



The Difficulty Is the Point

On what artificial intelligence owes the students who are behind

For most of human history, knowledge was scarce. Books were rare, teachers were few, and the hard part was getting access to the material at all.

We solved that. Every lecture, every textbook, every proof is now a search away.

And attainment gaps did not close.

Because access was never the bottleneck. The bottleneck is what happens in the hours between hearing something and understanding it. That work is slow, uncomfortable and unavoidable, and no technology has ever removed it. What technology can do is decide who has to do it alone.

The wrong promise

When generative AI arrived in education, the promise was that every student would have a tutor. What most students got was an answer machine.

An answer is not an education. Reading a fluent explanation feels like understanding, and that feeling is a measurement error. The student closes the laptop confident and sits the exam unprepared. We have built, at enormous expense, a machine for manufacturing the illusion of competence.

Worse, it is not distributed evenly. The students who benefit are the ones who already know how to interrogate an answer and notice when something is wrong. The students who needed help most take the first answer and move on. A technology sold as the great leveller is quietly reproducing the divide it was meant to close.

What we believe

Rigour and accessibility are usually treated as a trade. Simplify the material and it stops being the course. Leave it as delivered and it excludes the students who cannot process a fifty-minute lecture in one pass.

We refuse that trade.

A student who is behind does not need the work done for them. They need the same work, arranged so they can actually start it. The difficulty stays. What changes is where it is spent: not on decoding a badly recorded lecture at two in the morning, but on retrieving an answer from their own memory, being wrong, and being shown precisely why.

This is not a philosophical preference. It is the best evidenced finding in the learning sciences. Producing an answer beats reading one. Being tested beats rereading. Being wrong with confidence, then corrected, produces more durable memory than being nearly right.

So we built a system that makes students produce answers instead of consuming them.

What we will not build

- We will not write a student's essay. Our system generates revision material from their own lecturer's content, and cannot produce anything submittable. That is a limitation we chose and will keep.
- We will not optimise for time in app. Streaks and points are off by default, because a neurodivergent academic told us they cause anxiety, and she was right.
- We will not charge students. The university pays, or nobody does.
- We will not train on their work.
- We will not claim more than we can evidence. Our own research paper names two features the literature contradicts, and we published it anyway. A company that cannot say what it gets wrong cannot be trusted about what it gets right.

Where this goes

Every university already owns the material its students need. It sits in lecture recordings nobody rewatches and slide decks nobody opens twice.

Our work is to turn that dormant archive into something a struggling student can actually use, without asking them to admit they are struggling first.

Making rigorous knowledge accessible.
Not easier.
