

Appendix: Additional analysis

The following material was developed from the same dataset as the main report and published as standalone analysis. It is collected here as supplementary evidence that extends several of the main findings.

[For articles that discuss, see the tag on the site](#)

Socioeconomic status and AI use

The competitive anxiety gap

Students eligible for means-tested funding are nine percentage points less likely to use AI on their assessed work than those who are not eligible. They feel clearer about the rules – 38 per cent of eligible students say the line between acceptable and unacceptable AI use is "completely clear" to them, compared to 28 per cent of not-eligible students. And they are far less worried about falling behind. Among students not eligible for means-tested funding, 62 per cent agree or strongly agree that not using AI puts or would put them at a competitive disadvantage. Among eligible students, it is 45 per cent.

The clarity gap runs deeper than mechanics. Eligible students who described uncertainty in the free text tended to recount specific, contained scenarios – they tried something, assessed it themselves, and moved on. Not-eligible students were more likely to describe a systemic fog – uncertainty not about a particular act but about the rules themselves, sometimes even when they had not used AI at all.

First-generation students are 11 percentage points more likely than continuing-generation students to find the AI boundary unclear. State and FE students are 15 points more likely than privately educated students to find it unclear. But bursary-eligible students reported near-zero incidence of having used AI in ways they were not sure were acceptable. Privately educated students, by contrast, were far more likely to have pushed the boundaries and far more confident about defending what they did.

The production-shortcut split

Among eligible students who do use AI, the most common use is planning, structuring, and producing. Among not-eligible students, the skew is toward using AI to understand concepts or materials – 32 per cent versus 17 per cent of eligible students. Students from lower socioeconomic backgrounds are more likely to experience AI as a production tool. Students from higher socioeconomic backgrounds are more likely to experience it as a learning tool.

The single largest differential in the class analysis is on "AI helps me learn more effectively" – 66 per cent of privately or grammar-school educated students agree, versus 34 per cent of state and FE students. That is a 32-point gap. First-generation students are twice as likely as continuing-generation students to say AI helps them produce work but does not help them learn.

Decision-making by class

Students from lower-occupation backgrounds are far more rule-driven when deciding how to use AI – 85 per cent consider the assignment guidelines, compared to 61 per cent of those from professional and managerial backgrounds. They are nearly twice as likely to factor in the risk of detection. They are six times more likely to look at what other students are doing.

They are also significantly more driven by circumstance. Among bursary-eligible AI users, 40 per cent say time pressure and deadlines are a major factor in how they use AI – compared to 17 per cent of non-eligible students.

Forty per cent cite assignment difficulty, versus 13 per cent. These describe students reaching for AI not out of laziness or ethical indifference but because the clock is running out, the brief is unclear, and the assignment is hard – conditions that track closely with the financial pressure to work alongside studying, the caring responsibilities that eat into study time, and the workload compression that follows from all of it.

Students from professional and managerial backgrounds, by contrast, are more ethics-driven – 74 per cent consider "what feels right ethically" versus 46 per cent – and more learning-driven, considering whether AI would actually help them learn at a rate 20 points higher. State and FE students are significantly more likely to worry about whether AI use would undermine their learning – 56 per cent versus 25 per cent of privately educated students. That 30-point gap cuts against any assumption that disadvantaged students are cavalier about AI. They are not cavalier. They are anxious about a different thing – not "will I get caught?" in the abstract, but "is this making me worse?"

The hidden curriculum and AI

First-generation students are 17 percentage points more likely to agree that there are things you need to know to succeed that nobody tells you. State and FE students are 29 points more likely than privately educated students to say they have learned from peers about how the system works. Privately educated students are 16 points more likely to agree that official guidance tells you everything you need to know.

This matters for AI because the rules governing AI use are, in most institutions, exactly the kind of knowledge that requires confident handling of ambiguity. If you already know how the system works – which rules are real, which are aspirational, which tutors care about what – you can push the boundaries of AI use with relative safety. If you are still working out which rules matter, caution is the rational response.

Seventy-nine per cent of privately educated students say they would feel confident explaining their AI use to a tutor, compared to 60 per cent of state and FE students – a 19-point gap. That may not be a gap in competence. It may be a gap in confidence – in the belief that you can have a conversation with an authority figure about what you have done and come out of it unscathed.

Assessment legitimacy

Students eligible for bursaries are 19 percentage points more likely to see a gap between what their course says it values and what it actually rewards. At the same time, they are 19 points more likely to agree their tutors are interested in how they arrived at conclusions, and 16 points more likely to say their assessments test thinking. That is a counterintuitive combination. These students perceive more individual care from their tutors and more structural misalignment in the assessment system at the same time. They can see what their course is trying to do – and they can see that it is not doing it.

What disadvantaged students want instead

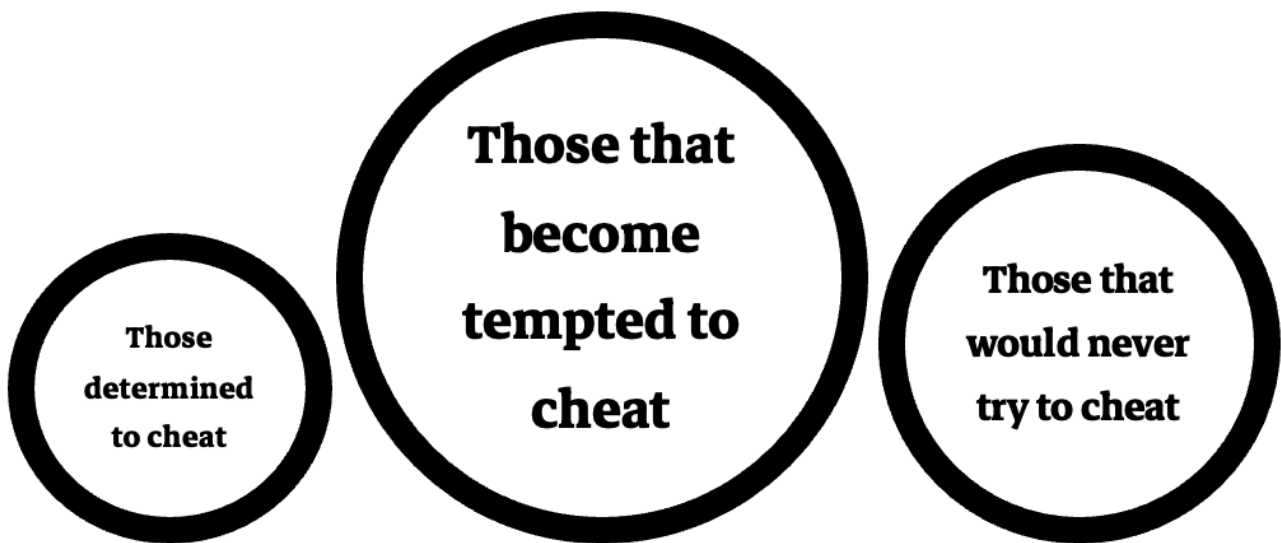
Eighteen per cent of eligible students called for practical, hands-on, applied assessment formats in the redesign question – compared to just four per cent of not-eligible students. That is the largest gap on any single theme in the redesign data. Eligible students also raised exam and memory-based concerns at a higher rate – 23 per cent versus 16 per cent – and were more likely to mention group work as something they would change.

The students most likely to experience AI as something that helps them produce but not learn are also the students most likely to ask for assessment formats where production without learning is impossible. They have identified the problem and they are describing the fix.

The three circles: who detection catches, who it misses, and who it harms

The framework

Students fall into three broad groups in their relationship to AI and academic integrity – those determined to use AI regardless of rules, those who become tempted under pressure, and those who have chosen not to use AI or to limit their use to what they consider ethically defensible. The enforcement apparatus – detection tools, declaration forms, misconduct procedures – lands on these three groups in an almost perfectly inverted relationship to the one it intends.



The determined group

Students described peers who submit AI-generated work without even checking whether it works. A software engineering student described classmates who "use the AI for assignments, to create a code, create an assignment. But most of the times, they are wrong. Some students don't even try to check if the assessment is right. They just go and submit it." The implication is that if the ones submitting broken code are sometimes detected, the ones who spend thirty seconds checking the output before submitting are invisible. A computer science student confirmed the logic from the other direction – "Realistically, you could get a passing grade in any written module. If you throw ethics out the window and get AI to write it all and just quickly correct the glaring mistakes." Detection tools built on perplexity and burstiness are trivially defeated by anyone who knows to add a persona or rephrase a few lines.

The anxious middle

Students using AI under pressure, often clumsily, often guiltily, are exactly the group most likely to be flagged. Their use is less polished, more detectable, and more likely to produce the patterns that both software and academics are trained to spot. These are students submitting work at 3am after a shift, using AI to get through a module they did not choose on a course that has not given them clear guidance. They are the ones most likely to end up in a misconduct hearing.

The competitive pressure that drags students into this group was visible in the focus groups. A postgraduate student described a peer – a diligent international student who had worked hard, avoided AI, and been awarded a lower second despite having been an exceptional school student. The peer told them they were "now forced to use AI" because without good grades the investment in their degree would be wasted. That student was dragged from the principled group into the tempted group by the grade competition itself.

The principled group

Students who have chosen not to use AI described being actively harmed by a system that cannot distinguish their work from the work of those who used AI to achieve a higher grade. A history and film student described the frustration – "I've done a piece of work without AI and then someone goes 'oh yeah I used AI and I got a 75' and I'm like – well I got in the 60s without using AI but you got a first with using AI, so why is the university not picking that up?"

She went further. Because of the unreliability of AI detection tools, she had started self-censoring her own writing – avoiding punctuation that might trigger a false positive. "I get scared to even use some grammar – so I completely don't use semicolons because I'm too scared that it's going to pick it up as AI." A principled student who has chosen to do her own work is afraid of her own punctuation. The detection system has made an honest student write worse – not because she has done anything wrong, but because the tools that purport to catch cheating cannot tell the difference between AI-generated text and a student who uses certain sentence structures.

The distributional result

The enforcement regime rewards both integrity and sophisticated dishonesty – but cannot distinguish between them in the output – while punishing financial stress, deadline pressure, and the kind of messy, guilty, half-hearted AI use that characterises students who are actually trying to do the right thing but cannot quite manage it. The result is an apparatus that catches one group and misses two.

Group work compounds the problem

Everything described above applies to individual assessment. Group work amplifies it. Students in year-long group projects described working alongside peers from different disciplines with different relationships to AI, all assessed together. A student who avoids AI because the academic regulations are "such a grey area" is graded alongside a student who has been told by their discipline to embrace it.

The burden of managing other people's AI ethics in a group project falls on the students who have thought hardest about those ethics. One student described the experience bluntly – "Oh it was bad enough when you just had someone lazy, but someone in my group just straight up pasted ChatGPT into the doc and didn't even check it. And I'm sat there like – do I snitch? Do I just fix it myself at 2am? Why is that on me? I'm not the AI police." And the anxiety runs both ways – "I've got groupmates who are terrified to even use AI for brainstorming because they think they'll get thrown off the course. So now I'm spending half my time going 'no honestly it's fine for this bit' and the other half actually doing the project."

The academic judgement problem

The academic judgement framework that protects assessment decisions from challenge was designed for a world in which a marker with a red pen could point to copied sentences. It held, just about, when Turnitin could trawl the internet for matching text. It does not hold when a marker is asked to determine whether a student used AI based on the frequency of em dashes, whether the subheadings use title case, or whether the essay has a suspicious fondness for a particular word. Extending the scope of unchallengeable "academic judgement" to cover AI-misconduct determinations puts students in an impossible position – they cannot challenge the decision because the university has deemed it a matter of academic expertise, but the "expertise" is often an algorithmic guess dressed up as professional judgement, and the student has no way to prove a negative.

Module choice, intellectual stimulation, and disengaged AI use

The stimulation gradient

Among students who said their course was "not very often" intellectually stimulating, 73.6 per cent reported using AI tools on their assessed work. Among those who found it "very often" stimulating, that figure fell to 57.3 per cent. But the more telling difference was not whether students used AI but how. Less stimulated students were more than twice as likely to use AI for drafting and writing (13.8 per cent versus 6.0 per cent), and significantly more likely to use it for planning and brainstorming (34.5 per cent versus 22.8 per cent). They were not reaching for AI to understand their subject better. They were reaching for it to get the thing done.

Compulsory modules and the absence of choice

In the focus groups and free text, a connection between compulsory modules, absence of choice, and heavier AI reliance came up repeatedly. Students described a targeted withdrawal from the parts of their degree that they had no agency in – not a general disengagement, but a specific disconnection from the reasons they enrolled. A second-year history undergraduate described a core module on historiography she found abstract and disconnected from the periods she had chosen to study – "I didn't even really understand what argument I was supposed to be making, so I described the essay question to Copilot and asked it what the response should look like. Then I basically wrote around that. I don't think I learned anything." A business management student described a compulsory statistics module – "I just needed to get through it. I don't use stats in any of my other modules – I'm never going to need to know this. So I basically let it do the working and I formatted it to look like mine."

The sector context

Module choice is contracting at exactly the moment this research shows that restricted choice is one of the drivers of disengaged AI use. A Universities UK survey published in May 2025, drawing on responses from 60 institutions, found that 46 per cent had removed optional module choices in response to financial pressure – up from 29 per cent the year before. OfS polling published earlier in 2026 found that 83 per cent of students who had noticed cost-cutting at their institution felt a gap between the experience they believed had been promised at enrolment and what they actually received. OfS sustainability reporting recorded a third consecutive year of declining surpluses and liquidity, with 43 per cent of institutions expecting a deficit for 2024–25.

A cascade, not a single symptom

When intellectual stimulation is tracked across three successive polling waves, low stimulation does not just predict different AI use. It predicts almost every negative outcome measured, across a gap of 20 to 36 percentage points on each one – whether students feel part of a learning community, whether they trust the assessment to test understanding, whether they believe staff are accessible and responsive, whether the course matches what was advertised, and whether they feel confident about their career.

Students who find their course unstimulating are more than twice as likely to agree that you can get good grades without understanding the material (55 per cent versus 23 per cent), and nearly three times as likely to say there is a gap between what their course says it values and what it actually rewards (64 per cent versus 23 per cent). Career confidence scores drop more than two full points on a ten-point scale. Life satisfaction and sense of purpose drop with them.

Disengaged AI use is one symptom of something wider. When students are stuck in content they did not choose and cannot see the point of, they do not just reach for AI – they also check out of their learning community, lose confidence in their future, stop trusting the assessment, and conclude that gaming the system is a more rational strategy than engaging with it. The structural squeeze on choice produces a cascade of disengagement, and AI usage is one way of handling it without it damaging their prospects.

Five types of moment students described as genuine learning

Across all the focus group sessions, the question "can you think of a moment where you genuinely learned something?" produced strikingly similar clusters. Five types appeared with the most frequency and force.

1. Verification-pressure accumulation

Students described later verification preparation as the most consistent trigger for genuine learning – not the exam or viva itself, but the intensive period of revision that preceded it. The pressure of having to know rather than just having read created conditions where understanding clicked in ways it had not before.

"Every single thing that I've said I've learned is because there was an exam coming up. When you're really studying for an exam, that's when I learn a lot." – Undergraduate, medicine

This type was the most widely shared across disciplines and was almost always reported as assessable – students could demonstrate what they had genuinely understood. It raises a specific question about AI – if the pressure of verification is doing pedagogical work, what happens when students offload revision to AI?

2. Practical application

Students across vocational and creative disciplines described a qualitative shift when theory became something they had to do rather than describe. The moment of application – executing a clinical skill, performing a scene, writing functional code – produced a feeling of understanding that reading and writing about the same material had not.

"I remember learning about it in theory, but actually going out and doing it practically – that's where my passion and learning strikes more." – Undergraduate, nursing

"The performance made me feel like I'd learned more than the actual coursework part of it." – Undergraduate, drama

Students in this cluster described the assessment as the site of learning, not just its demonstration – a distinction that sits in tension with much academic assessment design.

3. The seminar room

Small-group discussion – seminars, tutorials, informal debate – appeared as the most consistently named site of genuine learning, particularly for humanities and social science students. Students described something different happening in that space – ideas becoming real through dialogue, understanding sharpening through challenge.

"The seminars were where the learning happened – you actually got to debate and protect your ideas." – Undergraduate, English literature

"In seminars I feel like I'm learning a lot more because of the energy and the vibe in the room – we're bouncing our ideas off each other and it kind of feels more real." – Undergraduate, history

A recurring observation was that this learning is rarely assessed – it happens in the space that counts for nothing officially. Several students noted the frustration of that gap directly.

4. The synthesis moment

Several students described a distinctive late-stage experience – a moment near the end of a long or substantial piece of work, usually when ordering or presenting accumulated material, where disparate understanding suddenly cohered. This was not the same as completing work. It was the feeling of a semester clicking into place.

"Near the end of the diss, when I'd mostly written my points and was trying to order them – everything came together. I'm like, whoa, I've understood everything I've done." – Undergraduate, history

"Only really in that process of putting it together as a presentation did I actually start to click and understand what I'd actually been writing about for 10,000 words." – Postgraduate, history of art

This type points to something that AI-assisted writing may specifically displace – the sense-making that happens through the laborious process of assembling and ordering ideas is itself the learning, and cannot be outsourced.

5. Deep independent construction

A smaller but striking cluster of students described learning that was anchored to a moment of full intellectual ownership – when they had built something themselves, without shortcuts, and knew it completely. This was most often retrospectively identified through retention – the work they could still explain months or years later was the work they had constructed rather than assembled.

"I did the traditional research – going to the library, downloading articles, going through them. I got a distinction. It's been almost one and a half years. I can still tell you what it was all about." – Postgraduate, international student

The contrast students drew with AI-assisted work was not about quality at submission – it was about what remained afterwards. Deep construction left a trace in memory and confidence. Production for submission did not.