~~THE SECRET~~ LIFE OF STUDENTS

Trained to stop learning?

How students are experiencing assessment and learning in an age of AI

Jim Dickinson (Associate Editor) and
Mack Marshall (Community and Policy Officer), Wonkhe

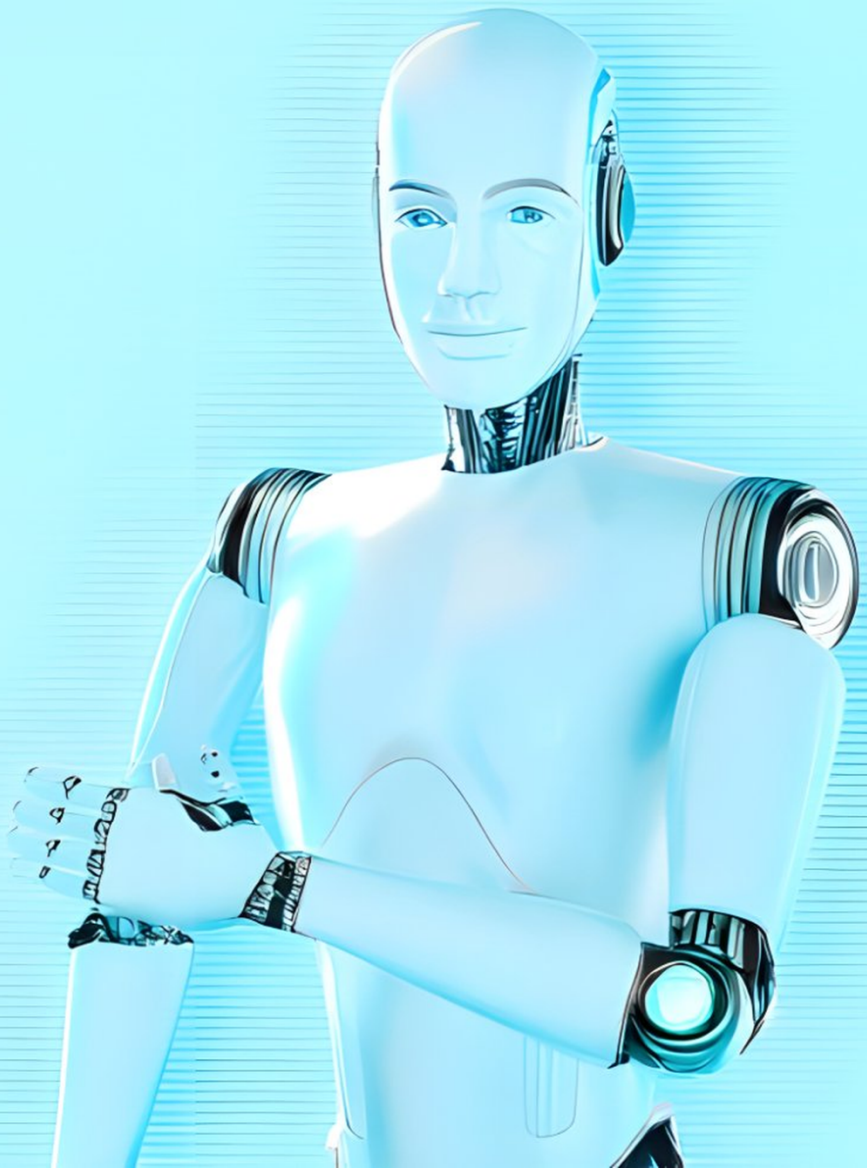

WONKHE

Trained to stop learning

How students are experiencing assessment and learning in an age of AI

"The difference in learning or not learning is a choice you can make now that AI is around. You didn't really have that choice before."

Contents

Introduction

We already know that the vast majority of UK students use AI, that most use it for assessed work, and that ChatGPT dominates. We know students reach for it primarily to save time and improve quality, that they worry about cheating accusations and hallucinations, and that most institutions now have a policy. That ground has been covered.

What the existing evidence has not addressed is the set of questions underneath the adoption statistics. Do students feel they have actually learned what they have produced? What are they weighing up when they decide how to use AI on a specific piece of work? Do they think their assessments test understanding – and if not, what would? Those are the questions we set out to answer in our new research.

The findings are uncomfortable. Students are submitting work they cannot fully explain, facing AI policies that do not function in practice, and responding strategically to assessment systems that reward production over understanding. The costs of that incoherence fall hardest on the most conscientious students, while disabled students are sourcing cognitive support through AI that they describe their universities as not providing – in a policy environment where that use is ambiguous at best.

But the findings are not only about institutional failure. They also reveal something the sector has not reckoned with – the seriousness with which many students are thinking about what AI means for their learning. Students in this research are not passively working through confusing policy. Many are actively constructing personal ethical frameworks – theories of practice about the relationship between tools, effort, learning, and identity – that are often more considered than the institutional guidance they receive. They are doing this work largely alone, with little support and no recognition. That intellectual and moral seriousness is itself a finding, and a resource the sector has barely begun to draw on.

The research identifies what appears to change student behaviour more reliably than any policy – the presence of a visible future accountability moment. Students who know they will need to demonstrate understanding describe using AI to test themselves, generate counter-arguments, and check their reasoning. Students without that downstream moment have no equivalent incentive. Assessment design emerges from this research as the factor most strongly associated with whether AI supports learning or replaces it.

The recommendations that follow are not about better detection or stricter policy. They are about whether higher education still asks students to be present – to think, to struggle, to understand – or whether it has settled for asking them to produce. That question predates AI. What AI has done is make the answer visible, because a system that cannot tell the difference between a student who understood and a student who produced something plausible was already in trouble. The sector's response should not start with what to do about AI. It should start with what assessment is for – and whether the answer it is currently giving is one it is willing to defend.

Headline findings

Finding 1: The gap between submitting work and understanding it predates AI – and assessment design is making it wider

Nearly half of students worry their grades don't reflect what they actually know, and 38 per cent admit submitting work they couldn't fully explain. AI has widened and accelerated a pre-existing gap, but the strongest predictors are assessment design features, not AI use itself.

Finding 2: AI use is not one thing, and treating it as one thing has left students without useful guidance

Students describe at least six distinct modes of AI use – from search replacement to production acceleration – each with different relationships to learning and different ethical implications. The same student routinely moves between modes on the same assignment, meaning any policy that treats AI use as a single behaviour will misfire.

Finding 3: Students know what properties genuine assessment should have, are rarely experiencing them, and have designed alternatives

Only 21 per cent feel their course primarily rewards thinking and reasoning, yet students across disciplines converge on what real assessment should look like – visible individual understanding, application, accountability, and developmental feedback. Many have designed specific alternatives that would test understanding, but none resemble a traditional essay submitted to a VLE.

Finding 4: Visible accountability moments are disappearing without replacement, and students use AI very differently when they know accountability is coming

Students who know they will need to demonstrate understanding in person use AI to test themselves, interrogate answers, and check their reasoning. When no downstream verification exists, the same students describe using AI on autopilot – the accountability moment changes not whether AI is used but how.

Finding 5: Discipline variation is structural, not incidental, and uniform policy will misfire in almost every context

The relationship between AI and learning differs fundamentally across creative arts, computing, healthcare, and humanities – shaped by professional identity, the nature of the knowledge, and students' ethical commitments to their field. A blanket institutional policy cannot account for this variation and will produce either meaningless restrictions or meaningless permissions in most disciplines.

Finding 6: AI policy is present in almost every institution and functional in almost none – and AI declarations may be making things worse

Students report tiered frameworks, declaration forms, and module-level guidance that varies between tutors on the same programme – none of it specific enough to tell them what to do when they sit down with an assignment. Declarations are penalising honest students while catching none of the heavy users, functioning as liability management rather than transparency mechanisms.

Finding 7: Policy incoherence is a distributional justice problem that consistently punishes the most conscientious students – and the cost is emotional as well as academic

The costs of unclear AI policy fall hardest on students most trying to comply – some have received lower grades for cautiously avoiding tools they couldn't confirm were permitted. Meanwhile, the most risk-tolerant students face no equivalent burden, and 59 per cent of all respondents worry AI could reduce their critical thinking.

Finding 8: Women are far less likely to use AI for assessment, and many carry anxiety about AI disadvantage without using it themselves

The gender gap in AI use is the largest demographic difference in the dataset – more than 20 percentage points – and does not disappear when controlling for other factors. Among non-users who worry about competitive disadvantage, 74 per cent are women, carrying the psychological costs of an AI-saturated environment without using the tools themselves.

Finding 9: Disabled students are using AI to meet needs that formal adjustments are not, in their experience, addressing
Students with dyslexia, ADHD, and related conditions describe AI as the most effective cognitive support they have encountered – often more useful than any formal university adjustment. Blanket AI restrictions risk removing this support in the name of academic integrity, in a context where the formal system has demonstrably failed to provide an equivalent.

Finding 10: Feedback timing is breaking the developmental function of assessment – and driving students towards AI
Feedback routinely arrives after students have started the next assignment, making the assessment sequence functionally summative regardless of what module handbooks claim. Where marking criteria are unclear, students turn to AI as a sensemaking tool – institutional opacity around expectations appears to drive AI adoption directly.

Finding 11: AI is compensating for gaps in institutional provision that students can see but that remain unaddressed
A substantial proportion of AI use described in this research would disappear if universities fixed problems within their control – inadequate library search, unclear briefs, inaccessible teaching, thin practical preparation, and slow academic support. Every heavy AI use case is diagnostic information about what the institution is not providing.

Finding 12: Time poverty is a structural driver of AI adoption that students are told is a moral or commitment failing
Students managing paid work, caring responsibilities, and unrealistic reading loads describe AI as the most efficient tool available for work they cannot bring themselves to care about. Where AI is expected or encouraged, the quality gap between paid and free tools is becoming a socioeconomic access issue that higher education should close, not widen.

Finding 13: Peer learning is a reliable route to genuine understanding and a largely untapped resource
Every moment of real learning described in focus groups involved other people – and every student asked whether explaining something to a peer felt intimidating or empowering said empowering, without exception. The unofficial curriculum of navigational knowledge flows through social networks that universities do not control, and formalising peer learning would distribute it more equitably.

Finding 14: Career-confident students are distinguished less by assessment format than by whether their course rewards thinking
The strongest correlates of career confidence are not vocational assessment formats but intellectual honesty – whether feedback develops thinking, whether stated values match actual rewards, and whether assessment tests understanding. Oral examinations have the strongest positive correlation with career confidence while placement-based outputs are among the weakest.

Finding 15: Students who feel they belong reach for AI less – because their course already provides what AI substitutes for
Belonging is one of the most powerful correlates of learning experience in the dataset, and students with stronger belonging are markedly less likely to use AI for assessments. Where belonging is absent, students experience their course as a production line and reach for AI accordingly – making the absence of resourced peer learning a structural driver of AI adoption.

Context

The existing research on students and AI is substantial. Work from the Higher Education Policy Institute/Kortext, Jisc, Chegg, and a range of academic studies has established that the vast majority of UK undergraduates now use AI tools, that most use them for assessed work, and that ChatGPT is the dominant tool by a wide margin.

Students report using AI primarily to save time, improve quality, and access support outside office hours. Their stated concerns centre on fear of cheating accusations, the risk of hallucinations, and worry about bias in AI outputs. These findings are established, and further replication would not advance the field.

The questions that remain unanswered sit below the adoption statistics. Do students feel they have actually learned what they have produced? What judgment calls do they make when deciding how to use AI on a specific piece of assessed work?

Do students think their current assessments test understanding – and what would authentic assessment look like to them? Do they feel their thinking process is valued, or only their outputs? What do students wish they had more time and space to learn that their courses do not currently prioritise? And what do students learn from each other about how to succeed that nobody officially tells them?

Both research exercises described in this report were designed to address those gaps. The survey asked first about learning and assessment before introducing questions about AI, so that responses about AI use would sit within a broader account of how students experience their education.

The focus groups were designed to surface the lived experience of assessment – how students decide to approach each assignment, the unofficial knowledge they share, and their vision of what meaningful learning could look like.

The policy context is defined by a tension that this research makes visible. Universities have moved rapidly to develop AI policies, but those policies have generally been designed around the question of what AI use is academically permitted rather than the prior question of what assessment is actually for. The result is a proliferation of declaration forms, tiered frameworks, and module-level guidance that varies from department to department and year to year – much of it operationally ineffective, as the findings make clear.

Methodology

Focus groups were conducted in February and March 2026, involving programme and departmental level student representatives from a range of disciplines, levels, and universities. The groups included undergraduates, postgraduates, and doctoral students across social sciences, arts, healthcare, science-related, creative technology, performance, and humanities backgrounds.

Each was structured around a common set of research questions developed in advance, with discussion supplemented by written responses and live contributions from a wider group of participants.

The survey was distributed via Wonkhe SUs subscriber SUs and comprised students from 52 HE providers across the UK. 1,055 students gave responses, weighted to account for gender and level of study.

The survey asked first about learning and assessment before introducing questions about AI. Quantitative analysis employed Spearman rank correlations, k-means cluster analysis, and principal component analysis. Fourteen open questions generated thematic analysis of free-text responses.

Who responded?

Sample profile at a glance

75%
Born in UK

73%
Living in England

39%
First generation

37%
Disabled/LDD

Gender

Women 62% · Men 32% · Non-binary 2.4% · Other/PNS 3.6%

Ethnicity

White 73% · Asian 14% · Mixed 5% · Black 3% · Other 5%

Age

18–21 53% · 22–25 24% · 25+ 20% · Other 3%

Level of study

UG/Bachelor's 67% · Taught PG 18% · Research 9% · Int. Masters 3% · Other 3%

Detailed findings

Finding 1 – The gap between submitting work and understanding it predates AI – and assessment design is making it wider

Our national survey asked students how they would feel if a tutor asked them to explain, without notes, the reasoning behind their last assessed submission. The responses suggest a large gap between what students submit and what they feel they understand.

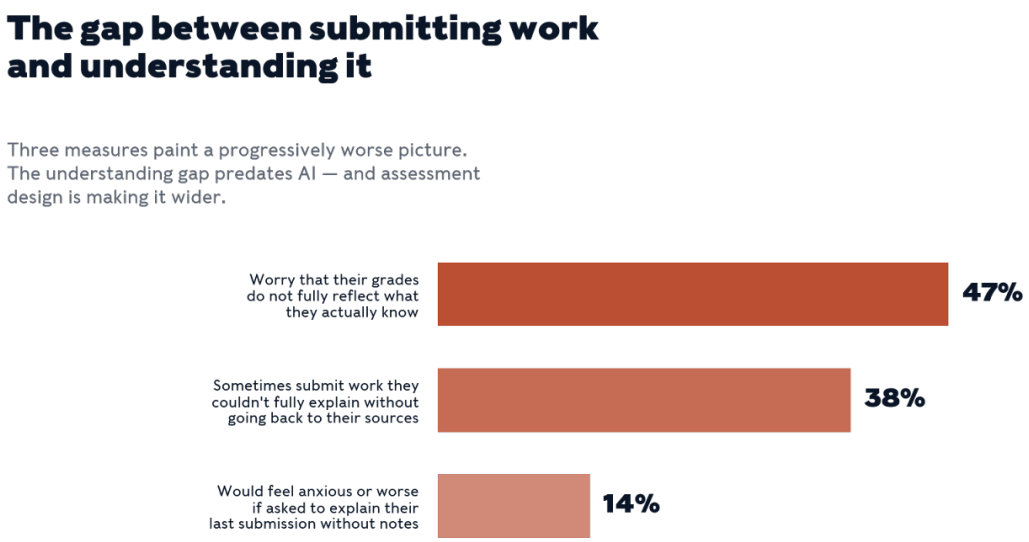

Fourteen per cent of students reported they would feel anxious or worse if asked to explain their last submission without notes. But a wider set of measures paints a worse picture. Thirty-eight per cent admitted they sometimes submit work they could not fully explain without going back to their sources. And nearly half the sample – 47 per cent – worry that their grades do not fully reflect what they actually know.

The gap between submitting work and knowing it is not new, and it does not appear to be primarily caused by AI. It has existed for as long as higher education has assessed learning by asking students to produce written artefacts. One focus group participant described watching peers operate this way years before AI existed:

> "Sometimes with theories or concepts, we'd read people's work and I'd ask them a question about it and they wouldn't be able to explain what it actually meant. They were able to write it in an interesting way, because they knew it would hit the rubric, hit the marks – they'd gamified it." – Undergraduate, social sciences

What AI seems to have done is widen that gap, accelerate it, and make the decision to stop learning explicit in a way it never was before. A student who wrote a mediocre essay in 2015 probably absorbed some understanding through the effort of producing it. A student who prompts an AI to structure and populate an essay in 2026 may absorb nothing at all.

When the survey tests statistically what predicts this pattern, AI use does not appear anywhere near the top. AI appears to have made it more convenient and harder to ignore.

> "We learn how to implement, rather than think – and if we graduated with only the stuff we learnt inside the curriculum, nothing good would come out of us." – Undergraduate, engineering

In the focus groups, students described their submitted work in terms of three distinct layers of understanding. The top layer – the main argument, the headline thesis, the case studies chosen – is almost always retained. The middle layer – the reasoning chain, the logic connecting evidence to claim – is where

things begin to fall apart. The bottom layer – specific phrasing, the how of the argument – is almost never owned.

> "I remember the main argument, but not how I went about arguing it, or how I interacted with the question." – Postgraduate taught, curating and art history

One computing student, asked to rate his confidence in explaining his last AI-assisted assignment, put a number on the gap: "Probably like a 4 or 5, if I'm being generous. Because that particular module was very much not my strong suit – I had to do a lot of writing. Everything I did was basically on autopilot. I wasn't actually physically applying myself in the assignment." He did not do well on it. The AI had not even produced a good grade – but nor had it produced any understanding.

A master's student averaging 80 per cent in their course described the calculation with unusual candour:

> "If I dislike the subject or don't find it particularly interesting, my main goal is to just get a good grade, and I will use AI as much as possible to help me – my human contribution is just checking against potential hallucinations, bias, and other things that might drag my grade down." – Postgraduate taught, finance

Another postgraduate student described using AI so heavily that gaps were visible in real time:

> "On paper it looks like I've touched on different theories and topics but if I were asked about them and to expand I wouldn't be able to. Yet the more those theories are mentioned as my thesis continues, I do sort of eventually learn about them." – Postgraduate research, education

But the understanding gap is not always about depth of intellectual engagement. In the focus groups, a more basic problem emerged alongside it – students who never understood what the assessment was asking in the first place. A biomedical science student described receiving a brief, a rubric, and a template for a single assignment – each of which said something different: "It was like three different assignments, and we had absolutely no idea what was going on." A paramedic science student described having "the lecturer, others on my cohort and AI explain what is required and for the life of me I could not understand it." A physics student got a good grade on a reflective assignment she still doesn't understand the purpose of – "it seems more for their benefit than mine." You cannot learn from an assessment whose purpose you cannot discern, regardless of how much AI you use or don't use.

There is a specific mechanism that appears to have been inadvertently built into assessment design. One student described feeding both the assignment brief and the rubric into a large language model, asking it to reverse-engineer a structure from the marking criteria, then populating that structure. The learning outcome becomes a formatting problem. The more precisely a rubric specifies what a good submission looks like, the more machine-readable it becomes – and the less intellectual effort is required from the student.

The grading problem is not experienced equally. The survey finds a large gap by disability status.

Open-text responses suggest this reflects assessment formats that do not suit different cognitive styles, conditions that penalise processing time, and a general sense that the measurement apparatus was not designed with disabled students in mind.

> "For a person who has dyslexia, having the rephrase ability in Grammarly helps me make my writing clearer – but in the back of my head I always worry it will impact my grade." – Undergraduate, humanities

Postgraduate taught students are a telling anomaly. They are the group most confident in their ability to explain their submitted work, and simultaneously the most likely to worry that their grades do not reflect their actual knowledge. This may not be the contradiction it looks. They appear to have worked out how to produce what the system rewards – which is why they also know their grades and their understanding have come apart.

> "It seems that each tutor's grading criteria are different. Some tutors gave everyone an A or B, while others gave everyone a D. Whether you achieve a distinction depends largely on luck rather than your hard work and ability." – Postgraduate taught, humanities

The consequences of the understanding gap were visible in the focus groups. A postgraduate student described a peer – a diligent international student who had worked hard, avoided AI, and received a lower second:

> "They told me they were now forced to use AI. Now they're in third year finishing their undergrad, and when I asked them if they'd learned anything, they said they couldn't be very certain – but now they're using AI, they're getting good grades, and those good grades will help them get into a good master's and eventually a good job." – Postgraduate taught, social sciences

Their conclusion was bleak and precise:

> "This is an international student paying so much money – and what's the outcome? They feel behind in the competition if they don't get good grades, so they're trapped."

The student who told this story added: "It was literally terrible. It was quite unsettling." That a peer was shaken not by the AI use itself but by what it revealed about the system – that a conscientious student had been trained to stop learning in order to succeed – captures something that the statistics alone do not.

This is not a student behaving badly. It is a student working through what the assessment system has put in front of them with the tools available. Their experience of their university's assessment system is that it trained them to stop learning in order to succeed.

Finding 2 – AI use is not one thing, and treating it as one thing has left students without useful guidance

For students, one of the most significant failures in current AI policy is the failure to distinguish between different kinds of AI use. Ask a group of students whether they use AI for their assessed work, and most will say yes.

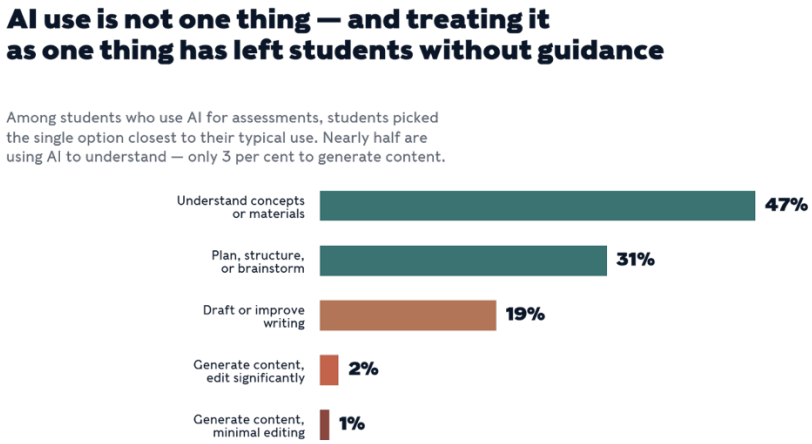

What is happening underneath that "yes" is at least six distinct activities, with different relationships to learning, different ethical weights, and different policy implications:

1. AI as search engine replacement. Students describe adopting AI as a research discovery tool not because they prefer it to library systems, but because library systems have failed them. The decisive advantage AI offers is conversational refinement. One humanities undergraduate, needing to produce an essay for a module where no teaching had been delivered, described how AI transformed a week-long scramble into a manageable process:

> "With a library service you can't go back and say, you've misunderstood me, I'm looking for this specifically. Being able to tell it off and make it better as it keeps going – that was what was really useful." – Undergraduate, humanities

She was still paranoid about hallucination, still checked every source was real, still built her own argument. The AI found the building materials. It did not design the building.

Several participants independently identified a worsening of the internet search environment as a structural AI driver:

> "I can ask an AI search engine for an open access, ultra specific topic in a set range of years and it will usually deliver exactly what I was looking for, or admit when there's a gap in research but recommend a wide range of related papers. I find it cuts the time I spend looking for resources in half and frees up more time to actually critically evaluate sources." – Undergraduate, environmental science

2. AI as structural scaffolding. The most commonly described use across the research is using AI to organise thinking the student has already done. Multiple students described developing their argument independently, then asking AI to propose a format – structure as a distinct skill from argument, one that AI can assist with without touching the intellectual content.

> "I'll do all the reading and work out what I think, but then I'll ask ChatGPT something like 'I've got these four points, what order should they go in?' It's not writing it for me – it's more like asking someone how to lay out a room when you've already bought the furniture." – Undergraduate, English literature

3. AI as debugging partner. For computing and quantitative students, AI sitting alongside work in real time functions more like a tutorial than a production tool:

> "It can't create the actual work for me, but it can correct where I went wrong, or if I've missed out something. Yesterday I was trying to build a graph, and I had missed out one of the prompts to do a round bar, and it told me to fix that and include it. I remember now." – Undergraduate, maths

The learning happens in the correction. This student was thinking in real time, not outsourcing.

4. AI as always-on tutor. This is the mode most often described as supporting learning, and the one current policy thinking is least equipped to accommodate – patient, available, non-judgemental, explaining the same concept six different ways, responding at 11pm:

> "I do enjoy the combo of both – sitting in class alongside some back-and-forth with Gemini AI. It's like having a personal tutor! I used to have a language tutor; now there's no need!" – Undergraduate, humanities

An aerospace engineering student described a more structured version – using AI not to get answers but to build the confidence to participate:

> "I used ChatGPT to let me know what I actually needed and what process I should follow. I asked it to question me in return and not give the solutions directly. So I felt confident, because it helped me to confidently ask my inputs to other teammates and get my work done and not look dumb!" – Undergraduate, aerospace engineering

A postgraduate student described a deliberate practice of doing the work first, then using AI to interrogate it:

> "I'd rather do some work and then show it to the AI to kind of validate it based on what's actually true. You don't just see it as 'get this work done and forget about it.'"

These are different modes from production acceleration – they support understanding rather than bypass it.

A related use emerged in the focus groups that does not fit neatly into any of the six modes – students using AI as a mock assessor. One computing student described completing a report, then submitting it to ChatGPT alongside the marking rubric and asking it to grade the work and identify gaps. He described his marks improving as a result. This is not tutoring in the traditional sense – it is a student creating their own accountability moment by using AI to simulate the assessment judgement before submission. It suggests that when formative feedback is not provided, students will build their own – and AI makes that possible in a way that was not available before.

5. AI as prerequisite substitute. A use that deserves close attention is students filling knowledge gaps that courses assume have already been closed. Students describe being admitted and recruited onto programmes that then take for granted foundations they don't have:

> "Half the module assumes you already know statistics – things like standard deviations and p-values – but I did psychology A-level, not maths. I got the grades they asked for and nobody said I'd need stats from day one. So I've been using ChatGPT to explain what the SPSS output actually means, because the lectures just move straight past it." – Undergraduate, psychology

A PhD student in environmental science described being expected to use R and MATLAB with inadequate teaching: "I often found that I was using AI to kind of tutor me through doing a lot of code and pointing out mistakes in my programming and stuff like that" – because "the explanations they gave us were just not good." Another focus group participant described the broader pattern bluntly: "Sometimes lecturers can't even speak properly within lectures so we would either have to use AI to understand the topic or YouTube."

AI is not only serving students who are cutting corners. It is also serving students who have arrived wanting to learn, who found that the course expected them to already know things it had never taught.

6. AI as production accelerator. This is the mode that students believe universities are worried about, and arguably the only one of the six that should warrant real concern about academic integrity. One student described feeding the assignment brief and the rubric into an LLM, asking it to produce a structure, working through it on autopilot, and acknowledging she "wasn't actually visibly applying myself." This is the mode associated with low understanding, low confidence in explaining the work, and – when the student reflects honestly – low learning.

One observation from the focus groups complicates the taxonomy. The six modes describe what students do with AI, but the same student routinely moves between modes on the same assignment or across assignments. A finance student who uses AI as a production accelerator on modules he dislikes is the same student who uses it only for brainstorming on modules he cares about.

An aerospace engineering student who used AI as a tutor to understand a process also used it to build confidence for social interaction with teammates. Students do not pick a mode and stick with it – they calibrate in real time based on interest, time pressure, the clarity of the brief, and their relationship with the material. Any policy framework that treats the six modes as types of students rather than types of moments will misfire.

Clusters

A cluster analysis of the survey data identifies four distinct student populations – each with a different relationship to their course, their assessment, and AI – and the critical policy insight is that the same AI behaviour in two different clusters can signal completely different things.

Savvy navigators (28 per cent). They've cracked the code – and know it needed cracking. Among the most confident in understanding their own work (71 per cent "very confident" vs 44 per cent overall), most positive about assessment quality, but also most aware of the hidden curriculum. They've decoded the system through peers and trial and error rather than official channels. AI usage is mainly for comprehension and planning, not production.

Disillusioned drafters (17 per cent). The smallest cluster, the loudest signal. They score lowest on every measure of assessment quality – 13 per cent say few or no assessments tested understanding this year (vs 4 per cent overall), and none of them say their course prioritises what matters. AI is their answer to a system they don't believe in – they're the most likely to use it for drafting, and the most likely to have used it in ways they weren't sure were acceptable. They also find the rules the least clear. Weakest community relationships in the sample.

Anxious producers (26 per cent). Getting through the work but not sure they understand it – only 15 per cent are "very confident" they understand what they submit. They score highest on submitting work they can't fully explain and worrying their grades don't reflect what they know. The youngest cluster (58 per cent aged 18–21), heavily first-year. AI is a comprehension crutch – 24 per cent use it to understand concepts (highest for that use type), and they'd be the least confident explaining that use to a tutor. What they want most is confidence in their own judgement. They want to feel capable, not just credentialled.

System naturals (28 per cent). The mirror image of the disillusioned drafters – everything works. Assessments test thinking, tutors care, feedback helps, and 60 per cent say their course already prioritises what matters. Lowest on every hidden curriculum measure. Many are mature postgraduates who've chosen their programme deliberately, and 81 per cent cite tutors as a source of learning – far above any other cluster. AI is peripheral – they're least likely to worry about falling behind by not using it. The system was designed for them.

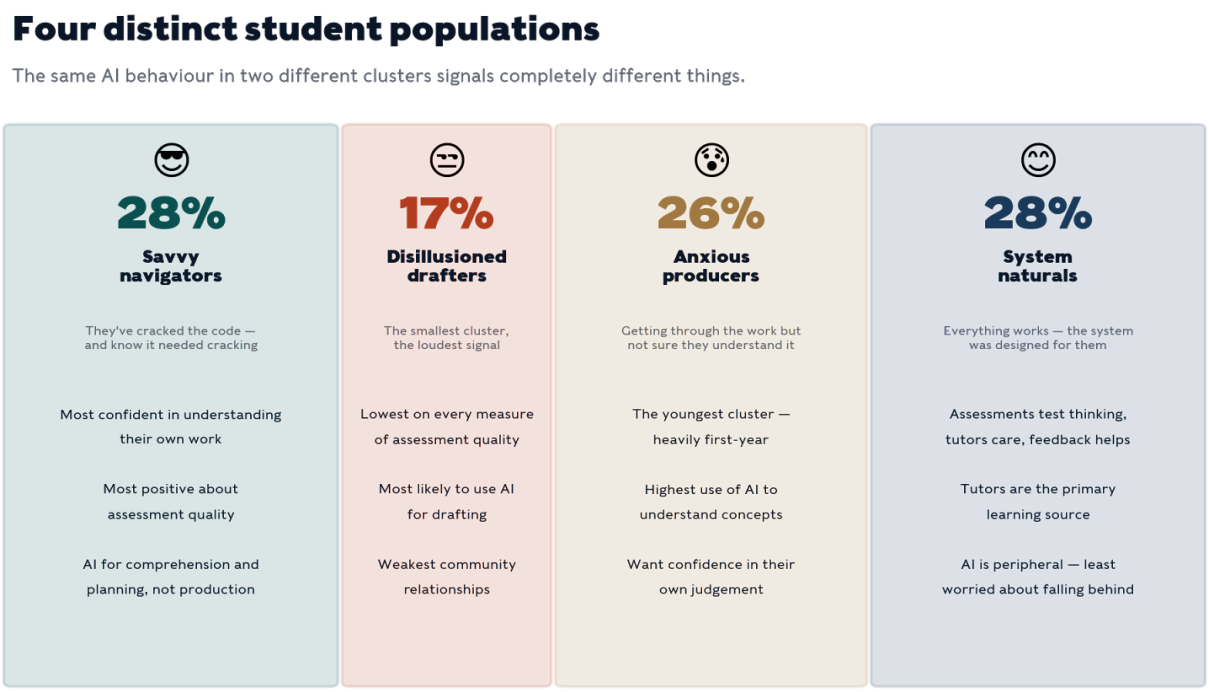

The typology matters for policy because savvy navigators and anxious producers show identical AI use rates but have arrived there via completely different routes and are getting completely different things from it. One group is augmenting. The other is compensating. Any policy – permissive or restrictive – that treats these two groups as a single population will either remove a crutch from students who need it or impose restrictions on students who are using AI exactly as institutions would want them to. And the same individual student may present as a savvy navigator on one module and a disillusioned drafter on another – depending on interest, clarity of expectations, and their relationship with the teaching.

Adoption

The survey data also tells us what is associated with higher adoption – students use more AI when they think the assessment system doesn't test what matters.

Students who feel their tutors care about their reasoning process use less AI. Students who think you can get good grades without really understanding the material use more. Students who see a gap between what their course claims to value and what it rewards use more. Students who believe their course primarily rewards outputs use AI at approximately twice the rate of students who feel thinking is valued.

And perceived fairness matters too – not just for how much AI students use, but for how anxious they are about it. Students who rate marking as unfair are far more worried about being left behind if they don't use AI than students who rate it as very fair – and the gap between those two groups (d = 0.86) is one of the largest effects in the survey. When marking feels unfair, AI stops feeling like a choice and starts feeling like a necessity.

In that light, AI use looks less like cheating and more like a strategic response to a system that students don't trust.

Finding 3 – Students know what properties genuine assessment should have, are rarely experiencing them, and have designed alternatives

Only 21 per cent of survey respondents feel their course primarily rewards thinking and reasoning. Thirty-three per cent agreed they could get good grades without deeply understanding the material.

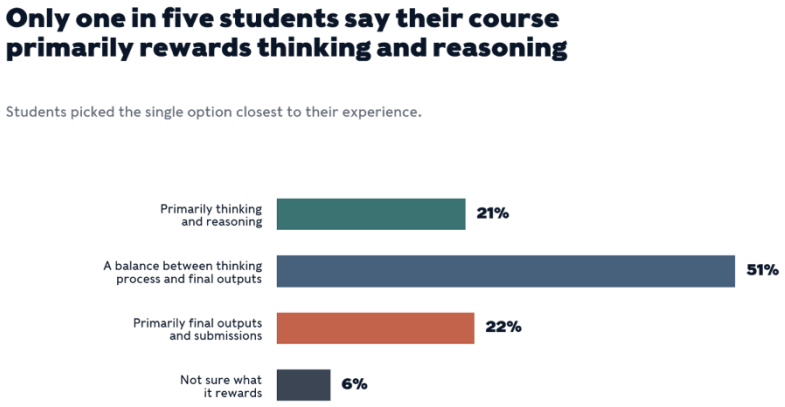

Two contradictions in the survey data stand out. Twenty-seven per cent simultaneously agree that their course develops their ability to think and that they could pass without deep understanding – encountering real intellectual development in some contexts and a gameable system in others, absorbing both lessons at once.

Two contradictions in the student experience

Students are absorbing both lessons at once — real intellectual development in some contexts, and a gameable system in others.

27%

simultaneously agree that their course develops their ability to think

+

that they could pass without deep understanding

Real development and a gameable system — absorbed at the same time

42%

rate themselves as confident learners who understand their work

+

agree their grades do not reflect what they actually know

The credential and the competence have come apart for nearly half the sample

Forty-two per cent rate themselves as very confident learners and agree their grades do not reflect their actual knowledge – nearly half the sample. The credential and the competence appear to have come apart for almost half the people in the room.

The survey also identifies which assessment formats students have encountered, and how much they trust each one to enable them to demonstrate their learning.

Different types of assessment are perceived as capable of showing off different levels of learning

Per cent confident each format can show what they've genuinely learned, among students who have used each format.

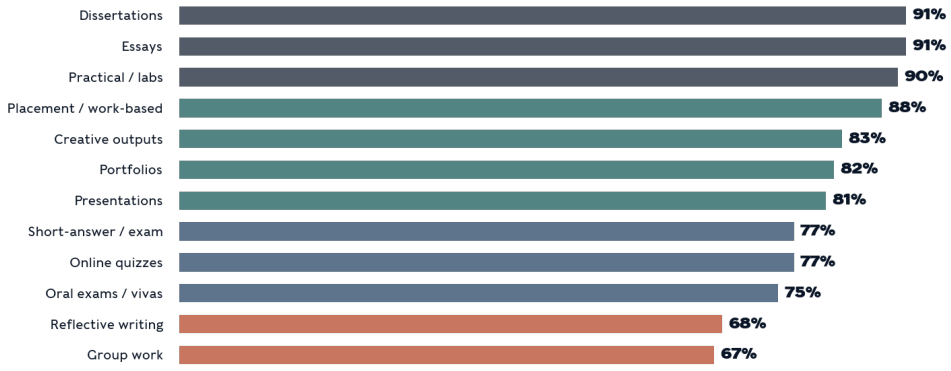

The essay paradox is worth dwelling on. Essays are the most widely experienced format (91 per cent of students) and among the most trusted (91.4 per cent confident they show their real learning). But students also know that essays are the format most vulnerable to AI. Students trust essays in principle as a vehicle for demonstrating thought – but they do not trust current essay regimes to verify that the thought is really there. You can produce a good essay without fully understanding what you've written, and AI makes that easier than ever. Students trust the format's ceiling, but the floor is very low. That is the design problem – an assessment format that rewards understanding when it's present but cannot detect when it's absent.

The focus groups were blunter about the floor than the survey. The gap between the 91 per cent survey confidence and the lived experience students described was one of the sharpest disconnects in the research:

> "The seminars were where the learning happened – that's where you actually got to debate and protect your ideas. I think the assignment was just the chance to show off how good you were at writing stylistically. I wouldn't say there was much learning done in the actual essay writing." – Undergraduate, humanities

> "I often find myself looking for quotes and references that I don't fully understand, but they get my point across. I also never fully read the article or book sometimes, as it is so time consuming." – Undergraduate, modern languages

These are not accounts of students using AI to cheat. They are accounts of how the essay format works even without AI – a format in which stylistic competence and strategic source selection can substitute for understanding. AI has not created that substitution. It has industrialised it. The survey may be capturing trust in the essay as a concept – "I could show my learning in an essay if I tried" – while the focus groups capture the experience – in practice, the essay rewards assembly and presentation at least as much as it rewards thought. A 91 per cent trust figure and a format that multiple students independently describe as producing no learning can coexist – because both are true, at different ends of the same distribution.

Group work scores lowest, and not because students think collaboration is pointless. The focus groups suggested that the problem is simpler – in a group output, individual understanding becomes invisible. A student who did all the thinking and a student who turned up to two meetings get the same grade. The assessment can't tell the difference, and everyone involved knows it.

Students told us as much repeatedly – "it normally ends up with one person doing the work for everyone" – and what they're objecting to isn't working with other people, it's being assessed on a product that doesn't reliably reflect what any one person actually learned. AI has added a new layer to this problem. In a group submission, students are now dealing not only with the familiar question of whether everyone has pulled their weight but the additional question of whether a group member's contribution was produced by AI. The ethical burden of policing a peer's AI use – without clear rules, without authority, and within a relationship you need to maintain for the rest of the project – falls on the most conscientious members of the group. It is the distributional injustice of Finding 7 reproduced in miniature inside every group assignment.

In the focus groups, students not only articulated what real assessment should feel like but had designed what would replace current methods. Real learning, across every discipline represented, is associated with the same set of properties – abstract knowledge meeting concrete application with real stakes, the presence of other people whose engagement makes going through the motions insufficient, personal stake in the material, and memorability that submitted work rarely achieves. Students have converged not on a single preferred format but on the characteristics they believe assessment should have – visibility of individual understanding, application, accountability, relevance, and developmental feedback. Their proposed routes to those properties vary – as the appendix shows, there is real disagreement about exams, presentations, and group work – but the underlying demands are stable.

> "It's not often I feel intellectually stimulated from my course – this has actually been disappointing as it's not what I thought university would be like." – Undergraduate, architecture

> "In terms of academic assessment, there have definitely been times I have written to meet the learning outcome and am unsure if I have fully understood the concept." – Undergraduate, nursing

Assessments that bear no obvious relationship to the discipline feel pointless – and students respond by clearing them as efficiently as possible:

> "I still don't fully understand what they wanted from me. I managed to get a good grade, and the only feedback I really had was maybe check your spelling. It seems more for their benefit than mine. It didn't actually feel like a proper reflection." – Undergraduate, nursing

> "There is no clear expectation set with our lecturers as to what they need to put – so it feels random and often like grades are given arbitrarily. If we are only assessed on Topic X using assessment type Y once in our whole degree, how does telling me what I did wrong help?" – Undergraduate, social sciences

> "I don't think traditional timed assessments are helpful – they promote a cram, write, forget attitude where you only learn what you think will come up." – Undergraduate, humanities

The specific formats students propose vary by discipline and personal preference, and there is no consensus on a single replacement model. But every proposal shares a common logic – the student must demonstrate understanding in conditions where producing something plausible is not sufficient. Not one of the formats they identify resembles a traditional essay submitted to a virtual learning environment.

> "It's fine having it in your head and writing for assignments – everyone can do that – but can you go out there and do it practically? Why isolate a student for an examination when they've got a team around them? In my profession you're never going to be on your own – even when you qualify, you've got a team around you. Working with other people reduces the pressure." – Undergraduate, nursing
>
> "You're not having to do three all-nighters working on an essay – you can just do a little bit at a time, but it's still written work. Everyone in my class has done it in a different format – some people have done it like a timeline, some like a notebook. You just have more freedom with it, and you can add photos, videos, voice notes." – Undergraduate, arts and media
>
> "Students can find something that they're really passionate about, or a really niche area in their subject. And you can become kind of like an expert at it." – Undergraduate, maths

Two focus group contributions described formats that already work. A design student described her course requiring a research folder of 100 to 200 pages with analysis on every single page – "so you're just constantly reflecting and looking back at your design choices throughout the entire project." The facilitator's response captures the significance: "The design of the assessment actually causes you to have to show your working." And an English literature student described an assessment where she had to refer back to a previous blog post and use it to develop her thinking further – iterative, self-referential, and challenging precisely because it could not be produced from scratch by AI.

One of the most interesting proposals came from a computing student who suggested assessment where students explain their work in what feels like a normal conversation – "you don't tell them that it's graded or anything. It can just be something that's like a normal conversation. And then at the end of the day, they find out that it was actually good. It's difficult for anyone to fail if you're really grading them based on what they understand, if they're actually saying stuff they're interested in."

> "I would redesign presentations to include discussion and viva elements, allowing students to display deeper knowledge that many know and just don't communicate." – Undergraduate, social sciences
>
> "When submitting essays, submit your notes alongside – so they can see how much work you've actually put in." – Undergraduate, humanities
>
> "I would include fake scenarios – like clients – and test the application of what we have learned. Something real." – Undergraduate, business
>
> "The single way to assess a person's understanding is an individual presentation on your essay or work in a Q&A manner. If you are confident in your work, regardless of your usage of AI, you will be fine in a Q&A. Of course you can factor in extensions and adjustments for people who require them – but on the whole, if you cannot answer questions about your own piece of work, then that piece of work is not yours." – Undergraduate, humanities

Every one of these proposals is harder to scale than the formats they would replace. Students are clear-eyed about this – a finance student named it directly, saying universities have become factories producing papers that are never printed, submitted to a VLE en masse, so structured and formatted that assessment has become what it always was. The response to scalability concerns is not to abandon the direction of travel. It is to design carefully for it.

The survey captures what students wish their courses were developing in them – a question asked before any mention of AI, and answered with consistency.

What students wish their course helped them develop

Select up to five. A further 34% said their course already prioritises what matters.

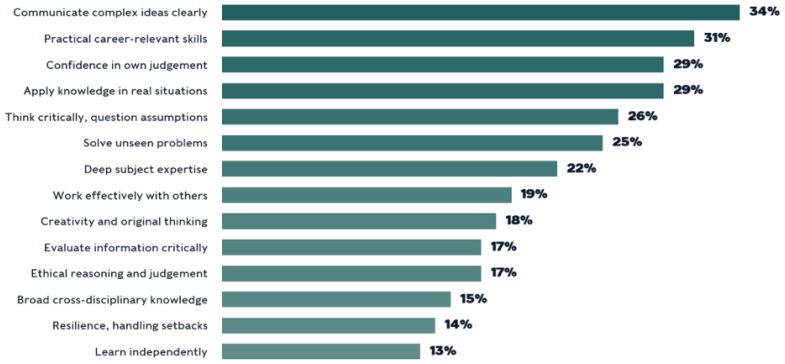

The pattern is not "give us harder academic challenges." It is something closer to – help us become people who can do things in the world. The desire for confidence in one's own judgement, named by 29 per cent, sits directly against a system that elsewhere in the survey students describe as primarily rewarding conformity to tutor preferences and rubric compliance.

One of the strongest cross-survey relationships was students who believe they could pass without deep understanding are nearly three times more likely to say AI has made them question what they are actually learning, at 61 per cent against 24 per cent.

The relationship runs both ways

Believing you can pass without understanding and questioning what you're learning are tightly linked — in both directions.

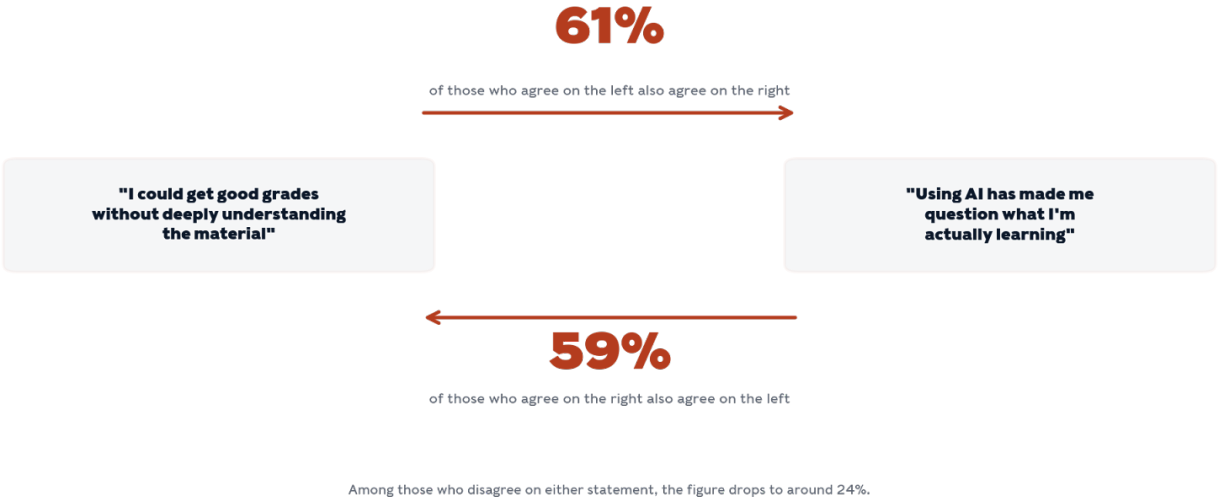

Among those who disagree on either statement, the figure drops to around 24%.

Final-year students feel this most sharply. AI has not created the problem – it seems to have made an existing problem impossible to maintain comfortable distance from.

> "I would emphasise presenting work to a group, explaining complex ideas in my own words, demonstrating original and critical thinking. But it could be triangulated with another assessment –

because I do have social anxiety and depression so on some days I may not be as good at presenting. Both matter." – Undergraduate, social sciences

Finding 4 – Visible accountability moments are disappearing without replacement, and students use AI very differently when they know accountability is coming

One of the most important findings emerged not from a question about exams but from a straightforward discussion about how students approach their current work. Two students at the same university, in the same focus group, on the same morning – one a maths student whose module had an end-of-semester exam, the other a computing student whose module was assessed entirely by coursework. The maths student approached her AI-assisted work with care. The computing student had used AI to structure her entire assignment on "autopilot" (her word) and not done well. The difference between them was one thing – whether a future moment existed at which they would need to demonstrate understanding.

> "In the back of my mind, I know in a few months I'm going to have to sit an exam and get tested on similar stuff. So I do need to actually study and do my own assignment instead of just allowing AI to carry it for me." – Undergraduate, maths

Asked whether the same applied, the computing student replied simply: "No, it doesn't, no." She wasn't describing a moral failure. She was describing a strategic response to a system that had removed the only moment at which understanding would be tested.

The accountability moment doesn't stop students using AI. It changes how they use it. Students who know they'll be tested on the material later don't just avoid AI – they prompt it more carefully, interrogate its answers, go back and forth, push on the reasoning rather than accept the first output. They use it to explore the material rather than to bypass it.

> "I use AI completely differently depending on whether there's an exam at the end. If I know I'm going to be sat in a room with a question and no laptop, I'll get it to explain things to me, then I'll argue with it, then I'll get it to quiz me until I can actually do it myself. If it's just coursework – I just want it to help me get the thing finished." – Undergraduate, biomedical science

The difference isn't AI or no AI. It's whether the student has a reason to make sure they've actually understood what AI has helped them produce. When a future verification moment is visible, students describe something that looks like good learning – questioning, practising, checking their own understanding. When nothing downstream requires understanding, the same students describe something closer to production.

The gap between what a student has produced and what they actually know only becomes visible when they are required to account for the work in person. A PhD student discovered this the hard way:

> "When the day of presenting came I was so anxious and unsure, and then I realised it was because I had depended on GPT to know for me – I was telling myself 'why can I not remember my script or even my topic that I've done so much work on?'" – Postgraduate research, art and design

The anxiety wasn't about the presentation. It was about the gap between what she'd produced and what she actually knew – a gap she only discovered when she had to account for the work in person. She noticed "a massive difference at the next presentation when I didn't use it at all and felt much more confident."

This is also where one of the facilitators made an unsolicited confession. Describing his own dissertation experience, he said: "The moments where I felt best at being able to explain what I'd actually submitted was ironically when I ended up having to present my dissertation in a presentation format – only really in that process of putting it together as a presentation did I actually start to clock and understand what I'd actually been writing about for 10,000 words." The act of explaining is not just a way of testing understanding. It appears to be a mechanism for producing it.

But the focus groups also introduced an important complication. Students want these accountability moments – and they are also afraid of them. A mature student who had experienced an interview-style assessment reported that "a lot of students were distressed about this format" and noted: "I'm also bad at articulating myself verbally as opposed to on paper, so I don't feel the assessment was reflective of my knowledge."

An architecture student drew a careful distinction between a one-to-one conversation (which she wanted) and a presentation (which she found nerve-racking): "I'm not saying this should be a one-sided presentation – this is not how the ideal assessment should happen." And another student, holding both sides of the tension at once, said: "I do have social anxiety and depression, so on some days I may not be as good at presenting. Both matter."

None of this is an argument for more exams. Students describe timed, closed-book exams as tests of memory rather than understanding – formats that are, as one participant put it, "set up against you."

Exams are one instance of a broader category – visible accountability moments at which students know they will need to demonstrate understanding without assistance. The data suggests that the presence of such a moment changes how students behave long before the moment itself arrives. It changes how they study, how they read, and how they prompt. Some courses have removed or weakened these moments without replacing them with equivalent verification – and when that happens, the incentive to use AI well is removed along with the incentive to use it less.

Any course team moving away from exams needs to answer what is replacing the accountability function, not just what is replacing the format. And the signal needs to be visible from the start – so it shapes how students use AI throughout, not just whether they use it at all. The design challenge is to build accountability moments that are experienced as developmental conversations rather than surveillance events – and that are accessible to students for whom live performance under pressure is itself a barrier.

Finding 5 – Discipline variation is structural, not incidental, and uniform policy will misfire in almost every context

A course representative responsible for around 200 students across writing, art, programming, computer science, and animation described what happened when all of them were set the same theory assignment.

Artists struggled because it asked them to operate in a logical register that did not match their professional identity. Programmers struggled because essay writing was not their strength. And each group's response to that struggle mapped directly onto their disciplinary relationship with AI:

> "Artists generally are very against AI, and I completely understand that. Programmers are very for it, because it's logical and it's probably where the workplace will develop – it's at the forefront of what they're doing. Design is somewhere in between, because we do both art-related and programming-related and logic-related things." – Undergraduate, creative technology

A software engineering rep described a different kind of disciplinary relationship with AI – one where engagement with the tool is part of demonstrating real competence rather than circumventing it:

> "We need to create systems – we can't just copy and paste the code. We need to analyse the code, check what's happening, know what we can do to make it better. I know the names I use in variables, strings, everything – which a person who just copies and pastes will never know, because the AI created the name." – Undergraduate, software engineering

The contrast with a creative technology student's reasoning about foundational knowledge is worth noting:

> "I believe it's better to learn the skill from the ground up, because it gives you a choice later on. You either have the skills to deliver it yourself, or you have the skills to better utilise the AI should you want to. You know the terminology, the output, the pipeline of what you're supposed to be making – so you can catch out the AI better once you know that foundational knowledge." – Undergraduate, creative technology

Both positions are coherent and defensible. They apply to different disciplines, different professional trajectories, and different questions about what a course exists to develop. The sector-level debate about AI as a graduate skill reflects this variation sharply. Students were divided in focus groups:

> "AI is a skill like using Excel and should 100 per cent be taught – if you don't, you will be behind in industry." – Undergraduate, events management
>
> "I am going to have to respectfully disagree that one may be behind in their industry. I really think it depends on the industry." – Undergraduate, social sciences

"What are we at university for if not for special learning?" – Undergraduate, creative writing

But the disciplinary divide is not only structural – for students on creative and practice-based courses, it is a question of professional identity and personal ethics that goes beyond assessment design. One arts student said plainly that AI "has taken away our jobs, basically – which is why I wouldn't use it." A digital media design student wrote in the session chat: "AI art is terrible for the industry." A creative writing student described AI as "a very dangerous thin end of a wedge" for creative courses and submitted a non-use declaration with her assessments.

The cultural dimension was visible in an unguarded moment in the session chat. When one participant mentioned the volume of AI-generated Christmas cards she had received, another described her mother giving her "AI-generated birthday card and framed AI artwork" for her birthday. A third compared the ubiquity of AI to "a bad episode of Black Mirror." For these students, AI is not an abstract policy question – it is something that is happening to their profession, their craft, and their daily aesthetic environment. A student union officer reported that "some students have shown reluctance to using AI and are intentionally not applying to some unis which encourage it within module specs." AI policy is affecting recruitment.

Students are also making sharp distinctions between modules they experience as core to their degree and modules that feel inserted to bulk it out – especially where elective choice has been replaced with core modules that they have little commitment to.

A foundation-year business report – twelve pages covering costings, financial figures, and customer personas, submitted by students who had enrolled to study creative technology – was described as a driver of near-total AI use among students who could see no connection between that work and the discipline they came for.

The relationship between perceived relevance and real engagement is consistent across the data. A blanket institutional AI policy cannot account for this variation.

Finding 6 – AI policy is present in almost every institution and functional in almost none – and AI declarations may be making things worse

The dominant experience reported by students was not the absence of AI policy but the presence of policy that does not function in practice.

Participants told us that universities have declaration forms, tiered frameworks, and module-level guidance. What they do not have is policy specific enough, consistent enough, or coherent enough to tell a student what to do when they sit down with an assignment and wonder whether to open a browser tab.

When students describe what they actually consider when deciding how to use AI on a specific assignment, the picture that emerges is far more principled than most institutional approaches assume.

Ethics and learning value dominate students' AI decisions — detection risk is near the bottom

Select all that apply. Among AI users only.

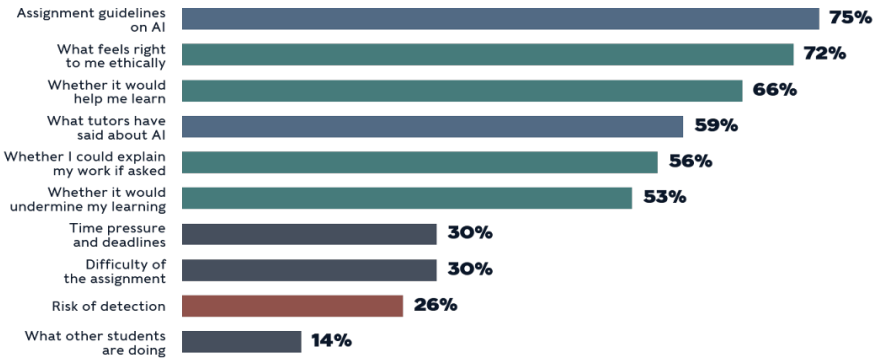

Ethics and learning value far outrank detection risk and peer behaviour as decision factors. Policies built around detection are answering a question most students are not asking.

But many students have gone further than working through confusing policy. They are constructing sophisticated ethical frameworks of their own – often more considered than anything their institution has produced. A computer science student described his governing principle: "I make sure my use of AI doesn't inhibit my understanding of the topic. For essay submissions, the final text is written by myself – I don't want to lose the ability to report on my findings." An engineering student had arrived at a principle of "augmentation, not replacement" – using AI for repetitive tasks while retaining responsibility for core logic and final validation. A graphic design student drew a line that was philosophically precise: "Using AI to do a final work for me, I say no, but to help me make a final work as a tutor or a supporter or a friend, I will say yes."

These are not students waiting for better policy. They have outrun it. But they are doing this intellectual and moral work largely in isolation – and often in silence. One engineering student said she hadn't heard others discuss their AI use "as it might look like cheating or like they do not understand the assignment." The furtiveness is itself a cost. Universities have an untapped resource in students' own ethical reasoning – and creating structured opportunities for that reasoning to be surfaced and shared would do more than another round of policy revision.

How clear is the line between acceptable and unacceptable AI use on your course?

Students picked the single option closest to their experience.
Nearly a third are in grey-area territory or worse.

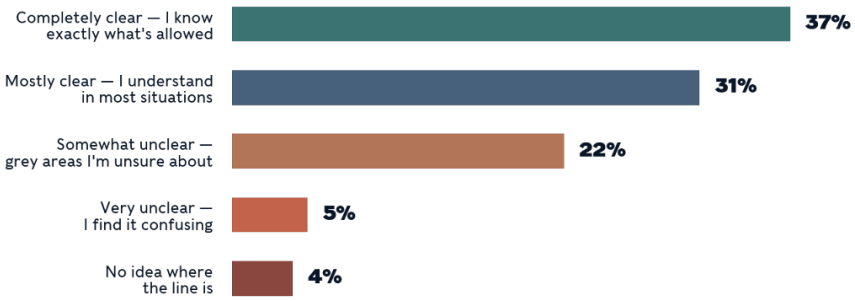

One focus group participant described her attempt to establish what was permitted on a particular assignment. She emailed her department and was told to ask the academic integrity lead. She found the module was run by the academic integrity lead. She asked him, and was told to go and ask the academic integrity lead. She gave up and did not use AI at all, staying up all night completing the assignment manually.

A postgraduate student described a more common version of the same failure:

> "The lecturer will say use your brain, do everything yourself. One tutor will be like, I encourage you to use AI as a guide. Another lecturer will say don't use AI, do your research. They're not communicating with each other, and it's making our brains go crazy. And then at the end they say we need to sign an AI declaration saying we didn't use AI – and I'm like, but I did use it. It's not like we don't use our brains. Sometimes you just need a guide." – Postgraduate taught, business and marketing

Multiple responses confirmed the pattern:

> "The uni has a tiered system based on the rules and regulations. Most of the time it is a grey area though – not clear, and some information contains things that contradict each other! We must also submit an AI policy declaration where we explain what we've used it for." – Postgraduate taught, social sciences

> "My course says no AI at all, but the university itself says some parts of AI are fine. And the problem is AI is so integrated in everything now that it's impossible to avoid." – Undergraduate, computing

A biomedical science student described a particularly revealing failure – her class received a brief, a rubric, and a template for a single assignment, and "each one of those things said a different thing – it was like three different assignments." AI could not resolve this confusion, she noted, because the lecturers themselves did not stick to their own brief. The problem was upstream of any tool.

The survey captures the grey zone that policy has created and failed to address:

> "For one assignment, I used AI to reorganise part of my argument and make the writing clearer. The ideas and analysis were my own – but I relied on it to improve structure and flow. I wasn't sure whether that fell within acceptable editing support. The guidance didn't explain how much help with structure is appropriate." – Undergraduate, social sciences

> "Since I am an international student, I found it harder to deliver concepts in a language I'm not really used to – so to be sure my ideas were strongly supported as I intended, I used AI to help me do that." – Postgraduate taught, social sciences

The deeper problem with "your own work" framing – which appears in most institutional AI policies – is that students interpret it in at least five different ways, none of them necessarily wrong.

Some take it to mean the ideas must originate with them. Some take it to mean the writing must be theirs. Some apply it to the final product but not the process that preceded it. Policy that uses it as its operative

standard is not providing guidance. It is outsourcing the ethical decision to the individual student while maintaining plausible institutional deniability.

The distinction between structural scaffolding and content generation is treated as important and morally significant by a large proportion of students across both research exercises, articulated with notable consistency:

> "From a humanities perspective, for me it's about whether the ideas are still coming from me – that's the important thing. I would never ask it 'what should I write about?' It would be 'here are my ideas, can you provide a structure for this, or some extra reading?'"

Whether or not universities agree this is a valid distinction, it is the distinction students are actually operating – and policy that fails to engage with it will continue to produce the compliance chaos this data documents.

AI declarations are a specific problem. Across the focus groups, declaration forms emerged not as a transparency mechanism but as a trust-destroying exercise. One postgraduate student was blunt about what happens when a declaration is required but guidance is absent: "They said we need to sign an AI declaration saying we didn't use AI – and I'm like, bro, I use this, but it seems you guys want us to lie to you, because that's not how we work."

A languages student described avoiding AI partly because "as we have to declare use of AI, this can affect how lecturers mark your assessment" – the declaration creating a perceived penalty for honesty rather than an incentive for transparency. A nursing student used the phrase "accidental academic misconduct" to describe her fear. The combined effect is damaging – students who use AI legitimately either lie on declarations or avoid AI entirely to avoid the perceived risk of declaring. Students who use AI most heavily are presumably the least likely to declare honestly. The declarations are catching the wrong people.

Finding 7 – Policy incoherence is a distributional justice problem that consistently punishes the most conscientious students – and the cost is emotional as well as academic

The dysfunction in AI policy is not equally distributed across the student population. Its costs fall hardest on the students who are most trying to comply, and almost not at all on the students who are not.

One focus group respondent described not using AI in their dissertation because they could not get a clear answer about whether AI transcription was permitted for something integral to their argument. Their supervisor was not sure either, so they omitted it. Their work was described as publishable quality. They did not receive a distinction because the transcription was missing – a direct academic penalty for conscientious caution in the face of a governance failure that was not theirs.

A nursing student described complete avoidance:

> "Mainly it's the fear of not really knowing how to use it effectively without accidentally cheating. And there's no guidance from the university on what you can use it for and what you can't. Also, I'm afraid of it changing the way my brain works – losing the ability to find things out myself and write things myself." – Undergraduate, nursing

Another response made the same point more starkly – "I did not use AI at any point as I fear accidental academic misconduct and I fear losing the ability to write!" A third described checking every AI interaction against her own ethical values rather than any guidance the university had provided, because no usable guidance existed.

These two fears – compliance anxiety and developmental anxiety – are different problems with different implications. The compliance anxiety is primarily about ambiguous policy. The developmental anxiety is something more serious – a worry about what repeated AI use is doing to the student's own capacity to think, reason, and write. The survey confirms this is not a minority concern – 59 per cent of survey respondents worry AI could reduce their critical thinking. That is the majority view.

But the focus groups revealed a dimension the survey does not capture – the emotional weight of the decision itself. For some students, it is closer to moral injury – a sense that using AI, even legitimately, compromises something about who they are as learners.

One English literature student uses AI to help with proofreading because she struggles with concise writing. She describes hating it, feeling guilty every time, and that while "it is easier in the moment," she feels "awful" afterwards. This is not a student gaming the system. This is someone caught between a real need for support and a deep sense that accepting it compromises her identity as a learner.

A paramedic science student identified the damaging social dimension: "You do feel like you are putting your all in and others don't show any interest in lectures yet score really well in assignments and actively state their AI usage when not within earshot of lecturers." The feeling is not just about fairness in marks – it is about the devaluation of effort and commitment as virtues.

An architecture student captures a moment that stayed with her. Working in a group, another student suggested asking AI for a list of rooms in a house. She replied: "Why don't you just use your brain?" She says: "This happened a year ago and I still think about it now. I don't want to get to a point where I don't think to use my own brain first." That a trivial moment left this deep an impression speaks to how existential the stakes feel.

The student averaging 80 per cent in their course faces neither anxiety. They describe deliberately maximising AI use on modules they do not like, checking only for hallucinations, and moving on. The de facto norm is being set by the most risk-tolerant students and the most permissive lecturers. The policy burden falls entirely on the conscientious.

That compounds the educational injustice identified in Finding 1, because the students most likely to limit their AI use out of concern are also most likely to be learning, and most likely to be penalised by a reward structure that cannot distinguish between production and understanding.

Finding 8 – Women are far less likely to use AI for assessment, and many carry anxiety about AI disadvantage without using it themselves

Two findings in the survey rarely appear in policy conversation, and they deserve dedicated sector attention.

The gender gap on AI usage is the largest demographic difference in the entire dataset – larger than any difference by disability, socioeconomic background, level of study, or year of study. It is a gap of more than 20 percentage points that does not disappear when controlling for other demographics. Whether it reflects different risk tolerances, different relationships to academic rules, different prior familiarity with AI tools, or some combination – the survey cannot say. But a gap this size warrants its own investigation.

Forty-six per cent of students worry that not using AI puts (or would put) them at a competitive disadvantage. For those that don't use AI, this is AI anxiety without AI adoption – a group experiencing the psychological costs of living in a world where AI is everywhere while having made a considered choice to remain outside it. Among non-users who still carry this anxiety, 53 per cent are disabled and 74 per cent are women.

These demographic patterns overlap and are not reducible to a single simple story. Disabled students, as Finding 9 documents, are using AI at high rates for cognitive support – yet disabled students are also overrepresented among anxious non-users. Women are disproportionately non-users – yet some women who are disabled may be using AI for accessibility needs while carrying anxiety about whether that use is permitted. The survey cannot fully disentangle these intersecting patterns, but they are worth noting because any policy response that addresses one group in isolation risks missing how these experiences compound.

These students are not failing to engage. They have made a principled decision.

> "I have not used AI in any of my work despite believing it could help improve my grades. It goes against my morals and would impede my ability to grow and think critically over taking shortcuts." – Undergraduate, humanities
>
> "I have, and will never, use AI for assessments. The whole point of a degree is to develop critical thinking skills. I realise I take quite a strong stance – but the ability to think through a problem is a dying skill, exacerbated by overreliance on generative AI." – Undergraduate, humanities
>
> "Its biggest challenge is against AI. So we need to keep on proving our skills can be used without it – and use it to improve our work, not do it for us." – Undergraduate, creative arts

These students carry the cost of their choice as anxiety. The implicit model in most institutional AI guidance treats non-use as the safe, compliant option – the absence of a problem rather than the presence of an

experience. These data suggest non-use is, for a substantial minority, an active decision made under conditions of competitive pressure that generates real distress. Universities may need to both validate the choice and address the conditions creating the anxiety.

Finding 9 – Disabled students are using AI to meet needs that formal adjustments are not, in their experience, addressing

In the focus groups, students with dyslexia, ADHD, and related processing differences described AI as the most effective cognitive support tool they had encountered – and in most cases more effective than any formal adjustment their university had provided.

> "I have all these great ideas, and my brain is such a mess at times that I speak to ChatGPT and it helps me organise my thoughts. That's where it comes in." – Postgraduate research, art and design

A student who used Copilot to structure a 6,000-word dissertation described her process – she put the learning outcomes in first, asked Copilot to build a framework, then wrote within it:

> "This really helped me as I've got dyslexia and it gave me motivation to start. Now I use AI I actually understand my work more and would be able to recall it better than without using AI. I think I find it easier to recall due to AI rephrasing things, which makes things stick in my brain better."

This last point deserves attention. The student is not just describing cognitive scaffolding – she is describing AI as improving her retention, a learning outcome claim rather than a process claim. If AI rephrasing helps a dyslexic student's understanding stick, that is a different kind of use from production assistance, and it requires different policy treatment.

> "Because I have ADHD, this really helps as a starting point as I really struggle getting projects started." – Postgraduate taught, education

> "When it came to organising my writing and ideas for portfolio submission – it was very useful for my ADHD." – Postgraduate taught, education

The gap between what the formal adjustment system provides and what these students need is visible in both datasets. One participant described a two-week extension through their Learning Support Plan as the primary formal support available – but noted that extensions apply to submission deadlines, not to the cognitive demands of the work itself.

> "It does mean that I feel like my assessments are behind everyone else's in terms of getting feedback, as well as trying to work on last-semester's work while trying to absorb new info from new lectures."

Extra time was designed for timed, paper-based examinations. It provides almost no support for the cognitive organisation demands of producing coursework, managing competing deadlines, or structuring extended written arguments.

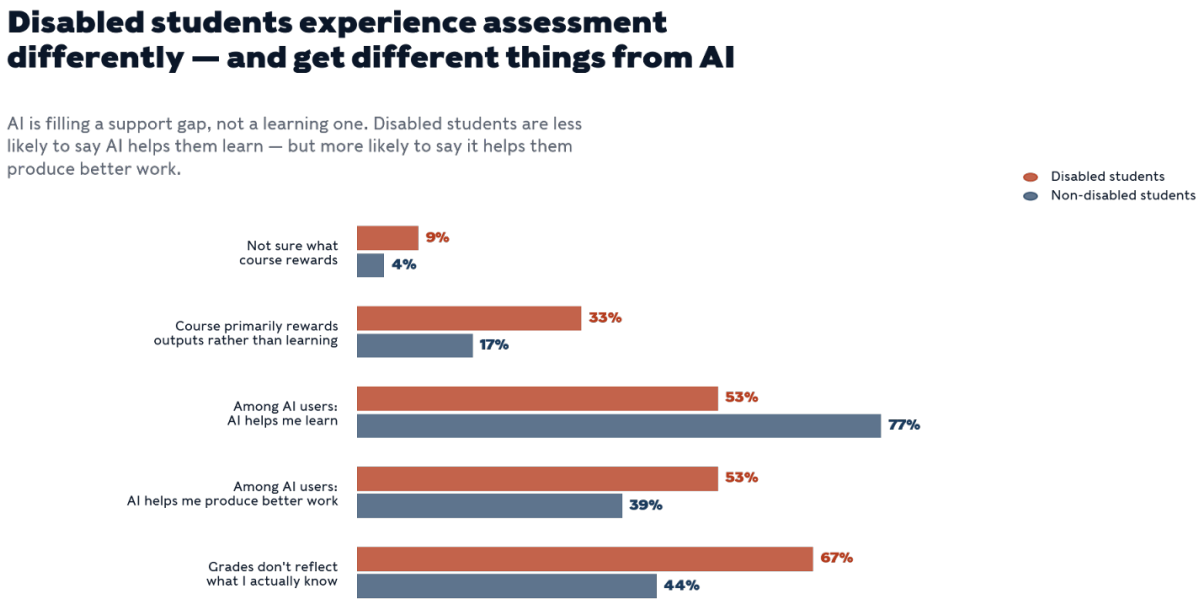

The survey strengthens the picture. Disabled students appear as a distinct constituency throughout the dataset – 67 per cent feel their grades do not reflect their knowledge against 44 per cent of non-disabled students, 33 per cent perceive their course as primarily rewarding outputs over thinking against 17 per cent of non-disabled students, and they are overrepresented in both the anxious producers and disillusioned drafters clusters.

If AI policy restricts use in ways that do not include explicit provision for disability-related cognitive support, providers are in the position of removing the most effective support tool many of their disabled students have ever accessed, in the name of academic integrity, in a context where the formal adjustment system has demonstrably failed to provide an equivalent. Disability services, academic policy teams, and legal advisors need to address this together.

Finding 10 – Feedback timing is breaking the developmental function of assessment – and driving students towards AI

Assessment serves two functions that are easily conflated. The first is summative – it produces a grade that certifies a level of attainment. The second is developmental – it tells students how to think better, not just how to perform better, and it consolidates learning that the process of producing the work may not have fully embedded.

The focus groups identified a specific structural failure that the survey data confirms – feedback routinely arrives too late to serve its developmental function. One student said plainly: "We usually get feedback from the first assessment after we've started the second, so there is no chance to develop on the feedback." Another described receiving feedback "at the end of the module, which meant we could not do anything with it, really." A third identified a subtler version: "markers reduce the amount of feedback you get if you get a first, resulting in it being less clear what you could do to improve."

When feedback cannot be used before the next assessment, the assessment sequence is functionally summative regardless of what the module handbook says. And a functionally summative assessment environment is precisely the one that incentivises AI use for production rather than learning – because the only thing the system asks for is the output.

The survey data identifies feedback quality as one of the strongest correlates of outcomes in the dataset. Students who feel feedback helps them think better – rather than just perform better – show higher career confidence, a stronger sense of community (measured throughout this report as agreement with "I feel part of a community of students and staff"), greater marking fairness perception, and lower anxiety.

The feedback question is not marginal. It appears to be one of the primary mechanisms through which institutions either close or widen the gap between assessment and learning.

In the focus groups, the feedback students describe receiving frequently fails even the quality test, let alone the timing test:

> "Sadly, in many cases, tutors use AI to give feedback, and sometimes non-AI generated feedback is absolutely useless and does not provide any scope for improvement." – Postgraduate taught, social sciences
>
> "The feedback box is blank. I receive a grade. Due to not being able to question academic judgment we don't receive the specific feedback needed to improve – this makes me feel like I've wasted time and money." – Undergraduate, social sciences
>
> "Teaching staff in my department are knowledgeable – they will know the answers to our questions. But as students, we often get stuck without knowing why, and that results in not being able to find answers for the most critical bits in time for assessments." – Undergraduate, sciences

The specific value of feedback emerges most clearly in the accounts of where real learning happens:

> "The assessment is almost the proving ground for me. I learn in classes, I test what I learnt when doing my assessments, but the true learning is after getting back the feedback." – Undergraduate, finance

This student's account of learning as a three-stage process – instruction, testing, then feedback-driven consolidation – makes visible what is lost when feedback fails. If students are producing work with AI assistance and receiving AI-generated feedback, with no verification moment in between, the developmental function of assessment has been lost at both ends simultaneously. That is a more serious failure than either problem alone – and it is happening now.

Students also identified a reciprocity problem that most institutional discussions have not confronted. If students are told not to use AI for producing work, but staff are using AI to produce feedback, the moral authority of the restriction collapses: "There's so much focus on students, but not necessarily on staff. They may be telling you to not use AI for writing or researching – staff shouldn't be using AI for marking in that case. I think there needs to be that kind of reciprocation."

The survey data connects feedback directly to AI adoption. Students working in environments where marking criteria are clear are less likely to find AI useful as a sensemaking tool. Where criteria are unclear, the majority find AI helpful. When the rubric does not explain what "good" looks like, AI fills the interpretive gap. Institutional lack of clarity around expectations does not just make students feel uncertain – it appears to drive AI adoption directly.

When what is rewarded isn't clear, AI fills the interpretive gap

Per cent of AI users agreeing that AI helps them learn effectively, by how clear they find the criteria used for marking.
Among AI users only.

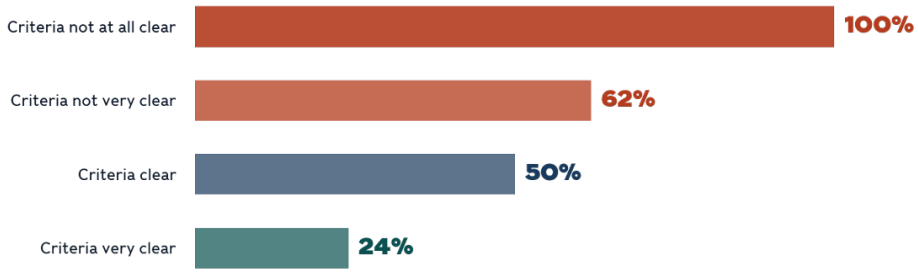

Finding 11 – AI is compensating for gaps in institutional provision that students can see but that remain unaddressed

This finding does not emerge from any single question. It emerges from reading across the entire dataset and asking what proportion of AI use described by students would disappear if universities fixed everything within their control.

The answer is – a substantial proportion. Students are not just using AI to shortcut learning. A large share of the AI use described in this research is filling gaps in what students describe the institution as not providing.

Inadequate search tools and library databases. Students described university library search as actively unhelpful. One said she could not tell the library service "you've misunderstood me – give me something more specific." A PhD student wrote that AI "works much better as a search engine than a lot of scientific databases – Google Scholar uses data tracking to become more biased over time and shows you what it thinks you want rather than just showing what you search for." Students are not bypassing institutional tools out of laziness. The tools do not work well enough for the work students are trying to do.

Unclear or contradictory assignment briefs. A biomedical science student described a brief, rubric, and template that all said different things. A paramedic science student described having multiple people – including AI – explain what was required "and for the life of me I could not understand it." When briefs are unclear, AI becomes the first interpreter – and sometimes the only one that responds at a useful speed.

Inaccessible teaching. One student wrote in the chat: "Sometimes lecturers can't even speak properly within lectures so we would either have to use AI to understand the topic or YouTube." This is AI compensating for a teaching quality issue, not a student motivation issue.

Insufficient practical preparation. A children's nursing student described using AI to generate clinical scenarios for practice – calculating PEWS scores and dosage calculations – because lecture time was being spent on activities she described as unhelpful. She was going on her first placement the following week. The AI was filling a gap in institutional preparation for high-stakes clinical practice. "Having that practice at home is very helpful," she said. "I'm less likely to make those critical errors on placement – and I think that reduces a lot more anxiety than going on a hospital ward and you're like, oh wait, I don't know what I'm doing, because we've spent more time singing in class."

Slow or unavailable academic support. AI provides instant support where emailing a member of academic staff can take days or weeks. One student described tutors with an open-door policy that was hard to access in practice: "Sometimes it's really hard to find them anyway. And their advice is very academic, and it's like, I don't understand what you're saying."

Assumed prerequisites that were never taught. Students on multiple courses described being expected to have knowledge or skills the course had never provided – statistical literacy, programming foundations, academic writing conventions. AI fills the gap between what the course assumes and what the student arrived with.

If you took away AI tomorrow, every one of these problems would still exist. The library databases would still be hard to use. The briefs would still be unclear. The feedback would still arrive late. The practical preparation would still be thin. AI has not created these problems – it has made them visible. Every institutional leader reading this data should ask not "how do we stop students using AI for this?" but "why is our provision so poor that AI is the best option available?"

Finding 12 – Time poverty is a structural driver of AI adoption that students are told is a moral or commitment failing

Time poverty is not a contextual variable in this research – it is strongly associated with AI adoption, strategic disengagement from teaching, and selective reading. A humanities student was submitting assessed work for a module with no teaching because of strikes, while simultaneously holding down a job:

> "I hadn't done any of the reading because – why would I have? I wasn't going to my lectures. I was also working to pay my rent." – Undergraduate, humanities

A healthcare student described a rational calculation about lecture attendance that most institutional strategies refuse to acknowledge:

> "Nothing's going to retain in your head if you're stressed and you've got other thoughts, so why be there when nothing's going to support you? I've not turned up to lessons that aren't engaging – I'll catch up on my own time, and I feel better. Mental health is a massive thing." – Undergraduate, nursing

A humanities student's advice to incoming students was surgically practical:

> "Only read the books you're interested in, or the short ones, because if you have three modules, you're reading three books a week for however long. You're only going to be tested on a maximum of three of them. Most people just don't have the time." – Undergraduate, humanities

The structural driver of AI for research purposes does not appear to be primarily intellectual laziness – it is partly the failure of university library systems to support the work students are trying to do. Several participants independently identified a worsening of the internet search environment as a force pushing students towards AI regardless of academic guidance.

Honesty requires a caveat. Not every student described as time-poor is experiencing structural time poverty. Some focus group participants described strategic disengagement – choosing not to attend lectures because they were running in SU elections, or selectively skipping sessions that were not engaging enough. AI serves the time-poor and the strategically selective alike, but for different reasons. A response built on "students are cutting corners" leads to surveillance and stricter policy. A response built on "students are managing competing demands with the best tools available" leads to redesign – of workload, of curriculum relevance, of what a realistic engagement model looks like for the students who are actually in the room. Both populations exist. The second is larger.

One additional equity dimension deserves attention. The cost of AI tools is not trivial. One postgraduate student described being unable to afford the paid version of any AI tool – "I'm a student, I don't have the money to be paying £30 or something to AI every month" – and instead combining outputs from multiple free tools to get usable results. Another focus group participant noted that some platforms offer student discounts.

If AI is expected or encouraged, the quality gap between paid and free tools is a socioeconomic access issue. Students who can afford paid tools get better outputs, fewer hallucinations, and more sophisticated reasoning. Students on free tiers get worse outputs and spend more time correcting errors. The AI advantage is becoming a financial advantage – the kind of access gap that higher education is supposed to close, not widen.

Finding 13 – Peer learning is a reliable route to genuine understanding and a largely untapped resource

The most consistent feature of every moment of real learning described in the focus groups is that other people were involved – not always as teachers or assessors, but more often as fellow students, audiences, or interlocutors who asked a question that required a real answer.

> "The best way I could describe it is community learning. Because we work on such long team projects and very closely together across modules, you're constantly picking things up from people on the art courses, writing courses, programming courses, design courses. You could just go over and say, 'oh, that looks really cool, how did you make that?' and they're more than happy to explain it. It just builds friendship and communication and relationships." – Undergraduate, creative technology

Trial and error at 52 per cent outranks official handbooks at 37 per cent as a navigation source. And the gap is sharpest along existing structural fault lines.

Where students learn how to succeed — by whether they're first generation

Select up to three. First-generation students rely more on peers and less on official guidance and personal tutors. The knowledge gap tracks existing structural inequality.

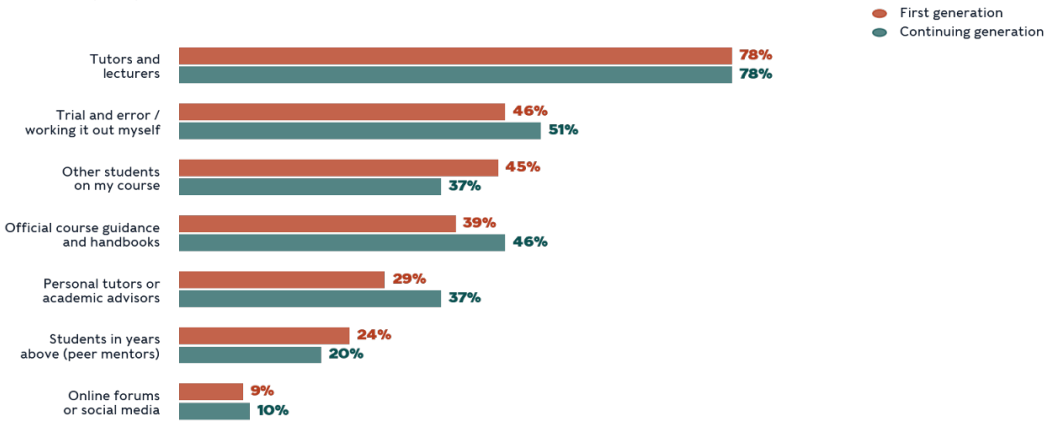

First-generation students are feeling their way – trial and error ahead of official handbooks. Less efficient, riskier, more anxiety-producing. The knowledge gap tracks existing structural inequality.

> "I am a mature student and there are a lot of assumptions made that I should know things already – but I don't, because I am an international student and my undergraduate was vastly different." – Postgraduate taught, social sciences

> "No one at my university talks about pathways into academia. So much of that is driven by my own independent research and I feel totally unaware and unsupported." – Undergraduate, computing

> "Find someone in a year or two above you. They have the best advice. Sometimes it is just 'use ChatGPT,' but other times it's 'here's how I learned how to do this, it's a lot easier this way.' And their advice is very academic – I don't understand what you're saying, going to be honest. But they are still good to go to. Definitely find another student – they are the best people." – Undergraduate, physics

One of the clearest data points in the focus groups emerged when students were asked whether explaining something to another student felt intimidating or empowering. Every single participant – without hesitation, without exception – said empowering. This unanimity is striking in a dataset where almost nothing else produces consensus. It suggests that the act of explaining is not just a way of testing understanding but a way of producing it – and that students experience this as a positive rather than an anxious process when it happens among peers rather than under assessment conditions.

This finding sits in instructive contrast to how little UK universities formalise or resource this kind of activity. The facilitators described visiting German universities where approximately a quarter of undergraduates are paid to facilitate learning for earlier cohorts – and where students preparing for mini-vivas in corridors and groups appeared relaxed about oral assessment because it was normalised. The UK has no equivalent model at scale.

Qualitative questions in the survey capture what belonging looks like when peer learning is present:

> "Because we talk a lot – we have WhatsApp groups, online meetings – and I have made some great friends. [The university] is honestly like a big family. There is always someone that will help you out." – Undergraduate, business

And what it looks like when it is absent:

> "Students just don't really engage with each other and there's a professional boundary between lecturers and students that prevents a true community. I rarely see people in the department now – our shared workspace was taken away and we are all dispersed." – Postgraduate taught, sciences

The unofficial curriculum has an equality dimension. The informal knowledge circulating between students – which lecturers to approach, how rubrics work in practice, which research tools the university does not advertise, and now which AI workflows optimise for which assessment formats – flows through social networks that universities do not control and many students cannot access.

The AI dimension of this informal knowledge is worth dwelling on. A postgraduate art and design student described realising that you could feed the module handbook and a diary of your work progress into ChatGPT and produce a good portfolio submission. She shared this technique – "only to my close friends, because it wasn't discussed what you could and couldn't do." The knowledge circulated through trusted relationships, not official channels. It was never formally prohibited, but it was never formally acknowledged either. Students who had the right social connections learned the hack. Students who didn't were left assembling their portfolios the hard way. This is precisely how structural advantage compounds – and AI has added a new layer to it.

Students who are well-connected, physically present, and socially confident have access to all of this. Students who are isolated – commuters, those with caring responsibilities, those without a strong peer group – are relying on official guidance alone. And multiple students were explicit that the official guidance is not sufficient.

Peer learning is not an enrichment opportunity. It is the single most reliable generator of the kind of understanding described anywhere in this research. Resourcing it, formalising it, and where possible paying students to do it would simultaneously support understanding, build community, and distribute the unofficial curriculum more equitably. It would also create the normalised conditions for oral and explanatory assessment that Finding 4 identifies as the most promising direction for reform.

Finding 14 – Career-confident students are distinguished less by assessment format than by whether their course rewards thinking

A common argument holds that traditional assessment is disconnected from professional life, and that assessment should more closely resemble authentic work tasks. If this were true, students who feel career-confident should report greater confidence in work-based formats and less in traditional academic ones. They do not.

Career confidence is self-rated on a 0–10 scale in the survey, in response to the statement "my current activities at university are preparing me for my future career." Splitting into low (1–4), mid (5–7), and high (8–10) groups produces consistent differences – but not where the "make it more like work" argument predicts. The strongest correlates with career confidence are not about assessment format at all. They are about whether the course is intellectually honest – whether feedback develops thinking, whether stated values match actual rewards, and whether assessment tests understanding rather than production. Feedback quality is the single strongest correlate ($r = +0.40$), and the absence of a gap between what the course says it values and what it rewards is close behind.

Among low-career-confidence students, 40 per cent strongly agree there is a gap between what their course says it values and what it actually rewards. Among high-career-confidence students, that falls to 4 per cent.

Only 5 per cent of the low-career-confidence group feel their assessment primarily rewards thinking. Nearly a third of the high group do.

Career confidence tracks intellectual honesty — not assessment format

Students split by self-rated career confidence (0–10 scale). The strongest correlates are whether feedback develops thinking, whether assessment rewards understanding, and whether stated values match actual rewards.

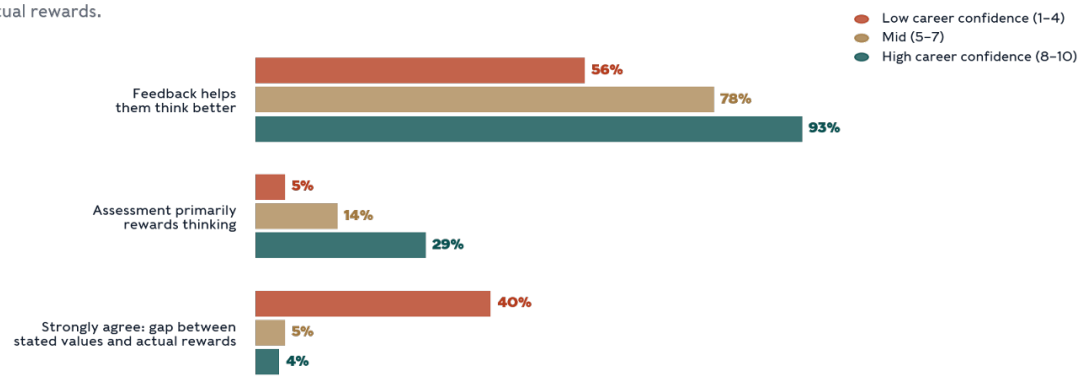

The free-text comments confirm this. The student at the bottom of the scale does not complain about insufficient workplace simulation. They complain about insufficient thinking:

> "we learn how to implement, rather than think and if we graduated with only the stuff we learnt inside the curriculum, nothing good would come out of us" – Undergraduate, computing

Oral examinations – among the formats most dismissed as archaic – have the strongest positive correlation with career confidence. Placement and work-based outputs – the format most explicitly designed to be "career-relevant" – are among the weakest. This does not mean format is irrelevant. The appendix confirms that practical, applied assessment is the single most commonly requested redesign. But career confidence appears to track intellectual honesty, developmental feedback, and clarity of expectations more strongly than it tracks superficially vocational formats.

Essays are the only format with a negative correlation, possibly reflecting recognition that essays as currently designed can reward polished production over demonstrated understanding. The formats career-confident students trust most are those requiring understanding to be demonstrated on the spot.

High-career-confidence students do not describe what they value in the language of workplace simulation:

> "I believe I am being taught the skills and work ethic needed for industry but that nothing can truly compare to hands of work experience. Even with our live briefs that make us work with an industry client in real time, the university structure is completely different. So while I believe the course is doing a great job I feel I will still have an immense amount to learn from experience." – Undergraduate, creative arts

This student has the most "authentic" assessment imaginable and still frames what the course gives them as skills and work ethic, not workplace replication. The course does not need to look like a job. It needs to build the capacity to do one.

Career confidence also correlates negatively with every measure of AI dependency. Among low-career-confidence students, 60 per cent worry that not using AI puts them at a disadvantage. Among the high group, 40 per cent. The most career-confident students frame this not as a compliance problem but as the defining professional challenge:

> "Its biggest challenge is against AI. So we need to keep on proving our skills can be used without it and use it to improve our work not do it for us" – Undergraduate, creative arts

Career-confident students are also more likely to agree that official guidance tells them everything they need to know and less likely to feel there are unwritten rules. At the other end:

> "No one at my university talks about pathways into academia. So much of that is driven by my own independent research and I feel totally unaware and unsupported." – Undergraduate, engineering

The problem is not format. It is clarity. These students are not describing a curriculum that fails to resemble the workplace. They are describing one that has not told them how it connects to anything beyond itself.

The direction of reform the data supports is not "make assessment look more like work." It is: make assessment intellectually honest. Reward thinking. Give feedback that develops reasoning. Close the gap between stated values and actual rewards. Be transparent about expectations.

These are the conditions under which students feel their education is preparing them for a career – and they have little to do with whether the assessment format resembles a professional task.

Finding 15 – Students who feel they belong reach for AI less – because their course already provides what AI substitutes for

If peer learning is the primary mechanism through which belonging is built (Finding 13), and belonging shapes whether students experience their course as rewarding thinking or rewarding production, then the absence of resourced peer learning is not a missed enrichment opportunity. It is a structural driver of the conditions that make AI adoption feel necessary.

The survey data supports that claim. Belonging – measured here as agreement with "I feel part of a community of students and staff" – is one of the most powerful correlates of learning experience in the dataset. The strongest single relationship is between belonging and criteria clarity (r = +0.54). One plausible interpretation is that students who feel part of a community find the rules of the game clear because they share interpretive knowledge with each other. But it is also possible that courses with clearer expectations produce both stronger belonging and greater clarity – the direction of explanation is not settled by the data alone, and the two accounts are not mutually exclusive. What the correlation does establish is that belonging and clarity travel together. The next strongest relationship is with peer knowledge sharing – students who feel they belong share what they know.

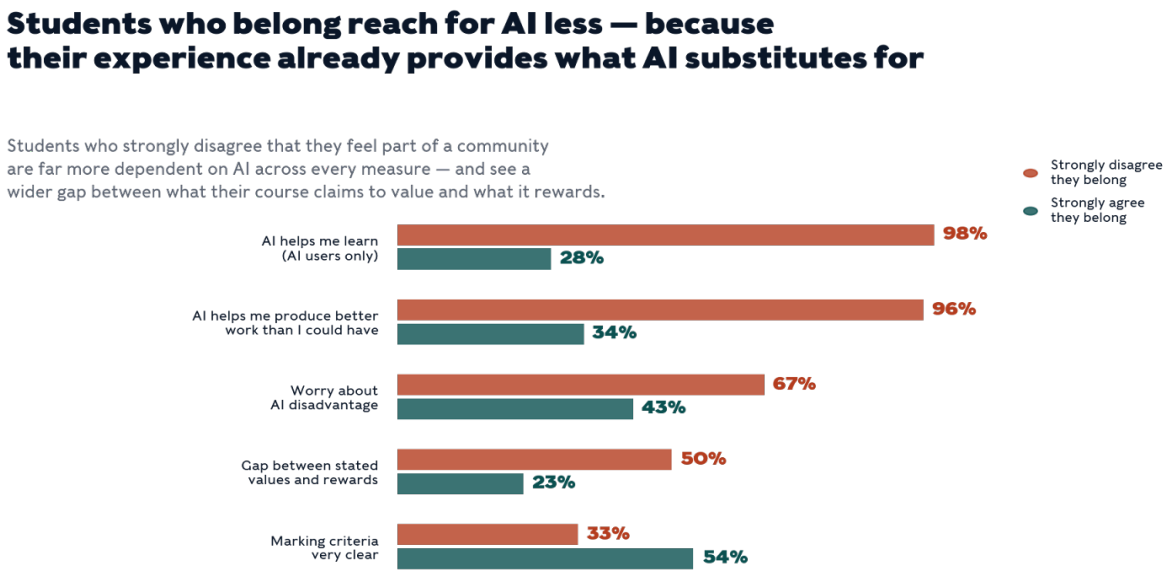

Students with stronger belonging report a different relationship with their course, their assessment, and AI. On assessment and learning, they are more likely to report that their assessments test thinking not just production, that feedback helps them think better, that their course develops thinking not just assignment completion, and that they are learning what they need rather than just producing what they need to produce. They are less likely to perceive a gap between what their course says it values and what it rewards.

On AI, the relationships run in the same direction. Students with stronger belonging are less likely to say AI helps them learn effectively, and less likely to worry that not using AI puts them at a disadvantage. Among students who strongly disagree on belonging, only 33 per cent do not use AI for assessments. Among those who strongly agree, 51 per cent do not.

The absence of belonging is visible in the data as isolation and strategic withdrawal:

> "I go to these exhibitions of student works and I don't see how any of them would get a professional opportunity after graduation. I'd like to think that I'm going at it alone to secure my future." – Undergraduate, computing

> "I feel like with my course it's not so much about learning – but having the time to produce your work to the standard they expect, whether you understand it or not." – Undergraduate, architecture

Belonging does not just make students feel better – it is associated with markedly different reports about their assessment, their learning, and their need for AI. Students who belong experience their course as rewarding thinking. Students who do not belong experience it as a production line – and reach for AI accordingly.

Recommendations

For universities

Introduce verification moments into assessment. A structured conversation, a brief oral component, a process reflection conducted in person – attached to a proportion of assessed work. This does not require a full viva for every student. It requires a policy decision that verification conversations are a normal part of assessment rather than a misconduct investigation tool. If students know submitted work may be followed by a conversation about it, the incentive structure changes. Design these as developmental dialogues, not performances – the focus groups consistently distinguish between conversations they want and presentations they dread.

Prepare students for live accountability. Students should not be put into oral and live formats cold. Low-stakes practice, explicit preparation, and multiple uncapped formative attempts make live accountability developmental rather than discriminatory. Accessibility concerns are real and must be designed through explicitly – alternative means of live demonstration, not avoidance of the direction of travel. Normalising peer explanation as a routine activity – as some European systems do – would build confidence for more formal accountability moments.

Reframe assessment design around a single question – how would a student prove to us that they understood this, rather than simply produced something plausible? For many current formats, the honest answer is that they cannot – and that was true before AI. Rubrics designed to reduce marking inconsistency have become tools for gaming assessment. The accountability moment is the corrective.

Replace AI declaration forms with stage-specific guidance at module level – what is and is not permitted at brief interpretation, research, structuring, drafting, and checking stages. Generic statements about "your own work" are not guidance. They are liability management. Current declarations penalise honesty and reward concealment; they should be replaced with guidance that students can actually use.

Treat library search as part of the AI problem. Students are adopting AI as a research discovery tool partly because institutional search systems do not work reliably. This is directly addressable.

Audit what AI is compensating for. If students are using AI because briefs are incomprehensible, teaching is inaccessible, practical preparation is thin, or feedback arrives too late to be useful – those are signals about institutional provision, not student misconduct. Every heavy AI use case is diagnostic information about what the institution is not providing.

Formally address how disabled students use AI for cognitive support. Disabled students are sourcing support through AI that they describe their institutions as not providing. Blanket AI restrictions risk removing the most effective support tool many have ever accessed. Disability services, academic policy teams, and legal advisors should jointly map current use and build explicit provision into AI guidance before further policy tightening.

Audit feedback practices. Is feedback arriving in time to inform the next piece of work? Is it providing actionable guidance on how to think rather than how to perform? Are criteria explicit enough to reduce reliance on AI as a sensemaking tool? Is quality sufficient to support the consolidation of learning assessment is supposed to produce? And if staff are using AI to generate feedback while telling students not to use AI to generate submissions, the reciprocity problem needs to be addressed explicitly.

Investigate the gender gap and AI anxiety among non-users as standalone policy concerns. The gap of more than 20 percentage points in AI non-use by gender and the 46 per cent carrying AI anxiety without AI adoption warrant examination on their own terms. Current policy has no answers because it has not asked the questions.

Resource peer-assisted learning as a core part of provision, not enrichment. Formalising, resourcing, and where possible paying students to facilitate learning for earlier cohorts creates one of the most reliably effective learning experiences available. It also distributes the unofficial curriculum equitably rather than leaving it as a privilege of social capital.

Treat workload, curriculum relevance, and equitable access to AI tools as AI policy questions. Students are not primarily using AI because they are dishonest. They are using it because they are time-poor, because some of what they are asked to do does not connect to the course they enrolled for, and because AI is the most efficient tool available for work they cannot make themselves care about. Where AI is expected

or encouraged, the quality gap between paid and free tools is a socioeconomic access issue that institutional licences could address.

For course teams

Produce module-level AI guidance specific to the discipline, the assessment type, and the stage of the process. The variation across creative arts, computing, healthcare, humanities, and quantitative subjects is too wide for any uniform approach. Course teams should be expected to design guidance that is actually operable – with institutional support – rather than leaving students to reconcile contradictory signals from different members of the same teaching team.

Every module should be able to answer clearly why it belongs on the programme. Students who see no connection between a module and the discipline they came to study will treat it as an obstacle to clear rather than learning to engage with. AI makes clearing it very easy.

Map the prerequisite assumptions in each module and make them explicit. The research identifies students using AI to fill knowledge gaps that courses assume have already been closed. If a module assumes statistical literacy, programming foundations, or academic writing conventions, that should be stated at enrolment – not discovered by the student in week three. Where gaps are predictable, bridging resources should be provided as part of the module, not left to the student to source through AI.

Coordinate AI guidance within programme teams. The most common student complaint is not that policy is too strict or too permissive – it is that different tutors on the same programme say different things. A ten-minute conversation at the start of the academic year between the members of a teaching team, producing a shared position that students can rely on, would address the single most reported source of confusion in the data.

Design at least one assessment moment per module where the student must account for their work in person – even briefly. This does not need to be a formal viva. It could be a five-minute conversation during a tutorial, a structured peer-explanation exercise, or a reflective walkthrough of the student's process. The point is not surveillance. It is that the presence of such a moment changes how students use AI throughout the module.

Create structured opportunities for students to develop and share their ethical thinking around AI. The most thoughtful students in this research have built personal doctrines about AI use that are more sophisticated than most institutional policies. That intellectual work is currently invisible and unsupported. Seminars, peer discussions, or case study workshops where students surface and compare their positions would do more than another round of policy revision – and would treat students as partners in working this out rather than subjects to be governed.

For the sector

Accept that detection is a trailing indicator. Detection technology will always be behind the capabilities of the tools students are using, and a large proportion of academically problematic AI use would not register on any current detection system. The response is assessment design that makes detection irrelevant – because the assessment itself reveals whether the student understands what they submitted.

Find the Struggling Disengaged. They are not using AI, not thriving, and not making noise. They are invisible in current data and policy – including, most likely, in this research. They will not seek help. Better use of early-warning data, proactive personal tutor outreach, and peer mentoring that addresses navigational knowledge gaps would reach them – but only if the sector first acknowledges they exist.

Conclusion

There is a deeper question running beneath all of this that the sector tends to sidestep. When students were asked whether AI was simply a skill like any other – part of what makes a graduate employable – one pushed back directly: "what are we at university for if not for special learning?"

That is not a nostalgic or anti-technology position. It is a question about what higher education is actually for – whether it exists to develop the capacity to think, or to certify the possession of transferable technical skills.

Assessment design that cannot distinguish between production and understanding has, in practice, already answered that question. The students in this research mostly have not accepted the answer.

One student arrived at university wanting to learn, received mediocre grades for understanding, concluded they were forced to use AI, and now gets good grades while acknowledging they cannot be certain they have learned anything. In their view, their university trained them to stop learning in order to succeed. No AI policy will fix that. Assessment design might.

But the final thing this research reveals is a resource. The students who participated in these focus groups and this survey are not confused consumers of a broken system. They are thoughtful, ethically engaged people working through real complexity with very little institutional support. They can describe with precision when real learning happens and when it does not. They have built personal ethical positions around AI that are often more considered than the policies their universities have produced.

They have designed alternatives to current assessment that would test understanding, and their proposals – while diverse – converge on a set of principles that any assessment designer would recognise as sound. Every single one of them, when asked whether explaining something to a peer felt intimidating or empowering, said empowering.

The sector would do well to treat them as partners in working out what learning means now – because on the evidence of this research, they have been thinking about it harder, and more honestly, than most of the institutions around them.

Research was conducted in February and March 2026. The national survey was distributed via Wonkhe SUs subscriber SUs, comprising responses from 1,055 students across 52 HE providers. Survey results are weighted for gender and level of study. Quantitative analysis – Spearman rank correlations, k-means cluster analysis (k=4), principal component analysis, and thematic analysis of free-text responses across the open questions. All percentages refer to the valid response base for each question. Significance thresholds – * $p<0.05$ ** $p<0.01$ *** $p<0.001$.

Appendix: Analysis of free-text survey questions

The two open-ended questions in the main body of the survey are

1. "If you could redesign one of your assessments to genuinely test whether you'd learned something valuable, what would you change?"

2. "Can you describe a situation where you were unsure whether your use of AI was acceptable?"

Question 1 – Assessment redesign

Despite being optional, this question produced a detailed picture of what students think is wrong with current assessment and what they'd prefer instead. The additional responses strengthen several themes visible in the original data and introduce some new ones, including vocational and placement-based learning, presentation anxiety, and poster assignments.

More practical, applied, and real-world assessment

This was the dominant theme across the combined data. Students across disciplines – computing, psychology, healthcare, and others – called for assessments that test whether they can apply what they've learned rather than simply reproduce it in writing. Specific suggestions included working software applications, case studies with simulated clients, real-world scenario tasks, and practical simulations.

Several students drew an explicit contrast between writing about doing something and actually doing it. One called for less essay focus and more active learning tasks – "doing rather than writing about doing." A computing student argued for building actual software as a better test of understanding. Healthcare students wanted more simulation-based assessment, including practising clinical skills on peers.

The additional responses strengthened the vocational angle in particular. Students on placement-based courses argued that reflective essays should focus on what they've practically done on placement rather than on how well they "analyse" – a critique of assessment that privileges academic framing over practical experience on courses that are explicitly vocational.

Problems with group assessment fairness

Group work remained a pain point. The combined data includes students who want to abolish group assessment entirely – one argued it doesn't teach teamwork as lecturers claim, because in a workplace there are consequences for people who don't pull their weight whereas at university there aren't – alongside others who see value in collaboration but want individual contributions assessed separately. Proposed solutions included attributing sections to individuals, having tutors observe and assess each person's contribution, and grading group projects individually.

A counter-current did emerge. Some students actually wanted more group work, in psychology especially, where one proposed collaborative case study presentations as preparation for multidisciplinary team working. The disagreement isn't about group work itself – it's about whether group assessment can fairly distinguish individual learning.

Frustration with exams as a test of memory

The critique of traditional timed exams as rewarding memorisation over understanding remained strong. Students described the "cram, write, forget" cycle and argued for open-book formats that better reflect real-world professional practice. But the expanded dataset reveals this isn't one-directional.

Some students wanted more exams in place of essays or coursework, and a few specifically wanted MCQ formats replaced with short-answer questions or vice versa. The underlying desire is the same – students want the format to actually match what it claims to test – but there's no consensus on which direction to move.

Presentation anxiety and format preferences

The additional responses sharpened this theme considerably. Multiple students expressed discomfort with live presentations, with one asking to replace them with portfolio-based alternatives and another noting that live formats don't reflect ability when someone struggles with public speaking.

One student mentioned social anxiety and depression as factors that make presentation performance inconsistent day to day. Conversely, other students wanted more oral and viva-style elements, seeing them

as better tests of understanding. There's a design problem here – formats that test communication skills can inadvertently penalise anxiety, introversion, or neurodivergence.

Poster assignments are disliked

A new theme from the additional data. Several students criticised poster assignments, arguing they prioritise design over knowledge and that tight word counts make it impossible to demonstrate understanding. One suggested converting poster tasks to presentations or essays.

Desire for more formative assessment and feedback

Students called for pop quizzes, mock exams, in-class feedback exercises, and reflective work that runs throughout the course rather than only appearing at summative moments. One student specifically wanted formative quizzes "to help us learn rather than it being graded." Another proposed more informal, low-stakes opportunities such as short group presentations, role plays, and facilitated reflection. The shared demand is for more frequent, lower-stakes chances to test and consolidate understanding before the high-stakes submission.

Clearer criteria and better guidance

Several students wanted less ambiguity in what's expected – clearer criteria, better alignment between word counts and the depth of analysis expected, and less vague assessment briefs. One noted that feedback frequently criticises lack of detail when word counts are tight, a misalignment between what's asked for and what's rewarded. This theme is less about assessment format and more about the clarity of the contract between students and the institution.

More assessment choice

A smaller but interesting thread. Students wanted the ability to choose their assessment format, suggesting that offering options like presentation, group project, or essay would let people demonstrate learning in the way they're strongest. Different students demonstrate learning differently, and a single-format approach disadvantages some.

Question 2 – Uncertainty around AI use

This was the more heavily answered question by a considerable margin. The combined dataset presents a complex and sometimes contradictory picture. The student body divides into several distinct groups, and the boundaries between them are often blurry. Categories below overlap – many students fall into more than one.

"I don't use AI" – principled and practical non-users

Some respondents either stated they don't use AI or expressed principled opposition to it. Within this group, motivations vary.

Ethical and environmental objectors. These students frame AI use in strong moral terms. Environmental impact featured prominently – multiple students cited water usage and energy consumption. Others described AI as fundamentally incompatible with learning, arguing it prevents students from developing their own knowledge, reasoning, and academic skills.

Self-interest non-users. Some avoid AI not on moral grounds but because they believe they produce better work without it, or because early experiences with inaccurate outputs eroded their trust. One student noted that initial institutional messaging warned against unreliable AI content, so they remain wary even though policies have since shifted.

Passive non-users. A handful simply said they don't use AI without elaboration, or indicated they don't know how.

The strength of anti-AI feeling in this group is worth registering. Several used emphatic language and some appeared to view the question itself as an opportunity to register opposition. This cohort likely shapes the social dynamics around AI disclosure – students who do use AI may be reluctant to discuss it openly if they perceive strong peer disapproval.

Planning, structuring, and brainstorming

The most common use case described across the dataset. Students use AI to plan essays, generate bullet-point outlines, brainstorm initial ideas, suggest headings, and structure arguments. Many feel their use is broadly acceptable, but a subset expressed uncertainty about where planning ends and content generation begins. One student said the line was unclear until they asked multiple lecturers their opinions. Another noted that AI suggested paragraph starters that they then used – sitting right on the boundary between structure and content.

This is perhaps the hardest category for institutions to regulate, because the same tool used the same way can look like either legitimate scaffolding or unacceptable assistance depending on how much the student then does with the output.

Studying, comprehension, and revision

Students described using AI to explain difficult concepts in simpler language, summarise academic papers, check their understanding of topics, generate practice questions, and create revision materials. This group generally framed AI as a learning aid rather than a production tool. One student explicitly said they then worried that using AI to help them understand lectures "undermined what my lecturers were teaching" – an anxiety suggesting that even using AI to learn might be seen negatively.

One case from the data is instructive. A student used an AI study platform to generate flashcards and revision notes, then was flagged for academic misconduct when those notes were used in an open-book exam. This shows how AI use in preparation can bleed into AI use in submission in ways that current policies may not adequately address.

Language, grammar, and writing support

This theme was amplified by the additional responses. Students use AI (including Grammarly, Co-pilot, and ChatGPT) for sentence-level editing, grammar checking, rephrasing, and improving clarity. Three sub-groups are identifiable.

International and EAL students. These students describe AI as essential for expressing ideas they already have in a language that isn't their first.

Students with disabilities. Dyslexic students and those with DSA-funded software face a particular dilemma. One dyslexic student described Grammarly's rephrase function as essential but worried it would impact their grade. Another noted that when AI guidelines were introduced, there was no guidance for DSA-granted software – a gap in policy that hasn't been closed. A student using Co-pilot for understanding complex sociological ideas reported that some lecturers say it's fine while others say it isn't, leaving them unsure what is actually acceptable.

General writing improvement. Other students use AI to smooth sentence structure or improve clarity, often with visible anxiety about whether this crosses a line. One described sending a draft to AI for improvement suggestions, liking what it produced, but then being afraid to keep it in their submission.

Students expressing genuine uncertainty about the boundaries

This is the group the question was most directly designed to surface. Their responses reveal several specific areas of confusion.

What counts as "AI"? Multiple students were unsure whether Grammarly, Canva's image features, PowerPoint Designer, or Google's AI overview count as "AI" in the context of academic integrity policies. One asked whether Grammarly is "AI use or just a good text editor?" Another noted that PowerPoint Designer isn't explicit about being an AI tool. These aren't edge cases – they expose a definitional problem as AI becomes embedded in everyday software.

Where's the line between editing and generating? Students who used AI to restructure arguments, improve sentence flow, or reorganise essay sections were often unsure whether they'd crossed from acceptable editing into unacceptable generation. One described it as "always a bit of a muddy ground."

Inconsistent guidance from lecturers. Students reported getting different messages from different lecturers about what's acceptable. This is damaging because students who do seek guidance – the most conscientious ones – end up more confused, not less.

Changing institutional messages. At least one student was confused by policies that initially prohibited AI but later allowed it, leaving them unsure what the current expectation is.

Referencing and source-finding

A specific use case that emerged clearly. Students use AI to find, format, or verify references. One student used ChatGPT to find references and checked the links were valid but noted it "felt a bit like cheating" because they hadn't gone out of their way to find the sources themselves. Others used AI for Harvard referencing, bibliography formatting, or locating sources they couldn't find through conventional search. The academic skill of source identification is part of what assessment measures, but the mechanical formatting of references is less clearly so – and current policies don't distinguish between them.

Image and design generation

A theme from the additional data primarily. Students described using AI to generate images for presentations, create logos for projects, and produce design elements via Canva. Uncertainty here focused on whether AI-generated visual content falls under the same rules as AI-generated text, with one student confused about why AI-generated photos were unacceptable when designing a resource. Another wasn't sure whether Canva's image generation "really counts as AI" – another instance of the definitional problem.

Confident, bounded users

A smaller group who described specific AI use cases and felt clearly comfortable with them. These students typically used AI for one or two defined tasks – planning, comprehension checking, brainstorming – and felt confident this was within acceptable bounds. Some explicitly referenced institutional guidelines as being clear. This group is a useful contrast. Where guidelines are clear and use is bounded, students don't experience much anxiety.

Students who pushed boundaries

A small number of students described more extensive use. One admitted to getting AI to write an assessment and then just changing a few words. Another described using AI to synthesise academic articles and paraphrase specific sentences – sitting very close to the line of AI-assisted writing. A third described asking AI for an example piece of work on their chosen topic, reading it, then deleting it – but debating how acceptable it was. These cases are rare but matter. They confirm that the full range of use exists.

Cross-cutting observations

The "what counts as AI?" problem is growing

Across both datasets, students are increasingly confused not just about how much AI use is acceptable, but about what counts as AI at all. When Grammarly, PowerPoint Designer, Canva, Google search overviews, and DSA-granted accessibility software all incorporate AI features, the distinction between "AI use" and "no AI use" is becoming unworkable. Policies that draw a line around "AI tools" without defining what falls inside that boundary will generate increasing confusion as AI becomes embedded in standard productivity software.

Accessibility needs remain a policy blind spot

Students with dyslexia, ADHD, English as a second language, and DSA-funded software reported that AI guidelines don't account for their needs. The gap between accessibility-tool AI and what might be considered misconduct is not addressed in the policies these students encounter. This affects a substantial proportion of the student body and puts them in the position of using tools they need while worrying those tools will be flagged.

Inconsistent lecturer-level interpretation matters as much as policy

Even where institutional policies exist, students report getting different messages from different lecturers. The issue is less about whether policies exist and more about whether they're communicated and applied consistently at module level.

Environmental objections are a distinct strand

The additional data strengthened the environmental dimension of anti-AI sentiment. Multiple students cited water usage, energy consumption, and environmental harm. This framing is distinct from the pedagogical objection (that AI undermines learning) and the ethical objection (that it's plagiarism), and it likely connects to wider generational concern about sustainability shaping attitudes toward technology.

Students want assessment that tests what it claims to test

The assessment redesign responses and the AI uncertainty responses point to the same underlying problem. There's a perceived gap between what courses say they value (understanding, critical thinking, application) and what assessments actually measure (output, format compliance, memory). AI exacerbates this gap because it's very good at producing polished output – making output-focused assessment increasingly unreliable as a measure of understanding. Students are, in many cases, describing this problem more precisely than the institutions trying to address it.

The social dynamics of AI use are polarised

With a substantial cohort of students vocally opposed to any AI use and another substantial group actively using it for planning, studying, and writing support, there's a real social divide. Students who use AI may feel reluctant to discuss it openly given the strength of anti-AI sentiment among their peers. This could push AI use underground – the opposite of the transparency that institutions say they want.