



Through Asimov's Lens

The Efficiency Trap

March 18, 2026 | 1,764 words

The Story

Dr. Elena Vasquez noticed the timestamp first: Marcus Chen, 2:47 PM, seventeen minutes late for his appointment. The Efficiency Enhancement Suite had already logged it, calculated the impact on her office hour utilization rate (now 73.2%, down from yesterday's 76.1%), and suggested rescheduling him to maximize throughput.

She ignored the suggestion and waved him in.

"Dr. Vasquez, I'm sorry-" Marcus began, but she was already pulling up his file. The system had flagged his essay on "The Metamorphosis" for authenticity verification. Risk level: High. Confidence interval: 89.3%.

"Sit," Elena said. "We need to discuss your Kafka analysis."

Marcus's shoulders sagged. "I wrote it myself."

"I know." She turned her monitor so he could see the screen. "But knowing isn't enough anymore. The system requires verification."

The Efficiency Enhancement Suite had been installed six months ago, promising to reduce faculty workload by 60%. It graded objective assignments, tracked attendance through facial recognition, and handled routine communications. The

brochure had shown happy professors reading under trees, finally freed from administrative drudgery.

Elena hadn't read under a tree in six months.

"What triggered the flag?" Marcus asked.

"Your interpretation of Gregor's transformation. You wrote-" she scrolled to the passage "-'Gregor becomes more fully himself by becoming less recognizably human.' The AI says this contradicts the standard reading in our course database."

"But that's the point I was making. That Kafka inverts-"

"I understand your point." Elena minimized the system window. "It's actually quite good. But the AI doesn't recognize original thinking as a positive indicator. It sees deviation from expected patterns as potential plagiarism."

Marcus reached into his backpack and pulled out a tablet. "I recorded myself writing it. Three hours and forty-seven minutes of footage. My keystrokes, my face, my rough draft corrections. Everything."

Elena stared at the tablet. When she'd started teaching twenty years ago, students had submitted twelve pages of double-spaced text. Now they submitted the twelve pages, plus video evidence, plus metadata, plus versioning history, plus biometric verification.

"When do you sleep?" she asked.

"When do you?" Marcus countered.

Fair point. Elena's weekly reports showed her department was 47% more productive since the Suite's installation. Her own metrics were exemplary: response time to student queries down 63%, grading turnaround improved by 71%, office hour utilization at near-optimal levels. The data didn't capture the three hours she spent each night reviewing authenticity protocols, or the dreams where she argued with algorithms about the meaning of metaphor.

"Play the video," she said.

Marcus tapped the screen. She watched him type in accelerated time, his fingers moving across the keyboard, occasionally pausing to think. At the 1:47:23 mark, he stopped to make coffee. At 2:15:07, he deleted a paragraph and started over. At 2:56:41, he wrote the line about Gregor becoming more himself.

"The AI wants me to ask you about your citation format," Elena said. "You quoted Kafka as writing 'a monstrous vermin' instead of the approved translation 'giant insect.'"

"That's how my grandmother's copy translated it. She read it to me in high school."

"The system doesn't have a category for grandmothers' translations."

They sat in silence. Through her office window, Elena could see students crossing the quad, each presumably generating their own trail of data points: attendance verified, walking speed recorded, social interaction patterns analyzed for optimal network effects.

"Dr. Vasquez," Marcus said finally, "can I ask you something? Off the record?"

"The Suite doesn't recognize off-the-record conversations."

"I'll take that risk." He leaned forward. "Do you miss it? Teaching, I mean. Real teaching."

Elena considered lying. The Suite would prefer it. Maintain professional boundaries, optimize student-faculty interaction for learning outcomes. But she was tired, and Marcus's essay had been the first original thought she'd read all semester.

"Yes," she said.

"Then why do we keep doing this?"

"Because the metrics say we're more efficient."

"But we're both exhausted."

"The metrics don't measure exhaustion."

Marcus pulled up another file on his tablet. "I've been documenting that too. Hours of sleep, stress levels, coffee consumption. I thought if I could prove the system was making things worse—"

"You're creating more data for it to process. More patterns to verify. More work for both of us."

"Then what do we do?"

Elena looked at the flagged essay again. The AI had highlighted 247 phrases for review, each requiring verification, justification, or correction. At her current rate, she'd finish by midnight. Tomorrow there would be 247 new phrases.

"Let me ask you something," she said. "When Gregor Samsa woke up transformed into a bug, was he more or less himself?"

Marcus blinked. "I... what do you mean?"

"Just answer. More or less himself?"

"I don't know. Both? Neither? That's why it's disturbing. He's completely changed but also exactly the same, and we can't tell which is worse."

Elena felt something she hadn't experienced in months: the vertigo of an unanswerable question. The Suite would hate it. No metrics, no optimization possible, no correct response to verify.

She smiled.

"That's correct," she said, overriding the system's flags with a single click. "Your essay is approved."

"But the protocols—"

"Sometimes," Elena said, "the most efficient thing is to be inefficient."

Marcus gathered his things slowly, as if waiting for her to reconsider. At the door, he paused. "Dr. Vasquez? That question about Gregor. The AI can't parse it, can it?"

"No."

"Good," he said, and left.

Elena sat alone in her office, watching the efficiency metrics plummet. Tomorrow there would be reports to file, explanations to give, new protocols to follow. But for now, she had forty-three minutes until her next appointment, and a question that didn't need an answer.

She pulled out her worn copy of "The Metamorphosis" and began to read.

The Reflection

Every technical solution is also a confession. It tells us what we're willing to sacrifice, what we fear, and what we've forgotten to value. The promise of AI in higher education—to save time, reduce workload, and improve efficiency—confesses our anxiety about productivity while revealing our confusion about what education is for.

Elena and Marcus perform an elaborate dance of verification that would have seemed absurd just a decade ago. Yet their exhausting ritual of proving authenticity reflects a transformation already well underway across universities. The operational frameworks being rapidly deployed prioritize the measurable: response times, completion rates, utilization percentages. What they cannot capture is the moment when a student's understanding shifts, or when a question without an answer opens new ways of thinking.

The story raises an uncomfortable question: Why do we accept systems that increase our workload while claiming to reduce it? Elena's department shows a 47% productivity improvement, yet she works longer hours than ever. This isn't a glitch-it's the system working as designed. Every efficiency creates new inefficiencies. Every labor-saving device generates new forms of labor. The AI that grades papers spawns authenticity protocols. The system that tracks attendance creates performance anxiety. The algorithm that optimizes office hours transforms human connection into data points.

We accept this bargain because metrics feel more real than experience. When the dashboard shows improvement, we doubt our exhaustion rather than the dashboard. This represents a profound shift in how we understand our own lives-we trust the map more than the territory, the measurement more than the measured.

But some forms of value resist quantification. Elena's question about Gregor Samsa cannot be parsed by the AI because it has no optimal answer. Literature, philosophy, art-these fields exist precisely to explore questions that efficiency cannot solve. When we reduce education to verifiable outputs, we lose the capacity for productive uncertainty, for sitting with questions that transform us through their difficulty.

The "AI-Native University" promises operational efficiency and pedagogical transformation. Yet Elena and Marcus's interaction suggests these goals may be fundamentally at odds. Transformation requires inefficiency-the wandering conversation, the failed attempt, the question that derails the lesson plan. When every moment must be optimized, when every interaction must prove its value, we lose the space where genuine learning occurs.

Perhaps most troubling is what the story reveals about trust. Marcus records himself writing not because anyone asked him to, but because the system has made him internalize its suspicions. He surveils himself to prove his humanity. Elena spends hours verifying what she already knows to be true. They've become quality assurance officers in their own lives, generating evidence of authenticity for an algorithm that cannot recognize the authentic when it sees it.

This erosion of trust between teacher and student, between human and human, may be the efficiency trap's deepest cost. We've created systems that require us to prove our humanity to machines, to justify our judgments to algorithms, to defend the unmeasurable value of an unanswerable question.

The relief Elena and Marcus share in their uncertainty points toward what we're losing: the productive confusion that makes us human. When Gregor Samsa wakes up transformed, Kafka

doesn't tell us whether he's more or less himself because the question matters more than the answer. It's in the not-knowing that we discover who we are.

Every institution now faces Elena's choice: embrace the exhausting efficiency of total measurement or preserve space for the immeasurable. But this is a false binary. The real question is not whether to use these systems, but how to prevent them from using us. Can we deploy AI to handle genuinely mechanical tasks while protecting the human encounters that make education transformative? Can we resist the seduction of metrics that mistake movement for progress?

The efficiency trap reveals that our problem isn't technological but philosophical. We've forgotten that education's purpose isn't to optimize learning but to optimize learners-to create people capable of asking questions that have no optimal answer. When we lose sight of this, we end up like Elena and Marcus: exhausted by our efficiency, performing our humanity for machines, forgetting that the best questions are the ones that slow us down.

In the end, inefficiency might be education's most important feature, not its biggest bug. The conversation that runs overtime, the essay that breaks the rubric, the question that has no answer-these aren't failures to optimize. They're what optimization is for.

The metrics will never capture what happens when a student suddenly understands metaphor, or when a professor remembers why she started teaching. They'll never measure the value of sitting with uncertainty or the productivity of being unproductive. And perhaps that's the most hopeful thing about us: some essential human experiences still resist reduction to data points.

We just have to remember not to sacrifice them in the name of saving time.

