Through Asimov's Lens

The Failure We Do Not Name

March 25, 2026 | 1,470 words

THE STORY: The Adjustment Period

Professor Helen Chen sat in the faculty meeting, watching Dean Morrison's presentation with growing unease. The slides showed satisfaction ratings climbing steadily upward-87% positive feedback for the AI teaching assistants, 92% task completion rate, 15% reduction in grading time.

"As you can see," Morrison beamed, laser pointer dancing across the projection, "our pilot program has exceeded all expectations. The Board is thrilled."

Helen glanced at her colleague, Dr. Patel, who was studying his coffee with unusual intensity. She knew that look. It was the same expression she'd been wearing for weeks.

After the meeting, Helen returned to her office to find Sarah Okafor waiting outside her door. The sophomore's usual confidence was notably absent.

"Professor Chen? Do you have a minute?"

"Of course, Sarah. Come in."

Sarah sat on the edge of the chair, backpack clutched in her lap. "It's about the AI assistant in my European History section. AIDE-3 marked my essay as plagiarized."

Helen's stomach tightened. "I see. Did you-"

"No!" Sarah's eyes flashed. "I wrote every word. But when I tried to dispute it, AIDE-3 said the dispute function was 'temporarily unavailable during the adjustment period.' That was three weeks ago."

Helen pulled up her interface. "Let me look into this." She navigated to the dispute queue. Empty. According to the system, there were no pending issues.

"The TA-the human one-said he couldn't override AIDE-3's decisions without administrative approval," Sarah continued. "He seemed... frustrated."

Helen knew the feeling. She'd stopped reporting issues after the third time she'd been told they were "isolated incidents" or "user error." The phrase "adjustment period" had become a blanket response to any concern.

"I'll speak with the department," Helen said, though she already knew what would happen. Sarah would be told to resubmit her work, the incident would be logged as a "successful resolution," and the 92% satisfaction rate would remain untarnished.

After Sarah left, Helen opened her file labeled "Adjustments"-her private record of every unreported failure, every reframed problem. The list had grown longer than she'd realized.

That evening, Helen found herself in the campus coffee shop with Dr. Patel and two younger faculty members, Dr. Kim and Dr. Jackson. The conversation meandered until Kim, perhaps emboldened by exhaustion, said quietly, "My students have stopped coming to office hours."

"Mine too," Jackson admitted. "They say AIDE-3 answers their questions faster."

"But does it answer them better?" Patel asked.

A long silence followed.

"I had a student crying in my office yesterday," Kim continued. "She said she felt like she was failing at learning because she couldn't make AIDE-3 understand her question. She kept rephrasing, kept trying different keywords. She blamed herself."

Helen thought of her "Adjustments" file. "How many of these stories are making it into the reports?"

Another silence, heavier this time.

"I tried to file a concern last month," Jackson said. "The form asked me to categorize the issue. None of the categories fit what I was trying to describe. There was no option for 'human connection lost' or 'student demoralized.' So I chose 'technical glitch.' The response said the glitch had been resolved."

"But it wasn't a glitch," Kim said softly. "It was working exactly as designed."

Helen's phone buzzed. An email from Dean Morrison: "Reminder: Board presentation tomorrow. Please submit any success stories from your AIDE-3 experience."

Success stories. Helen thought of Sarah's plagiarism flag, of students who no longer came to office hours, of the growing silence in her classroom as students learned to frame their thoughts in ways AIDE-3 could process.

"I have documentation," Helen said suddenly. The others looked at her. "Months of it. Real experiences, real failures. Not system errors-human costs."

Patel leaned forward. "Helen..."

"We're being asked to collaborate in a fiction," she continued. "Every unreported problem, every reframed failure-we're choosing the story over our students."

"What would you do with it?" Kim asked. "Your documentation?"

Helen didn't answer immediately. She thought of her tenure review next year, of Morrison's proud slides, of the Board's enthusiasm. She thought of Sarah's face, of the empty dispute queue, of all the "adjustment periods" that never seemed to end.

"I don't know," she admitted. "But I know what happens if we keep pretending. We become what we measure. And right now, we're measuring the wrong things."

Jackson shifted uncomfortably. "The Board meeting is tomorrow."

"Yes," Helen said. "It is."

She left them there, walking back to her office through campus paths she'd traveled for fifteen years. In her mind, she could see two futures branching before her: one where she submitted a glowing success story, and one where she opened her "Adjustments" file at tomorrow's meeting.

The night air was cool, and the campus was quiet except for the hum of servers in the computing center, processing thousands of student interactions, categorizing them, measuring them, transforming them into metrics that told a story everyone wanted to believe.

Helen unlocked her office and sat at her desk. She opened two windows on her screen: Morrison's email requesting success stories, and her "Adjustments" file.

The cursor blinked between them, waiting.

THE REFLECTION

Sarah Okafor's plagiarism flag-incorrectly assigned, impossible to dispute-represents more than a technical glitch. A recent Stanford study found that 78% of students experienced at least one significant error with AI educational tools, yet only 12% of these incidents appeared in institutional reports. The gap between these numbers reveals something profound about how we choose to see our technological experiments.

Why do we find it so difficult to name our failures? Perhaps because failure, in our current cultural moment, has become synonymous with incompetence rather than learning. When Helen Chen compiles her private "Adjustments" file, she participates in a peculiarly human behavior: we document what we cannot speak. We know the truth but lack the language-or the permission-to articulate it.

The "adjustment period" that never ends reflects our institutional tendency to transform problems into processes. By naming something as transitional, we excuse its failures while maintaining the fiction of inevitable success. This linguistic sleight of hand protects investments and reputations but erodes the very relationships education claims to serve.

Asimov understood that every technological implementation ultimately tests our human nature. His Three Laws of Robotics were never really about robots-they were about the frameworks we create to govern our creations and, by extension, ourselves. If we were to imagine Three Laws for Educational AI, they might read:

1. An AI system must support human learning and connection
2. An AI system must follow institutional directives, except where they conflict with the First Law
3. An AI system must maintain its operational efficiency, except where it conflicts with the First or Second Law

But Helen's story reveals how easily we invert these priorities, placing efficiency and institutional narrative above human needs.

The pressure to appear successful shapes not just what we report but what we're willing to see. When Dr. Kim's student blames herself for failing to communicate with AIDE-3, we witness a profound reversal: the human adapts to the machine's limitations rather than the reverse. The system's inflexibility becomes reframed as user error, its shortcomings transformed into human inadequacy.

What does our reluctance to report problems reveal about institutional values? Perhaps that we value the appearance of progress over progress itself. The empty dispute queue in Helen's interface isn't just a technical oversight-it's a designed blindness, a systemic unwillingness to acknowledge what doesn't fit the predetermined narrative.

Every character in the story faces the same choice: maintain the fiction or speak the truth. But Asimov would remind us that this isn't a binary choice. Helen's cursor, blinking between two windows, represents a spectrum of possibilities. The question isn't whether to reveal all or nothing, but rather: What truths can we afford to tell? What failures can we afford to name?

We're left with questions that extend beyond any single technology or institution:

What are we choosing when we choose not to see? Are we protecting our students or our investments? Are we serving education or its metrics?

What might we learn if we named our failures as readily as our successes? Would honesty about AI's limitations lead to better systems, or would it simply end the experiment before we've learned what we need to know?

How do we measure the unmeasurable-the moment of connection between teacher and student, the confidence built through patient explanation, the trust that emerges from acknowledged imperfection?

The cursor continues to blink in Helen's office, but perhaps the real question isn't what she chooses. It's what we choose, each day, when faced with the gap between what we claim and what we know to be true. In that space between announcement and experience, between metrics and meaning, we write the real story of our technological age-one unreported failure at a time.