued that much cheating springs from a perceived lack of value in the curriculum itself; police the symptom and the disease persists.

The pre-read's constructive answer is the “AI Mirror”: persistent, contextual systems that observe a learner's curiosities, frustrations and patterns over time to build a developmental map, allowing examinations to give way

to reflective portfolios, personal narratives and developmental journals. The ambition is to teach students to read their own trajectories and discover what the pre-read calls their distinctive “human value streams” — the particular ways, whether through creativity, empathy, collaboration or systemic reflection, in which each person turns latent capacity into contribution. The hurdle is institutional inertia: persuading systems built to trust a number to trust instead a richer, more qualitative, and unavoidably more human picture of a person.

4.0 Metacognition and the de-agency trap

If one theme bound the briefing together, it was the danger the pre-read named most starkly: human de-agency. The argument runs that if educational AI is engineered only to maximise efficiency and engagement, it will smooth away the very friction — the struggle to remember, to synthesise, to judge — out of which an independent mind is forged. The true existential risk of the Intelligent Era, on this reading, is not machine domination but the quiet erosion of human agency, producing students who are technically competent yet psychologically passive, having outsourced original thought to the machine.

Recent research lends this warning unusual weight. A 2025 study from MIT’s Media Lab, “Your Brain on ChatGPT,” used EEG to track participants as they wrote essays with a large language model, with a search engine, or unaided. Those leaning on the model showed markedly lower neural engagement, weaker memory for what they had just written, and a diminished sense of ownership of the work. The researchers describe an accumulation of “cognitive debt.” The study is a small, pre-peer-review preprint and its authors are appropriately cautious; tellingly, they also found that students who first grappled with the material unaided and only then brought in AI showed the strongest engagement of all. The lesson is not abstinence but sequence: *the machine is a powerful finisher and a poor substitute for the initial cognitive work.*

The panel’s practitioners had reached the same conclusion by instinct. Mark Bridle stressed that students’ fluency with these tools already outpaces staff oversight, and that schools must therefore make metacognition explicit — teaching young people to understand and manage their own learning, so they engage with friction rather than flee it. John Neill put it most memorably: the core task is to “learn to learn,” and using AI to become passive is corrosive. The pre-read’s proposed safeguard is a single disciplined principle — *systems suggest, but humans decide* — operationalised by deliberately teaching students to interrogate AI outputs, recognise bias and misinformation, and treat cognitive effort as something to be cultivated, not avoided.

“The machine is a powerful finisher and a poor substitute for the initial cognitive work.”

5.0 From factories to studios, from lecture halls to communities of inquiry

Both the pre-read and the discussion converged on a deeper point: once AI absorbs routine information transfer, schools and universities can no longer justify themselves as distributors of content or gatekeepers of

credentials. They must become deliberate environments for human formation. The pre-read framed this as two parallel transformations: schools from factories into studios, universities from lecture halls into communities of inquiry, pointing to project-based models such as High Tech High, and to Cambridge's Digital Education Futures Initiative and its emphasis on "expanding dialogic space," using technology to provoke deeper collective questioning rather than passive absorption.

Mark Bridle's account of Bootham School embodied the studio ideal in a single phrase: *technology adoption must be "people-first,"* answering genuine human and cultural needs rather than deployed for its own sake. Headmaster Deneal Smith reinforced the point, describing a deliberately flexible curriculum — accelerated sciences, room for Latin and astronomy — alongside a defence of unfashionable but effective tools: pen and paper for memory, and the classroom as an irreducibly social, communal space. Drawing on visits to United States schools, he noted that anxieties about cheating are strikingly similar across the Atlantic, but that the American system felt more "liberated," leaning harder on oral presentation and spoken explanation — precisely the human-facing skills that resist automation.

Professor Alan Brown supplied the unsentimental counterpoint about why universities find this so hard. He described a business model in which loss-making departments are quietly subsidised by the high margins of business schools — large, often international MBA cohorts taught at scale — creating a financial incentive to pass students regardless of performance. AI, he argued, has merely surfaced long-standing pathologies: a lack of personalisation, teaching prized far below research in academic promotion, and students drifting away from impersonal lectures towards online resources while institutions respond by policing rather than reinventing. At board level, he found discussion too often stuck on the tactical — restricting tools, managing costs — rather than the strategic question of pedagogy. Others described how some are breaking the logjam from within: Trish Wheeler outlined work at Southern New Hampshire University to map core processes and operationalise agile course development within academic governance, while a contributor from the Hebrew University of Jerusalem described adopting a "product mindset" — building internal, AI-integrated knowledge platforms that bypass slow, committee-bound decision-making.

6.0 The lifelong imperative

If knowledge dates faster than ever, the logic of front-loading all formal education into a person's first two decades collapses. Craig Walker framed the problem precisely: university knowledge becomes obsolete quickly, demanding a model of continual, lifelong learning. Roger Camrass observed that this is also an opportunity — a route to sustainable revenue for institutions willing to serve learners over the whole of their careers rather than in a single intense burst. But the panel was clear-eyed about how far reality lags the rhetoric.

Tracey Jessup noted that while United Kingdom policy is shifting towards a lifelong-learning model, there is, as yet, little evidence of genuine market demand, and real uncertainty about how universities would deliver it. The policy landscape illustrates the gap. England's Lifelong Learning Entitlement — a flagship reform offering individuals a loan entitlement worth roughly four years of post-18 study, usable flexibly across modules and courses up to the age of sixty — has been repeatedly delayed; applications are now due to open in September 2026 for courses beginning in January 2027. History counsels caution about uptake: the Advanced Learner Loans it partly replaces saw demand fall sharply over the past decade. The same pattern appears

internationally; Singapore's much-admired learning-credit scheme has struggled with low usage, suggesting that the binding constraint is not funding but mindset — a deeply held sense that *learning is something one finishes*.

Micro-credentials, often proposed as the vehicle for lifelong learning, drew similar scepticism. Tracey observed that students are reluctant to take on further debt for credentials whose value is unproven, and that there is no settled market consensus on how to make them meaningful. Craig Walker argued they would need to be pan-European or global to match a mobile workforce. The most persuasive model on offer was older and more practical: John Neill described Unipart's "*learn in the morning, apply in the afternoon*" approach to operational excellence, and both he and Roger pointed to the integration of study with mandatory industry placements — as practised at Bath and within Unipart itself — as a proven route to genuinely employable graduates. In his vision of the future school, AI-driven platforms handle personalised, high-speed instruction, freeing teachers to act as coaches focused on character and collaboration.

7.0 The hybrid mentor and the human capability stack

That image — the teacher as coach rather than content-deliverer — is exactly the hybrid model the pre-read advanced. In it, the AI platform shoulders the operational burden of routine instruction, marking and scheduling, while the human mentor is liberated for the far more demanding work of ethics, reflection and the cultivation of latent potential. Crucially, this is not a diminished role but an enlarged one, and it depends on educators developing what the pre-read calls the "human capability stack": emotional intelligence, systemic thinking, ethical reasoning, and the cultivation of mindfulness and purpose — the developmental domains that resist automation precisely because they emerge from lived, embodied experience.

The candid difficulty, acknowledged on both sides, is that the workforce required for this does not yet exist at scale. Most teacher-training and doctoral programmes are still designed to produce subject lecturers, not developmental coaches or facilitators of reflection. Re-skilling an entire profession is a high-friction undertaking, and the panel did not pretend otherwise. Tracey Jessup added an essential equity dimension: genuine inclusion in an AI-rich system means meeting students where they are, not merely offering everyone the same nominal access. That entails actively teaching young people, especially those from less advantaged backgrounds, how to use these tools to their benefit, lest AI widen the very gaps it might have closed.

8.0 The mirror, not the master

Running beneath every strand of the discussion was a question about the technology's proper posture. The recurring image — the AI as a mirror — captures it well. A mirror reflects, reveals patterns, and shows us things about ourselves that we otherwise struggle to see; it does not decide on our behalf, and it does not replace the self that looks into it. The de-agency trap is, in essence, what happens when the mirror is mistaken for a master, when a tool built to reflect our thinking is quietly permitted to do our thinking for us.

Framed this way, the apparent tensions of the evening resolve into a single coherent principle. *AI in education should amplify human agency, not substitute for it*; it should personalise without isolating, accelerate without enfeebling, and inform judgement without supplanting it. The same principle disciplines every earlier theme: assessment that maps development rather than polices memory; pedagogy that teaches students to think critically about the machine; institutions that use automation to free teachers for formation rather than to cut

them; and a lifelong model in which the technology serves the learner across a whole life rather than capturing their attention for a season.

9.0 Toward a self-authoring generation

The pre-read closed with a contrast between two questions. The industrial paradigm asked a transactional one: how do we prepare human beings to serve economic systems? The Intelligent Era forces a more searching one: how do we help human beings become conscious, purposeful authors of their own lives? The June discussion suggested that the institutions we have inherited are, for now, far better equipped to answer the first than the second — constrained by financial models, regulatory caution, slow governance and the sheer gravity of habit.

Yet the mood was not pessimistic. Across a student revising with a chatbot, a school insisting on people-first technology, a university scaling formative feedback, and an industrialist describing learning fused with doing, the same conviction surfaced: the transition is navigable, provided we keep our eyes on the right prize. The danger is not that machines will do too much, but that, in letting them, we will ask too little of ourselves. The opportunity is the mirror image of that danger — an education that finally has permission to concentrate on what only humans can do, and to ask each young person not merely what they can produce, but who they intend to become.

That is the work ahead, and it will be the focus of the briefings to come. As ever, my thanks to our panellists and to the community whose contributions make these conversations what they are.

“The danger is not that machines will do too much, but that, in letting them, we will ask too little of ourselves.”

Roger Camrass

June 2026

About the programme

In 2025, Roger Camrass launched *AI and the Human Experience*, a global research programme supported by a community of contributors worldwide. Through bi-monthly briefings and a series of reports, it examines how artificial intelligence is reshaping every dimension of human life — and how we might steer it toward the greater good. Further detail, and the full library of reports, can be found at rogercamrass.com.

Roger Camrass

Research Director

Over half a century, Roger has advised governments on the strategic impact of technology, led a global multi-client programme, *Business in the Third Millennium*, and was research director of the Judge Business School survey into Corporate Innovation. He has helped create “innovation communities” between and within Fortune 1000 companies, and has held part-time academic posts at several universities and business schools, including the Judge Business School at Cambridge, Bristol University, the Ashridge School of Management and ESADE in Spain. Today he holds a visiting professorship in the humanities department of the Hebrew University in Jerusalem. He is married with three children and resides in Highgate, North London. See rogercamrass.com.

Sources and further reading

- Kos'myna, N. et al. (2025), “Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task,” MIT Media Lab (preprint).
- Institute of Student Employers (2025), graduate recruitment trends and forecasts, including the decline in tech graduate roles.
- McKinsey & Company (2025), analysis of UK job advertisements for graduate-type occupations since 2022.
- World Economic Forum, Future of Jobs Report 2025, on employer expectations of AI-driven workforce reduction.
- Stanford Digital Economy Lab (2025), employment effects of AI exposure among workers aged 22–25.
- UK Department for Education / House of Commons Library (2026), the Lifelong Learning Entitlement: scope, entitlement value and revised timeline (applications opening September 2026).
- Camrass, R. (2026), *Transforming Education for the Intelligent Era* (briefing pre-read) and *A New Vocabulary for the Intelligent Era* (April 2026 briefing).

