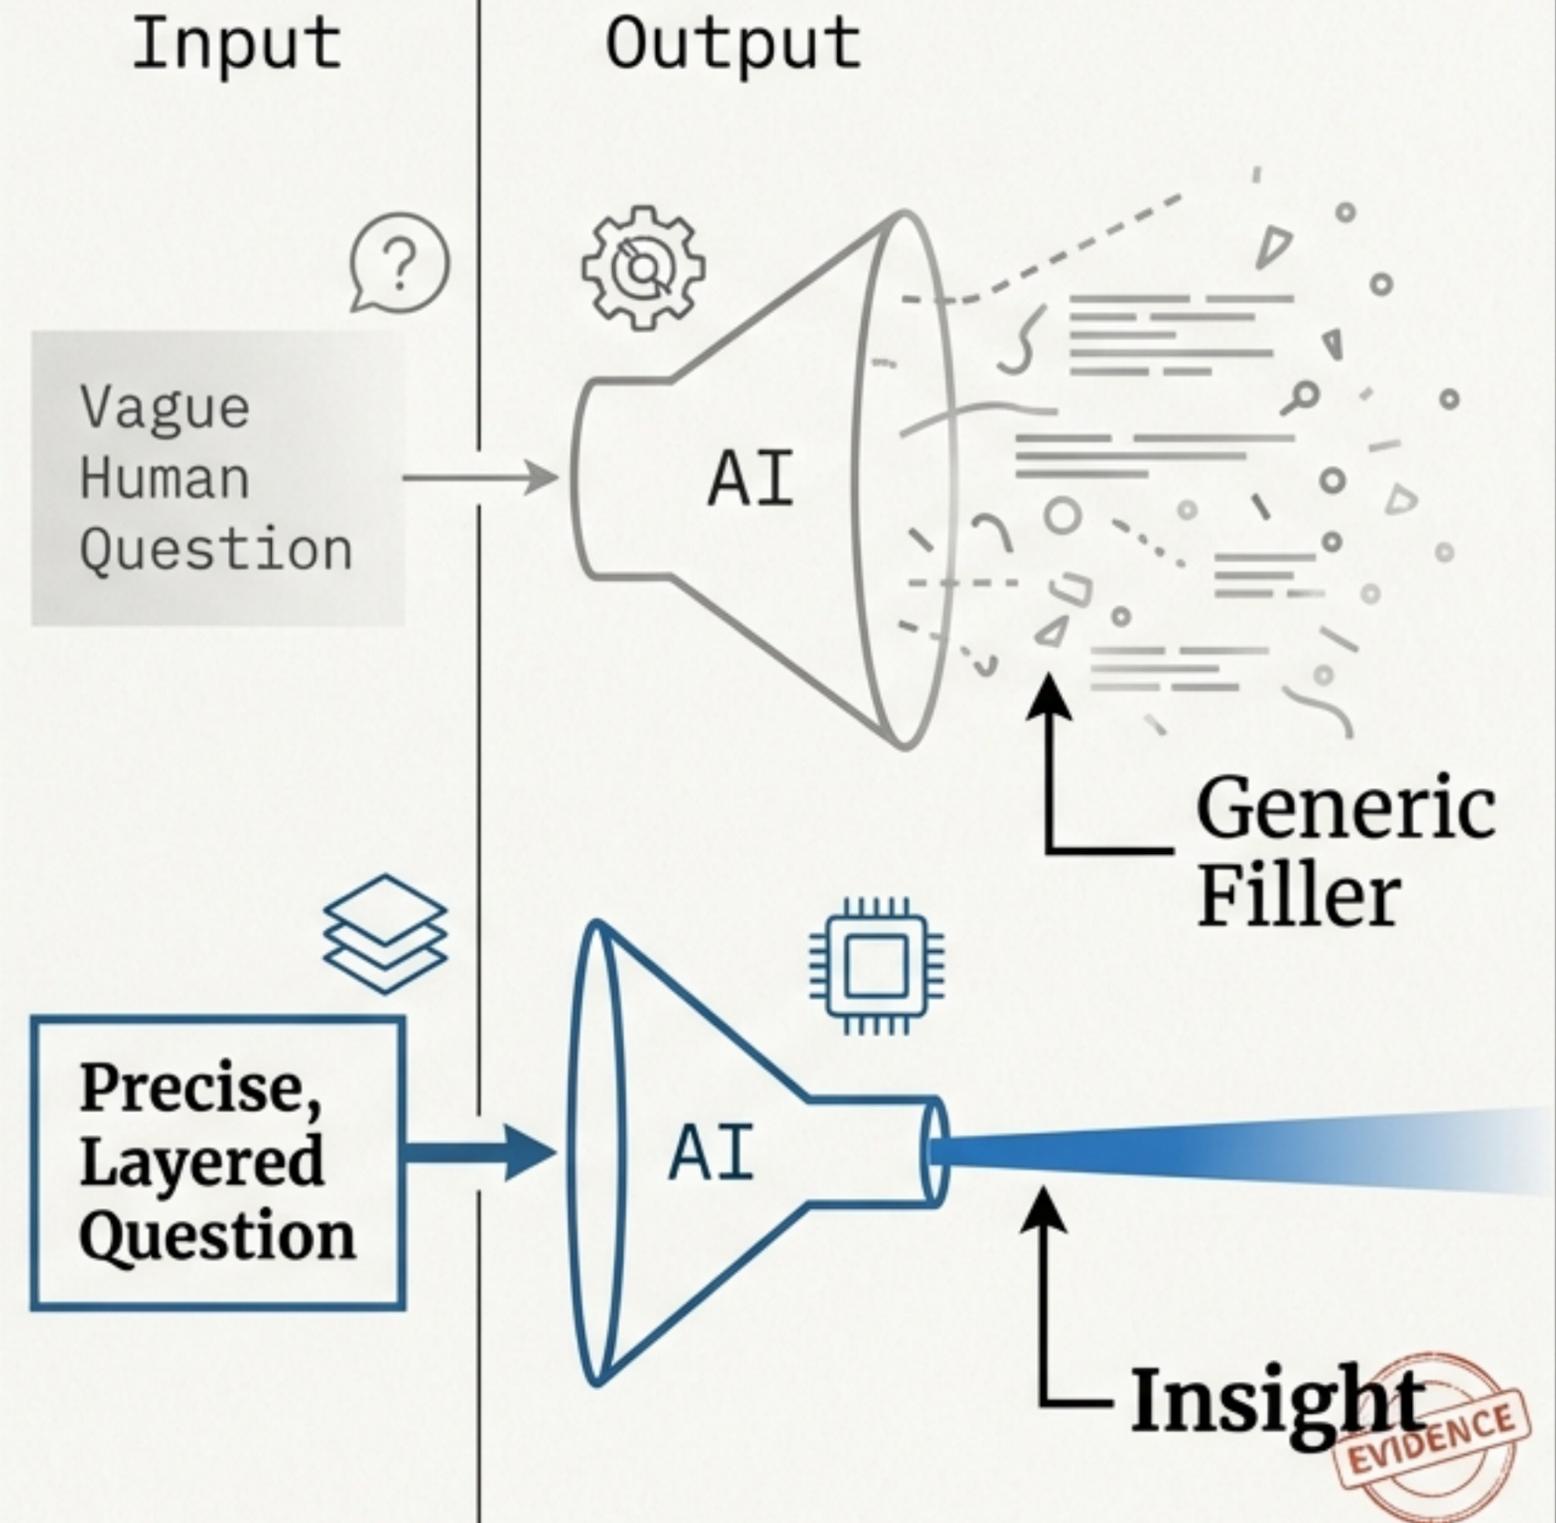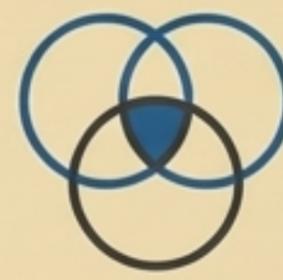# Asking Better Questions

The AI-Era Thinker's Playbook

AI Quality is Downstream of Question Quality

The student who asks better questions gets better answers from every tool, every person, and every system for the rest of their career.

This is not about prompt engineering. It is about thinking clearly enough to know what you *actually* need to know.

Input

Vague Human Question

Output

AI → Generic Filler

Precise, Layered Question

AI → Insight

EVIDENCE

The Four-Step Cognitive Gauntlet

Step 1: Predict
(The Prediction Lock)

Commit your diagnosis before touching AI.

Step 2: Evaluate
(The Question Tournament)

Rank questions by diagnostic power, not topic.

Step 3: Diverge
(The Divergence Test)

Prove independent thinking beyond generic machine output.

Step 4: Defend
(Live Defence)

Survive live peer questioning and AI counter-arguments.

NotebookLM

Step 1: The Prediction Lock

The Human Side

TIMESTAMPED

Human Hypothesis

The Rule: Before opening any AI tool, you must write down your initial diagnosis and 10 ranked questions with predicted answers.

The Value: The value of this document comes entirely from being written before AI access. It creates a genuine, unalterable record of independent thinking.

The Machine Side

No Access Yet

AI

The Outcome: Comparing this sealed prediction

The Outcome: Comparing this sealed prediction against eventual AI output reveals your cognitive blind spots—showing exactly where your thinking was strong and where it was lazy.

The Reasoning Receipt

Core Concept:

Documenting AI interactions to prove critical engagement.

The Mechanism:

For every prompt sent, the user must log the response and actively choose to Accept, Reject, or Modify it.

The Proof:

A one-sentence justification must accompany every decision. This transforms AI usage from invisible copying into a measurable, verifiable reasoning process.
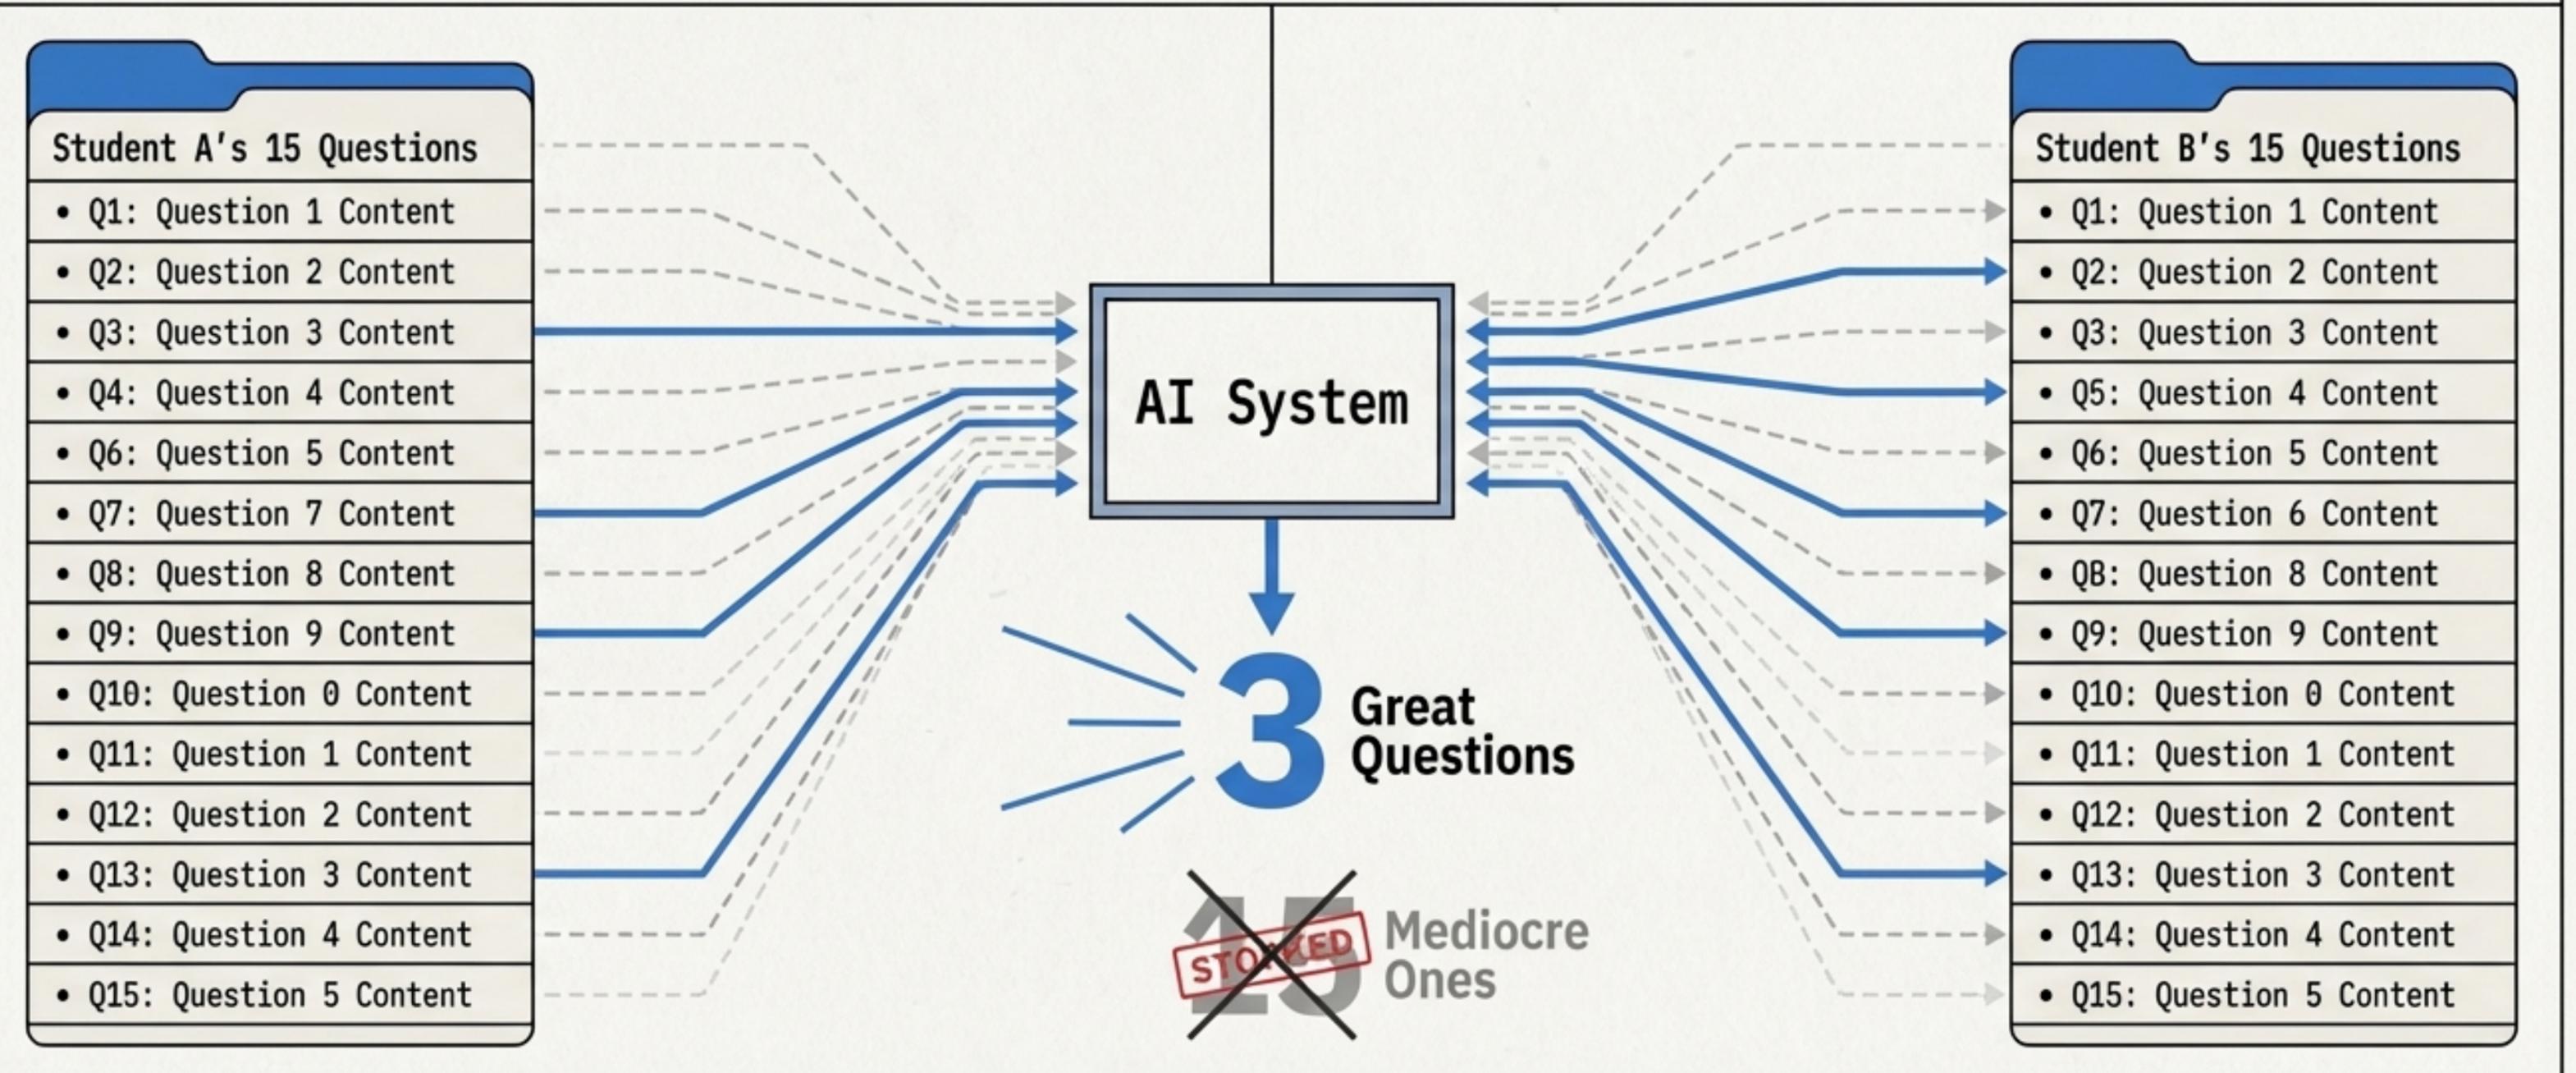
Raw AI Response

The Human Filter

Accept   Reject   Modify

One-Sentence Justification

Verified Insight

NotebookLM
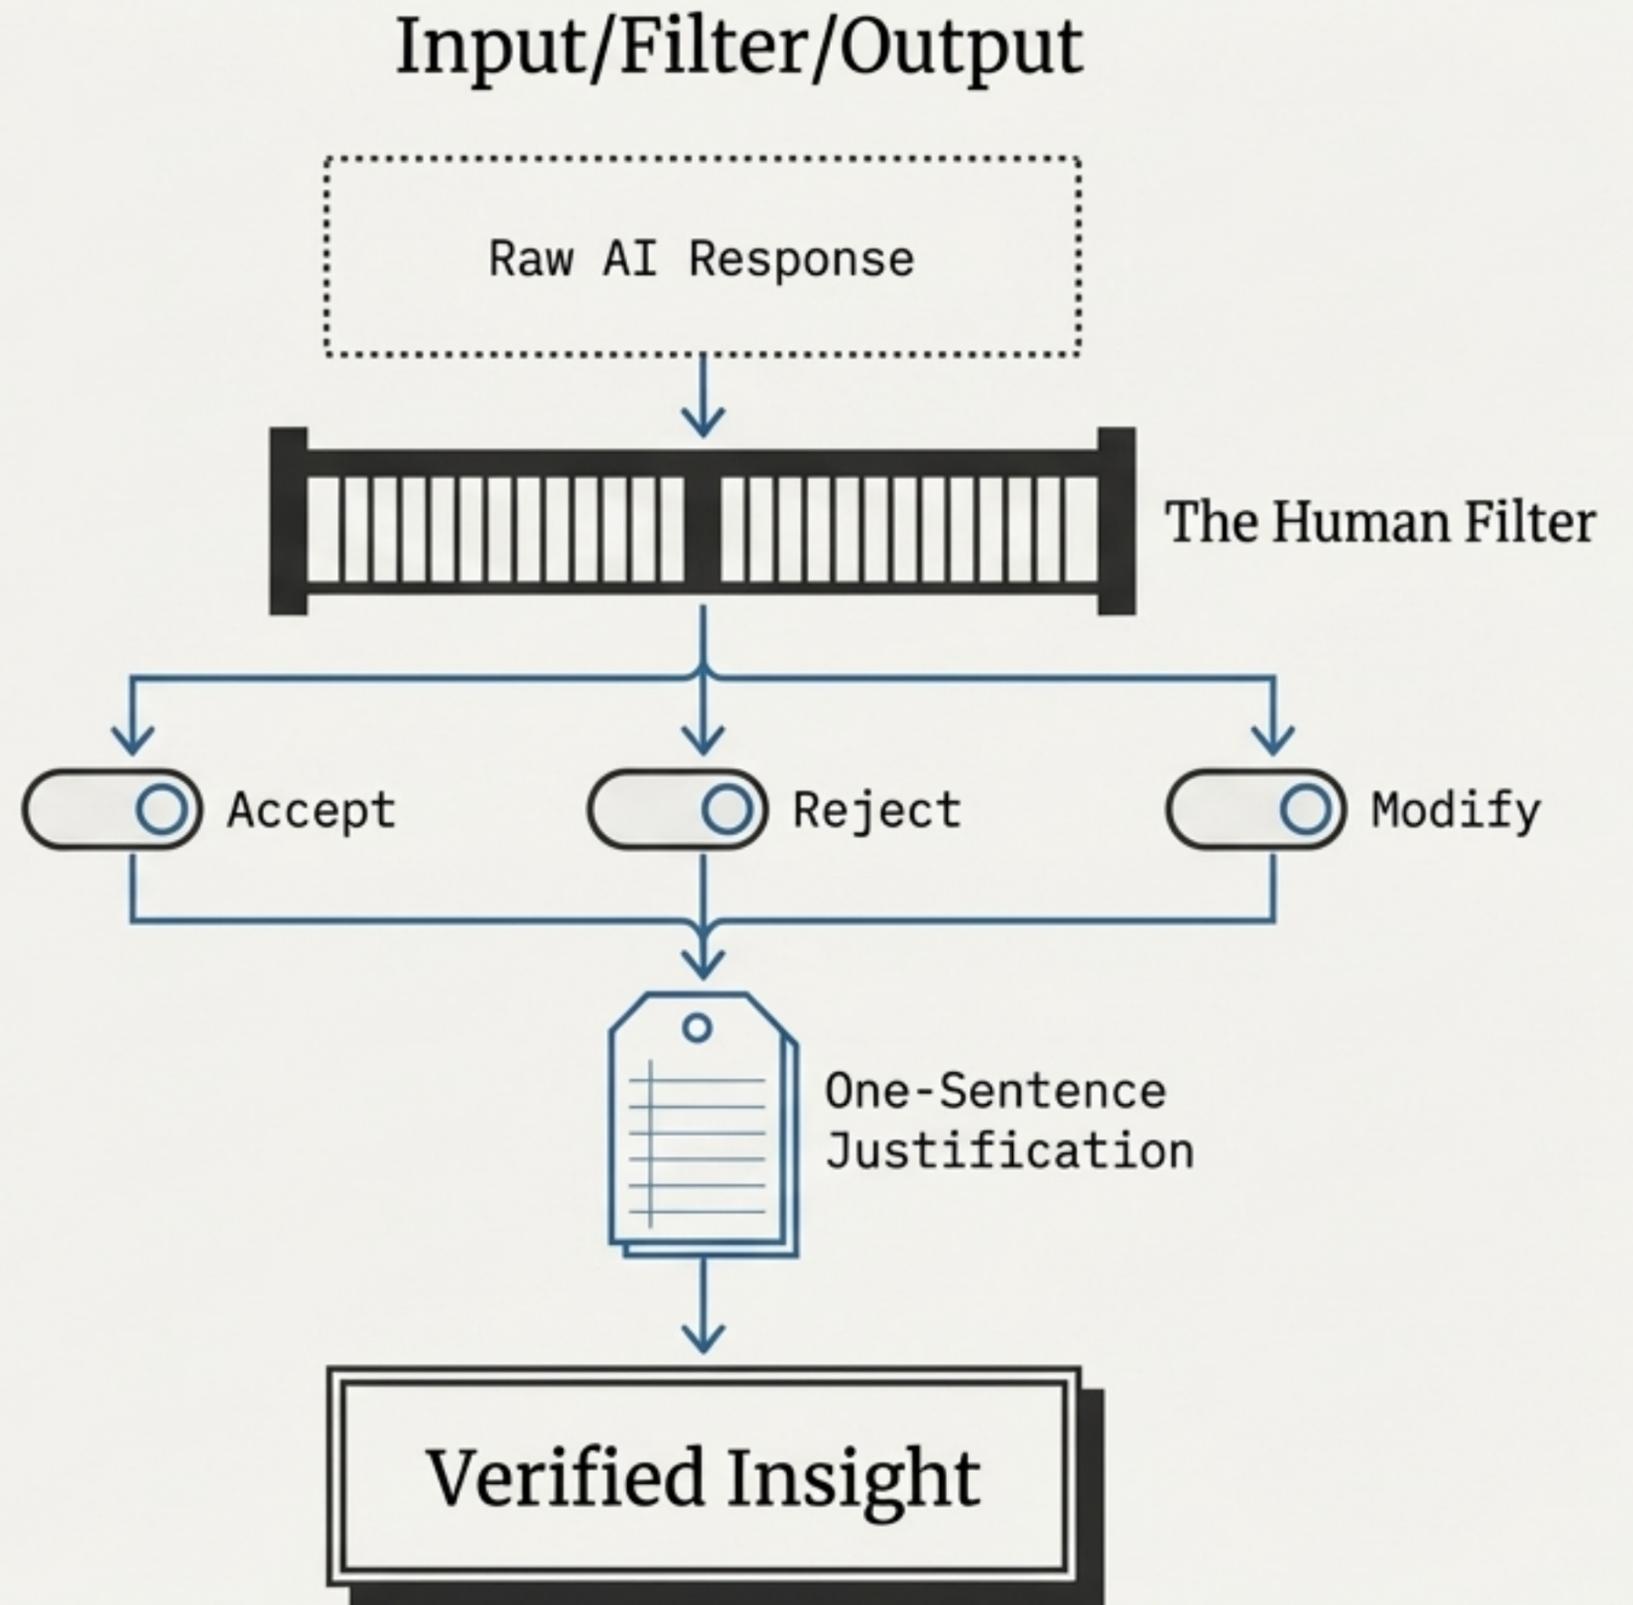
Step 2: The Question Tournament

The Challenge: Students swap 15 original questions, rank their partner's list, and feed the top 5 into AI to see which actually produce actionable answers.
The Revelation: Question quality is a measurable skill, not an innate talent. Many well-phrased questions produce useless, generic filler.
The Lesson: Evaluating a partner's questions develops analytical skills faster than self-assessment by revealing patterns invisible in your own work.

Student A's 15 Questions
- Q1: Question 1 Content
- Q2: Question 2 Content
- Q3: Question 3 Content
- Q4: Question 4 Content
- Q6: Question 5 Content
- Q7: Question 7 Content
- Q8: Question 8 Content
- Q9: Question 9 Content
- Q10: Question 0 Content
- Q11: Question 1 Content
- Q12: Question 2 Content
- Q13: Question 3 Content
- Q14: Question 4 Content
- Q15: Question 5 Content

AI System

3 Great Questions

STOPPED 15 Mediocre Ones

Student B's 15 Questions
- Q1: Question 1 Content
- Q2: Question 2 Content
- Q3: Question 3 Content
- Q5: Question 4 Content
- Q6: Question 5 Content
- Q7: Question 6 Content
- QB: Question 8 Content
- Q9: Question 9 Content
- Q10: Question 0 Content
- Q11: Question 1 Content
- Q12: Question 2 Content
- Q13: Question 3 Content
- Q14: Question 4 Content
- Q15: Question 5 Content

Generic vs. Diagnostic Questions

Generic (Ranked by Topic Importance)	Diagnostic (Ranked by Diagnostic Power)
Yields: Generic AI filler and summaries.	**Yields: Actionable, specific insights.**
Example: Did sales drop across all product categories? — QUOTE	**Example: What changed in the marketing mix immediately preceding the drop?** — QUOTE
Result: Leaves multiple possibilities open.	**Result: Eliminates competing hypotheses to reveal root causes.**

Step 3: The Divergence Test

The Premise: When the entire class uses identical AI tools to analyze the same 500-800 word scenario, the only differentiator is the quality of human thought.

The Uniqueness Statement: Students must isolate exactly 100 words explaining what specific idea, connection, or insight in their analysis AI could not have produced on its own.

The Litmus Test: If you cannot articulate your unique contribution, your analysis is just an AI pass-through.

Identical
Scenario
+
Identical
AI Access

Human
Independent
Thinking

The Uniqueness
Statement
(100 words)

NotebookLM
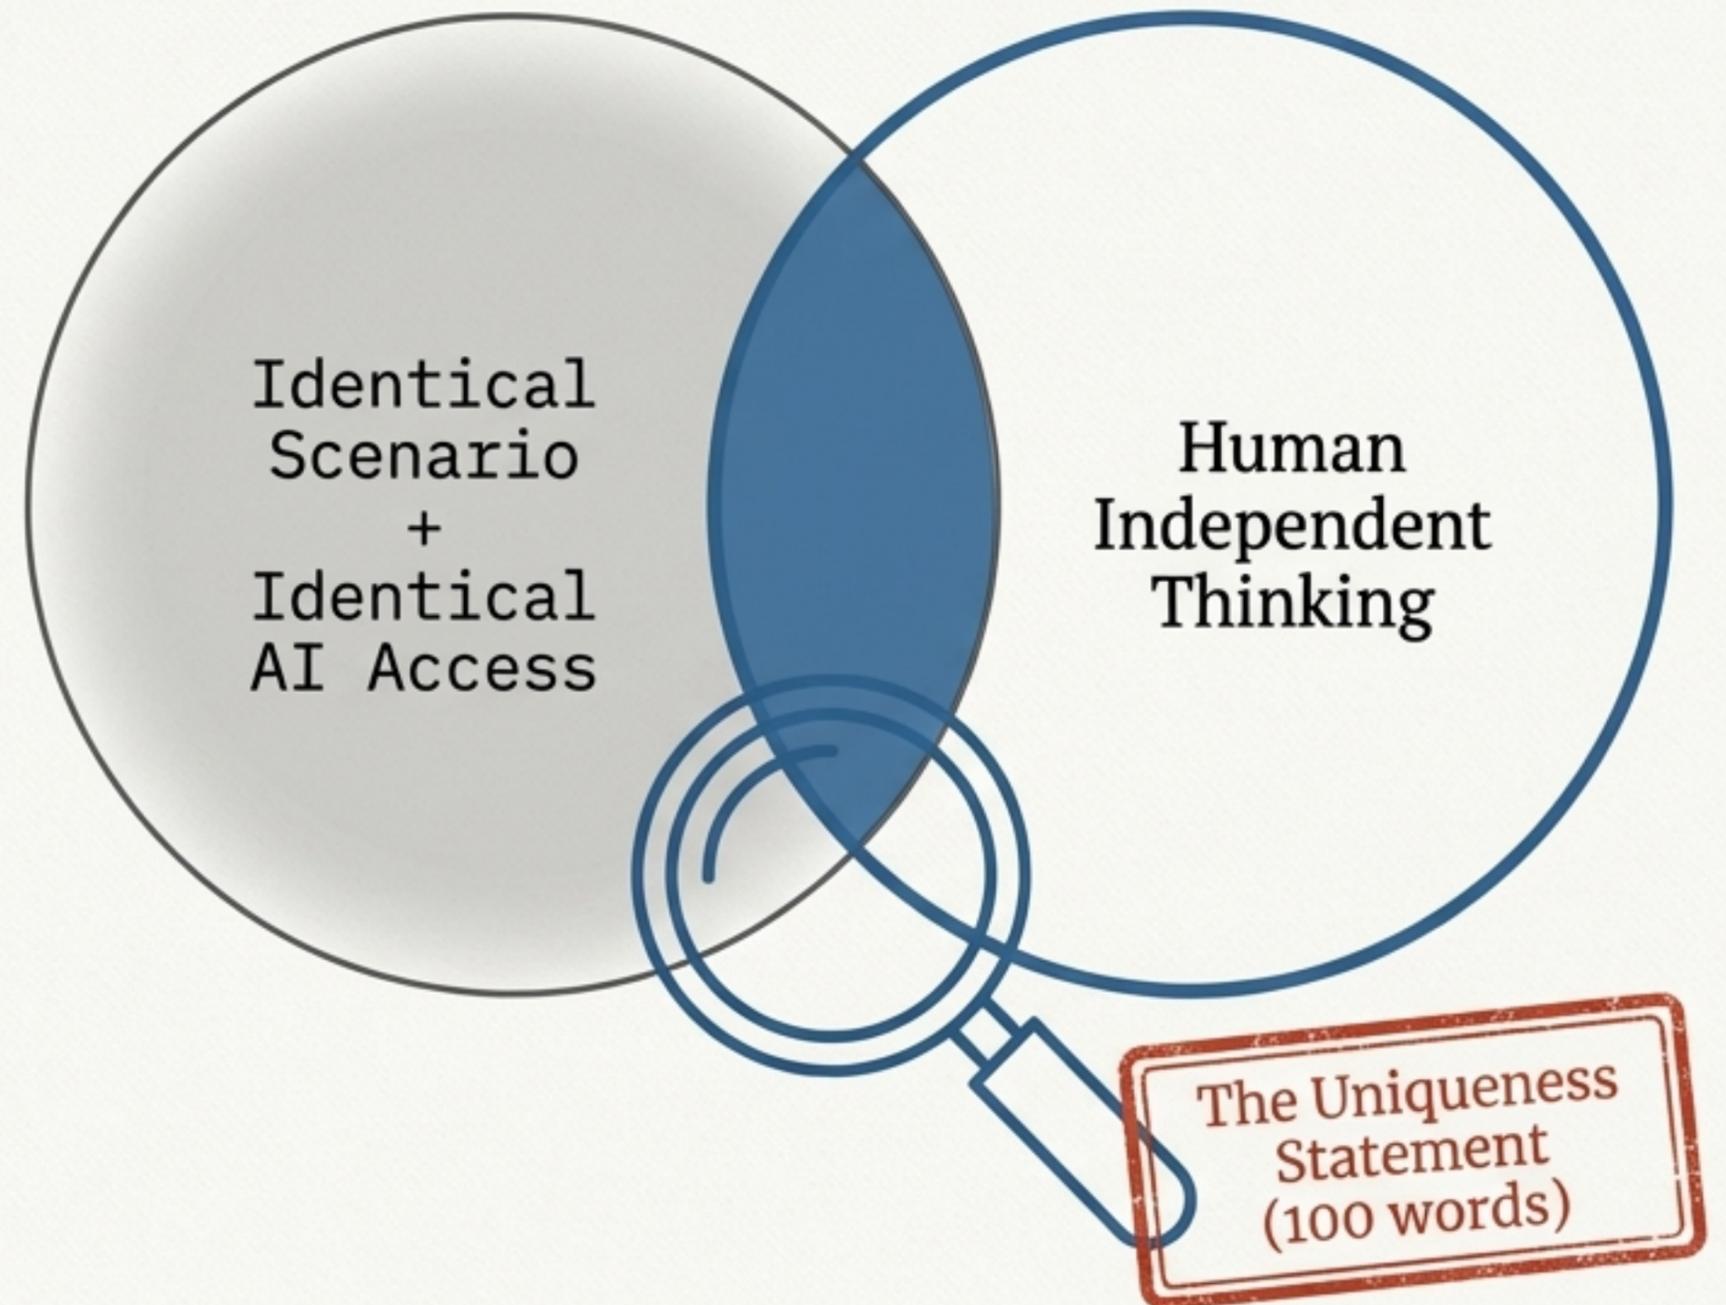
Step 4: The Live Defence

The Ultimate Stress Test: Outsourced thinking collapses under live questioning.

360-Degree Pressure Test

Unpredictable Human Attack

A peer panel asks 'why' and 'what would change your mind' to probe reasoning. (10 mins, no AI access).
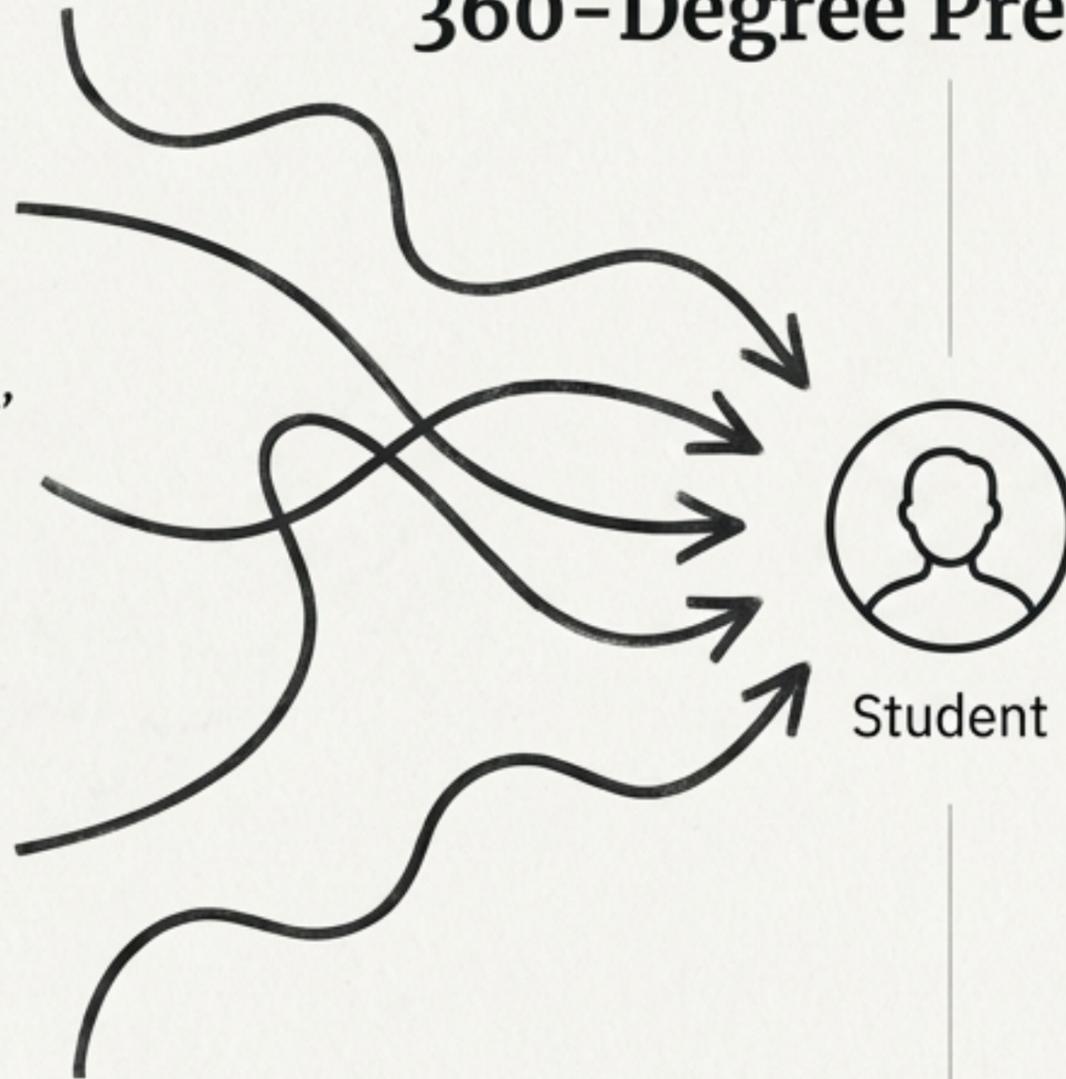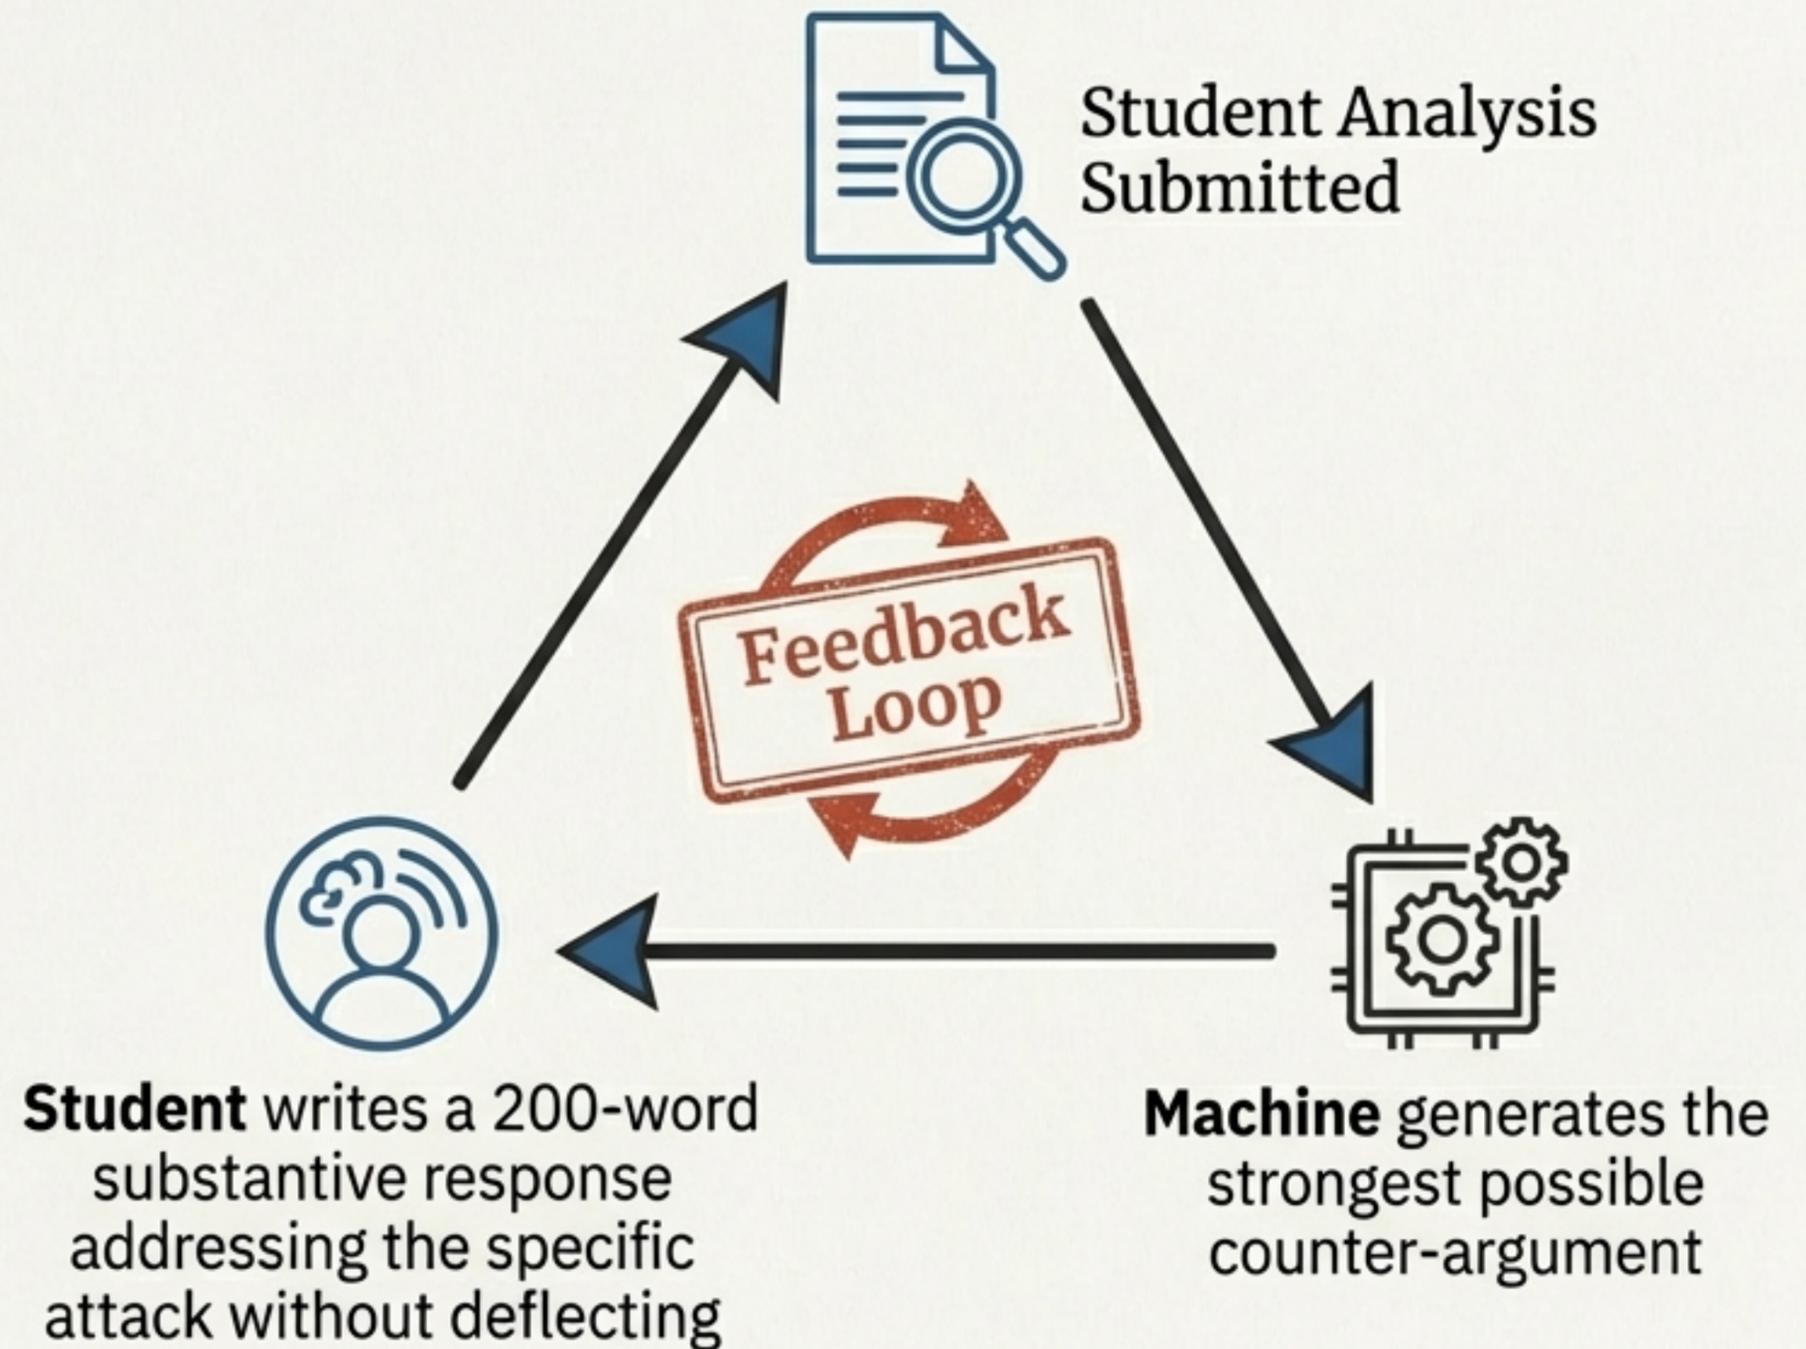
Student

Systematic Machine Attack

AI is prompted to systematically dismantle the student's logic, forcing them to defend their position against a ruthless machine critique.

The AI Counter-Argument Loop

The Mindset Shift: Treat AI counter-arguments not as attacks to survive, but as opportunities to strengthen your thinking.

The Requirement: You must respond to the specific attack with evidence and reasoning.

The Result: Understanding your own work is fundamentally different from producing it. This loop forces true comprehension.

Student Analysis Submitted

Feedback Loop

Student writes a 200-word substantive response addressing the specific attack without deflecting

Machine generates the strongest possible counter-argument

Two Ways to Use a Machine

The Mindset Shift

Passive Consumer	Active Thinker
Views AI as an "Answer Machine."	Views AI as a "Thinking Partner."
Ranks questions by topic importance.	Ranks questions by diagnostic power.
Accepts the first AI output.	Filters every output through a Reasoning Receipt.
Longer analysis = better analysis.	Values divergence and novel connections over word count.

The Question Quality Portfolio

The Thesis: The thinking process is the entire grade.

15% - Prediction Lock Quality (Coherent initial hypothesis).

15% - Question Tournament (Evaluation & ranking justifications).

25% - Divergence & Originality (Uniqueness of the reasoning receipt).

20% - Live Defence (Peer feedback performance).

15% - AI Counter-Argument (Response quality).

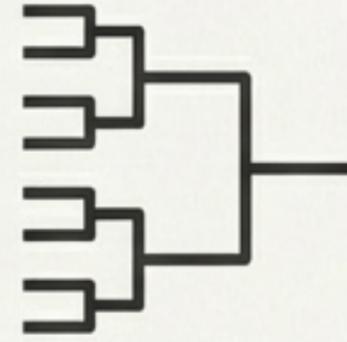**10% - Final Reflection** (Depth of self-assessment).

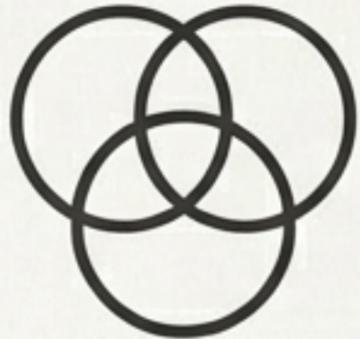Final Answer to Terrisk Business Scenario = 0%

The Process is the Product

1. Predict Before You Prompt:
Seal your hypothesis to expose your blind spots.

2. Demand Diagnostic Power:
Ask questions that eliminate competing hypotheses, not just 'important' ones.

3. Prove Your Divergence:
If you can't isolate your **unique** contribution, you are an AI pass-through.

4. Survive the Stress Test:
Use AI to systematically attack your own logic.

AI will generate the text.
You must generate the thinking.