THE AI_NATIVE MBA

HOW BUSINESS SCHOOLS CAN USE AI TO EXTEND MANAGERS' JUDGMENT, KNOWLEDGE, AND SKILL

DICK STROUD

SUMMARY

THE ESSAY ARGUES that the MBA is a revealing test case for how higher education responds to generative AI.

Because business schools sit close to the labour market, they cannot treat AI merely as an academic integrity problem or as another topic to add through electives, labs, and policy statements.

Leading schools are already introducing AI-related courses and initiatives, but the deeper challenge is whether the MBA will be rebuilt around the fact that intelligent assistance is no longer peripheral to managerial work, but part of its operating environment.

The pressure is not only pedagogical. AI also threatens part of the MBA's labour-market rationale by making polished managerial analysis, synthesis, modelling, and recommendation-building less scarce than they once were.

An AI-native MBA should use AI to extend what students can do, so that broader research, deeper testing, sharper comparison, and more accountable judgment become the standard for the degree. Students should prepare more cases, test recommendations against objections and risks, organise human–AI teamwork, use multilingual evidence, manage attention under overload, and defend their own reasoning through oral and minimally assisted exercises.

The essay takes seriously the risks that AI may weaken memory, concentration, critical thought, and the credibility of academic signals. Its answer is not defensive prohibition, but structured use: AI as coach rather than crutch. If business schools merely decorate existing programmes with AI language, the MBA value will weaken. If they redesign around judgment, verification, prioritisation, and responsible tool use, the MBA could show that universities can still redefine excellence when technology changes what competence means.

THE PROBLEM AI EXPOSES

This essay sits within a wider concern I already have about academia and artificial intelligence. Universities can continue to function, award credentials, and speak in the language of standards even as the meaning of those standards becomes less clear. In earlier work I described that risk as 'continuation without correction': institutions preserve procedure and prestige while the degree's value as a signal of effort, competence, and readiness slowly weakens. Generative AI sharpens the problem because it loosens the old link between visible academic output and underlying capability.

The question, then, is not whether AI should be permitted, declared, contained, or

detected. It is whether universities are willing to rethink what they ask students to do and what they certify as evidence of mastery. Much of academia still responds defensively, as though the task were to preserve older demonstrations of competence for as long as possible rather than redefine excellence under altered conditions.

The MBA is a revealing place to test that argument because it sits adjacent to the labour market. It still presents itself as a bridge between academic study and business practice but now does so in a world where intelligent tools are becoming normal in both domains. If business schools cannot rethink the MBA in light of AI, there is little hope that disciplines further from market pressure will adapt. Hopefully, the MBA becomes an early demonstration that higher education can still raise standards when technology changes what capability looks like.

THE MBA AS A TEST CASE FOR ACADEMIA

The MBA matters not simply because it is prestigious or expensive, but because its meaning depends heavily on companies' needs. It is sold as preparation for leadership, strategic judgment, organisational decision-making, and career acceleration. That makes it a useful test case as the nature of management changes. Recent employer and recruiter research is broadly consistent on the point: AI capability is rising in importance, but so are the human capacities needed to use it well, including analytical judgment, leadership, and adaptability. (World Economic Forum; GMAC)

That combination is why the MBA is such a revealing credential. The challenge is not to choose between technical fluency and human judgment. It is to educate managers who can work in an environment where both are required simultaneously.

The public story about AI is often framed in terms of loss: too much intellectual work becomes easy, derivative, or optional. There is truth in that concern. Yet from the standpoint of business education, another truth matters just as much: the market is not waiting for universities to decide whether this change is philosophically comfortable. Employers already treat AI capability as part of professional readiness, candidates increasingly expect practical preparation, and graduate management programmes are responding, even if unevenly. Accenture recently announced that senior-level promotions will be linked to a manager’s AI usage.

A recent report from the Graduate Management Admission Council (GMAC) clearly captures the direction of travel: most graduate management programmes have integrated AI in some way, and candidates are especially interested in hands-on experiences that help them apply AI to strategy and decision-making.

That makes the MBA different from much of the rest of academia. In many disciplines, the pressure to adapt remains indirect and slow. In business education, it is more immediate, because the qualification is judged not only by academic norms but by whether it still maps onto the expectations of firms, recruiters, and students paying substantial sums for career acceleration. The MBA, therefore, sits at the point where the weakening of the academic signal and the strengthening of AI in business life collide most visibly.

THE LABOUR-MARKET THREAT TO THE MBA

The threat to the MBA is therefore not only an assessment and course redesign problem. It is also a labour-market problem. AI does not merely make it easier for students to produce work that looks competent. It also makes it easier for employers to obtain some of the work that once helped justify hiring MBA graduates in the first place.

That point should not be overstated. The MBA has never been valuable only because its graduates could analyse a case, build a model, prepare a PowerPoint deck, or produce a confident recommendation. Its stronger claim has always been broader than that: judgment, leadership, organisational understanding, and the capacity to act under uncertainty. Even so, the degree's labour-market value has depended heavily on a certain scarcity of polished managerial analysis. MBAs were attractive partly because they could take business data, structure it, compare alternatives, model consequences, and translate complexity into a recommendation that others could understand.

AI can now provide much of this value. It can summarise markets, draft strategic options, produce initial financial analysis, generate counterarguments, compare cases, translate foreign-language sources, and turn rough ideas into polished managerial prose. These tools are not reliable enough to replace human judgment, but they are capable enough to challenge the old assumption that this kind of analytical fluency is inherently scarce.

The consulting sector reveals the consequences of this issue. Much of the work associated with junior consultants, research, modelling, synthesis, benchmarking, slide logic, and preliminary recommendation-building overlaps with the kind of work for which the MBA has long been a preparation route. Recent commentary on consulting has therefore begun to ask not only how AI improves productivity, but how it changes the structure of professional-services firms when work once performed by large junior layers can be partly automated or compressed. The implication for business schools is uncomfortable: if AI changes the economics of the firms that traditionally hired MBAs, it also changes the rationale of the degree that supplied them.

Experiments with consultants using GPT-4 suggest that AI can raise performance substantially on some tasks while damaging performance on others when users trust the system outside its area of competence. That is exactly the jagged pattern business schools should care about. The threat is not that AI simply replaces managerial intelligence. The threat is that it makes competent-looking analysis abundant while making sound judgment harder to identify.

Nor does the evidence on demand for the MBA qualification support a simple story of its demise. MBA applications and enrolments vary by region, format, and school prestige. Some full-time programmes have seen renewed application growth, possibly caused by the tightening labour market. The broader enrolment data suggest pressure on the traditional MBA and growth in specialist or more flexible business degrees. The point is not that the MBA is disappearing; it's that the old generalist MBA can no longer assume that its prestige, cost, and traditional demonstrations of competence will remain self-explanatory.

This strengthens rather than weakens the case for an AI-native MBA. If the degree continues to present itself as a place where students learn to produce polished analysis, it will be competing against systems that make polished analysis cheap. If, however, it becomes a place where students learn to govern that abundance, to know when analysis is thin, when

evidence is misleading, when options are strategically irrelevant, when a model is outside its competence, and when an organisation needs commitment rather than more output, then the MBA's value can move upward. The degree's future claim cannot be that its graduates are better at sounding clever. It must be that they are better at judging what should be done when sounding clever has become easy.

These claims need to be tested against what leading business schools are doing, not against an abstract idea of the MBA. If the threat is that AI changes the value of managerial analysis, then the first question is whether the programmes that define the high-status MBA are redesigning the degree around that change or merely adding AI to the existing architecture.

ADDING AI IS NOT REDESIGN

To make that question concrete, this essay works with a defined sample: Wharton, Harvard Business School and MIT Sloan in the US; and London Business School, Cambridge Judge and Oxford Saïd in the UK. Using the QS Global MBA Rankings 2026 as a simple selection rule is not the only measure available, but it provides a clean way to choose influential, employer-facing programmes. The six schools are not the whole market, but they are influential enough to show what the high-status MBA is currently thought to be for.

The American schools offer a familiar picture with different emphases. Wharton offers a broad general-management MBA built around core studies, majors, and a large menu of electives. Harvard Business School remains tightly structured around its Required Curriculum, now including Data Science & AI for Leaders. MIT Sloan combines a rigorous core with action learning and intensive applied work. The details differ, but the shared claim is clear: the MBA is meant to form managers, not narrow specialists.

The British schools show the same broad architecture with different accents. London Business School stresses an experiential programme built from fundamentals, a flexible core, electives, and global experience. Cambridge Judge compresses the MBA into an intensive one-year structure with projects, specialist subjects and live client work. Oxford Saïd also combines a core curriculum with electives in an accelerated format. Across all three, the family resemblance is stronger than the differences: the MBA is treated as a broad managerial education rather than a technical silo.

That shared architecture matters because it shows what elite MBA education still believes it is for. These programmes are designed to train people to interpret incomplete information, move across functions, work through other people, and make decisions under pressure. Case discussion, participation, projects, and team-based work are not decorative extras. They are central to the degree's claim that it teaches the use of knowledge in action.

This is why the arrival of AI in these programmes is significant but still incomplete. Harvard has made Data Science & AI for Leaders part of the required curriculum. Wharton offers an Artificial Intelligence for Business major. MIT Sloan has developed applied work such as its Generative AI Lab. London Business School, Cambridge, and Oxford all offer AI-related electives, with Cambridge going further in introducing students to GenAI tools at orientation.

None of that should be dismissed. Elite business schools are not ignoring AI. Even so, a pattern appears when one looks at where AI sits. It is usually a course, a major, a lab, an elec-

tive, or an orientation layer inside a largely familiar MBA. AI is being added to the degree, sometimes seriously, but the degree itself is not yet being reimagined around the fact that intelligent assistance is now available across almost every managerial task.

That distinction matters because teaching students about AI is not the same as redesigning the conditions under which they prepare cases, conduct research, analyse markets, build recommendations, work in teams, and defend judgment. A required course can increase literacy without changing what the rest of the degree assumes about essential work. On the evidence of the official curriculum pages, the leading schools are moving toward integration, but treating AI as an important topic within the traditional MBA rather than as the condition that should redefine the qualification itself.

The issue is not whether schools can add more AI to what they already do. It is whether they are prepared to raise the standard of what the MBA demands across the degree as a whole.

THE ANXIETY ABOUT AI AND LEARNING

Before arguing that the MBA should become AI-native, it is necessary to take the objections seriously. If AI makes polished output easier to produce, it becomes harder to know what a student actually knows, understands or can do without assistance. That matters because the value of a degree depends not only on the experience of study, but on the credibility of the signal it sends about competence and readiness. Generative AI puts pressure on that signal by weakening the old connection between visible performance and underlying capability.

The risk is real. Students may retain less if they know information can be instantly retrieved. Concentration may suffer if skimming, outsourcing and rapid switching become normal habits. Integrity becomes harder to assess because the boundary between a student's work and the model's contribution is no longer obvious. But the greatest danger is not simply that students may cheat. It is that plausible solutions are created so quickly that students stop before the real thinking begins.

That is also why AI could become so valuable if business schools use it properly. Its best use is not to give students faster answers, smoother summaries or more polished prose. It is to give them near-constant access to a relentless processor of thought: a system that can test the quality, depth, structure and accuracy of an answer before the student mistakes fluency for understanding.

A dated but useful analogy is *The Paper Chase*, the 1970s film about James Hart, a first-year student at Harvard Law School, and Charles W. Kingsfield Jr., the terrifying contracts professor who dominates his education. Kingsfield does not linger in the memory because he gives students answers. He lingers because he refuses to let them settle for answers that merely sound plausible. He probes, interrupts and redirects until the weakness in a student's reasoning is exposed. What first appeared to be understanding turns out to be half-memory, loose assertion or confidence without command.

That is the part of the ordeal worth preserving. Kingsfield's cruelty is not the model to copy, and no business school should confuse intimidation with intellectual seriousness. But the mechanism beneath it remains powerful: the presence of a formidable mind that pushes students beyond the point where they would naturally have stopped. Used badly, AI does the opposite. It removes the struggle, supplies the answer and lets the student leave too early.

Used well, it can become a civilised Kingsfield: always available, endlessly patient, but intellectually unforgiving. It can ask the next question, demand the missing evidence, test the assumption, require the counterargument and expose the weak joint in the reasoning.

This is close to Nelson Dellis's distinction between using AI as a coach rather than a crutch. Skills weaken when too much of the work is outsourced. Used lazily, AI removes the friction that strengthens memory and reasoning. Used properly, it restores that friction in a new form: not by withholding help, but by refusing to let the student stop thinking too soon.

The emerging research points in the same direction. The evidence remains limited, but it does not support the simple claim that AI either destroys thought or automatically improves it. Heavy dependence on AI can be associated with weaker critical-thinking effort; more structured use can shift cognitive work towards verification, integration and task stewardship. If critical thought now lies partly in checking, integrating and governing outputs, then those capacities have to be taught and assessed directly.

The concentration question sharpens the issue. Human attention does not expand simply because machine assistance does. AI enters an environment already marked by interruption, overload and context-switching. It can relieve that pressure, but it can also intensify it by generating more tasks, drafts and partial lines of inquiry than a person can handle well.

This is where the MBA differs from some other parts of academia. In certain subjects, the response to AI may be a defensive attempt to preserve older forms of solitary performance. In the MBA, that move is harder to sustain because the degree already claims to prepare students for a world in which decisions are social, pressured, incomplete and increasingly AI-assisted. Yet the objection does not disappear. A management degree that embraces AI carelessly may produce graduates who are more fluent, more polished and less intellectually reliable.

CROSS-TOOL JUDGEMENT

There is a further complication. Business schools are not merely integrating AI through neutral systems; they are beginning to choose among rival visions of what educational AI is for. Different platforms place different weight on governance, workflow, learning support and guided use. The distinction should not be overstated, but it matters. A university that licenses one system rather than another is not making a purely technical choice. It is also selecting a bundle of assumptions about pedagogy, work and what good human–AI performance should look like.

The stronger answer, although operationally difficult, is to teach across a range of tools rather than achieving fluency within a single licensed environment. An AI-native MBA should not train students as if 'AI' were a single entity. It is already a mixed ecology of systems, each with different strengths, interfaces, defaults and temptations. The educational task is not to make the student loyal to one of them. It is to make the student capable of choosing between them.

For instance, a student preparing a market-entry analysis might use ChatGPT as a research and workflow environment: drawing on web search, uploaded documents, connected sources, data analysis and a working canvas to test the structure of an argument, interrogate assumptions and turn evidence into a draft model or memo.

The same student might then take the argument into Claude for a different purpose: to

turn the analysis into a cleaner document, pressure-test the narrative, build a prototype, or explore whether the argument holds together as a self-contained piece of work.

The point is not that ChatGPT is 'for research' and Claude is 'for writing'. That would be much too crude. Both systems can research, summarise, draft, challenge and analyse. One may be better suited to assembling sources, interrogating live information, using connected data and building a working analysis. Another may be better at reshaping a long argument, creating presentation materials, or exposing whether the final answer reads as a coherent judgment rather than assembled material. The skilled student should not merely ask, 'What answer did the model give me?' They should ask, 'Why did this model give me this kind of answer, what did its interface encourage me to do, what did it leave out, and what happens when I take the same problem somewhere else?'

That is where real AI literacy begins. A capable MBA student should be able to ask ChatGPT for a first-pass analysis of a merger, then ask Claude to attack the logic of the same recommendation. They should be able to ask one system to produce the most persuasive board paper for entering a market, and another to write the investment committee objection. They should be able to compare the outputs, identify where both systems converge, notice where they hallucinate or overstate, and decide what still requires human verification. They should know when to use AI to broaden inquiry, when to narrow it, when to demand citations, when to ask for a model or table, when to switch systems, and when to stop using AI altogether.

That is the point at which the redesign problem becomes more than a simple adoption issue. A business school that merely licenses one platform and teaches students where the buttons are has not created an AI-native MBA. It has created software familiarity. The harder and more valuable task is to teach cross-tool judgement: how to compare systems, question outputs, combine tools and remain responsible for the final decision. In the world MBA graduates are entering, competence will not lie in being fluent with one model. It will lie in knowing how to govern several.

If the MBA is to become genuinely AI-native, the redesign has to show up in the ordinary work of the degree. It is not enough to add a course, a lab or an approved platform. The standard of performance must rise across preparation, assignments, teamwork, evidence, attention and assessment.

MORE PREPARATION, NOT LESS

If AI reduces the time required for first-pass preparation, business schools should not merely bank the time saving and leave expectations unchanged. They should increase what they ask students to prepare. Tasks such as summarising a case, identifying stakeholder conflict, or drafting a first-cut analysis are no longer as scarce, slow, or solitary as they once were.

In a case-method environment, where the aim is to prepare students for time-sensitive decisions with limited information, AI should allow schools to simulate that condition more aggressively rather than relax it. Students could prepare several cases, using different scenarios rather than one polished first answer. The educational difficulty moves upward, away from basic preparation and toward comparison, prioritisation, and flexibility under uncertainty.

That shift matters because managerial life rarely presents itself as one neatly bounded

problem at a time. AI gives business schools an opportunity to train students not merely to arrive with a plausible answer, but to discriminate well among several plausible frames under pressure and in public. That is closer to the kind of performance the MBA already claims to cultivate.

FROM PLAUSIBLE ANSWERS TO TESTED PROPOSITIONS

The second implication concerns depth. One of the great dangers of generative AI in education is that it makes plausible answers cheap. A memo can look coherent before it has been properly tested. A strategy can sound fluent before its assumptions have been examined. The temptation will be to treat this mainly as an integrity problem. That is too narrow. The larger opportunity is to make plausibility the new baseline rather than the standard.

In practical terms, major MBA assignments should become harder to satisfy. Recommendations should arrive already stress-tested. Students should be expected to use AI to surface objections, alternative strategies, stakeholder resistance, downside cases, regulatory complications, and execution risks before they submit or defend their work. The educational aim changes from producing a recommendation to producing one that has survived a structured attack.

This is one of the places where the MBA could most obviously outperform a more defensive model of higher education. In a business environment where AI use is spreading quickly, there is little value in graduating managers who can generate polished first answers but have never been trained to interrogate them properly.

HUMAN-AI TEAMWORK AS A CORE MANAGERIAL SKILL

The third demand follows from something elite MBAs already do well: they treat teamwork as central rather than incidental. Group work, participation, and collaborative projects are not side features of the degree. They are part of what the MBA claims to prepare students for.

Yet once AI becomes normal infrastructure, ordinary group work is no longer enough. Teams should be assessed not only on what they produce together, but on how they organise human and AI labour. Who used AI to widen the evidence base? Who verified the claims it produced? Who challenged generic outputs? Who integrated conflicting results into a coherent recommendation? These are emerging managerial skills, not marginal digital tricks.

That means the group assignment itself changes meaning. The familiar difficulties of coordination, fairness, and conflict remain. But the AI-native MBA should add another layer: whether a team can build a sound division of cognitive labour between people and systems. Teams that use AI badly may simply converge faster on the answer. Teams that use it well can prepare more ambitiously, challenge themselves more systematically, and integrate divergent lines of thought more effectively. Success in tomorrow's business world will depend on who uses AI most effectively.

BEYOND AN ENGLISH-LANGUAGE MBA

The fourth demand has often been underplayed in traditional MBA design. Elite programmes often speak about global outlook and cross-cultural leadership, yet the evidence

base for classroom discussion remains disproportionately English-language. Historically that bias has been partly practical. Working seriously across languages has usually been too cumbersome for ordinary assignments. AI changes that.

Machine translation and multilingual search now make non-English source work far more feasible at scale. In strategy, marketing, regulation, supply-chain risk, or consumer insight, students can increasingly be expected to incorporate material from outside the English-language layer of the internet. That does not remove the need for judgment. It makes that judgment more important.

The skill here is not translation alone but interpretation under translation. Language carries institutional context, idiom, power relations, and cultural cues that do not move cleanly across borders. A serious MBA should therefore treat multilingual evidence work as a taught and assessed part of global managerial preparation rather than as an optional convenience.

PRESSURE, PRIORITISATION, AND THE MANAGEMENT OF ATTENTION

The fifth demand is less flattering. AI makes it easier to run multiple workstreams in parallel, but that does not mean the human mind has become better at managing them. Modern work is already fractured and overloaded. AI can help, but it can also worsen the problem by allowing more tasks, more drafts, and more options to proliferate than a person can handle well.

The AI-native MBA should therefore make the management of attention a visible learning objective. Students should have to decide what deserves deep attention, what should receive AI support, what can be deferred, and what should be abandoned.

That would be closer to managerial life than the tidier academic pattern in which every task is treated as though it had an equal claim on attention. The manager who thrives in an AI-rich environment will not be the person who merely touches the largest number of outputs. It will be the person who can distinguish between motion and judgment.

COACH, NOT CRUTCH

The final demand is the one that stops the whole argument from collapsing into productivity worship. An AI-native MBA should use AI to extend what students can do, while also preserving some forms of effort that remain central to cognitive autonomy. A student who uses AI to recall every fact, frame every paragraph, and resolve every uncertainty is not extending their mind. A student who uses it to quiz, challenge, provide hints, and expose blind spots is doing something different educationally.

This is where the AI-native MBA should remain deliberately unfashionable in one respect. Not every exercise should be frictionless. Oral defence is still essential because it forces ownership of reasoning. Live problem-framing shows whether a student can impose structure before the model does it for them. Short, bounded, minimally assisted exercises reveal what remains when scaffolding is stripped back. Reflection is equally important, since students must explain not just what they concluded, but how they arrived there.

Taken together, these demands point toward a different understanding of the degree. The lesson is not simply that human skills remain important. It is that the MBA should now be

built around the combination of intelligent tool use and the distinctly human capacities needed to judge, prioritise, challenge, integrate, and lead. The AI-native MBA, if it exists beyond slogan form, should therefore be harder than the traditional MBA in exactly the places that matter most.

CAN BUSINESS SCHOOLS ADAPT?

It is one thing to say that business schools now have an opportunity to raise expectations and simulate business reality more effectively. It is another to ask whether they have the resources, organisational capacity, inclination and will to do so. On the evidence now available, the answer is mixed. AACSB's recent work is encouraging, but it also shows that coordination remains uneven and execution is harder than aspiration.

That pattern reflects the wider structural problem that prompted this essay. Universities often recognise the need for change yet default to continuity because it is administratively easier, reputationally safer, and more aligned with existing incentives. The danger is not simple denial. It is a partial adaptation: enough movement to prove awareness, but not enough to fundamentally alter what the degree actually does.

Evidence from business schools suggests many are at this first stage of adaptation. Recent surveys show an optimism gap between deans and faculty, thin governance structures, and, above all, limited training time relative to the scale of the pedagogical change being discussed. That is not the profile of a sector redesigning itself quickly and evenly.

Mandatory AI training remains uncommon, and GMAC is blunt that many institutions still lack the resources and support needed for effective AI integration.

Employers increasingly want graduates who can use AI in strategy and decision-making, while prospective students increasingly want practical preparation rather than abstract reassurance. The pressure on the schools is coming from both sides at once.

The contradiction is clear. The MBA is exposed early to business pressure, but it still inherits the university's slower reflexes. That makes adaptation possible, but not automatic.

Where institutions have both the will and the means, the MBA could become the place where academia learns to respond properly to AI: not by defending a pre-AI benchmark of solitary performance, nor by surrendering standards to convenience, but by making the degree more demanding.

The strongest programmes are likely to move first because they have the brand confidence, faculty depth, employer relationships, and financial flexibility to experiment. Others may adopt the language of AI-native education while making only incremental changes. If that happens, the MBA market itself may reveal a sharper split between schools that genuinely redesign around AI-extended managerial performance and schools that merely decorate the traditional model with AI-friendly rhetoric.

CONCLUSION: THE DEGREE SIGNAL, RENEWED OR WEAKENED

The larger question, then, is not simply what happens to the MBA. It is what the MBA reveals about academia. If institutions respond mainly by defending the status quo, they will not preserve trust; they will spend it.

The MBA matters because it is one of the few degrees that still claims, explicitly and

expensively, to prepare people for the frontier of managerial work. If it is to keep making that claim credibly, it cannot define success as performing as though AI were absent, nor as fluency in whichever licensed product a university happens to choose. The stronger claim is that graduates can think, compare, verify and decide well across the AI systems used by the workplaces they enter.

If the MBA remains tied to pre-AI demonstrations of competence, its value will weaken. The degree cannot justify itself by teaching students to produce polished analysis in a world where polished analysis is increasingly cheap. Its future claim has to be higher: that its graduates can test evidence, compare tools, defend judgment and decide responsibly when intelligent assistance is everywhere. For the MBA, and perhaps for academia more broadly, this is not a decorative reform. It is the difference between renewal and managed decline.