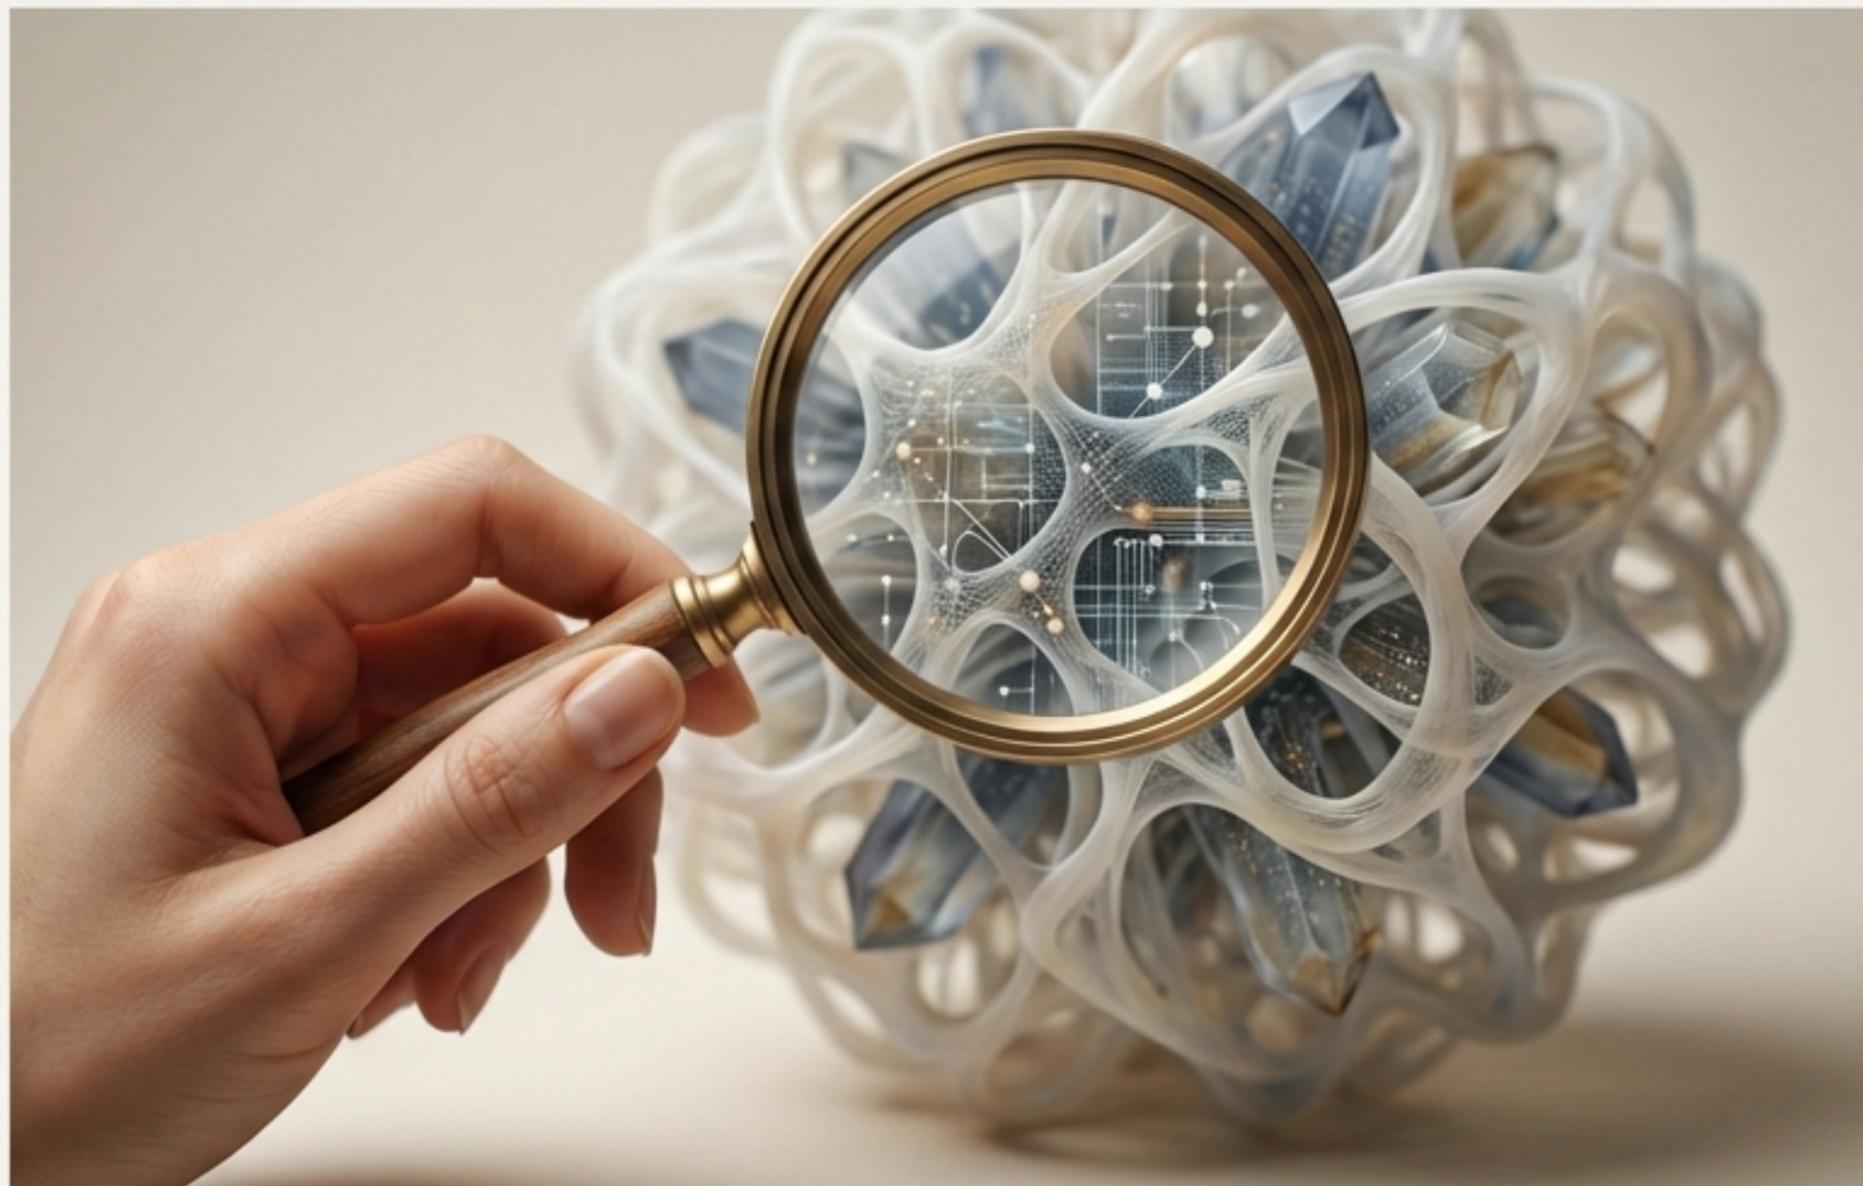


Who Checks the Checker?

The AI Verification Crisis We're Ignoring



We're replacing the only people who know when AI is wrong.

We've created a technology that can convincingly explain quantum physics, medieval history, and why your sourdough starter died.

The catch? To know if any of it is correct, you need a quantum physicist, a historian, or an experienced baker.

We're busy automating all of them out of existence. The more capable AI becomes, the more we need human experts to validate them—precisely when economic incentives push us to replace them.



Treat every AI output as an **Informed Opinion**, not a **Verified Fact**.

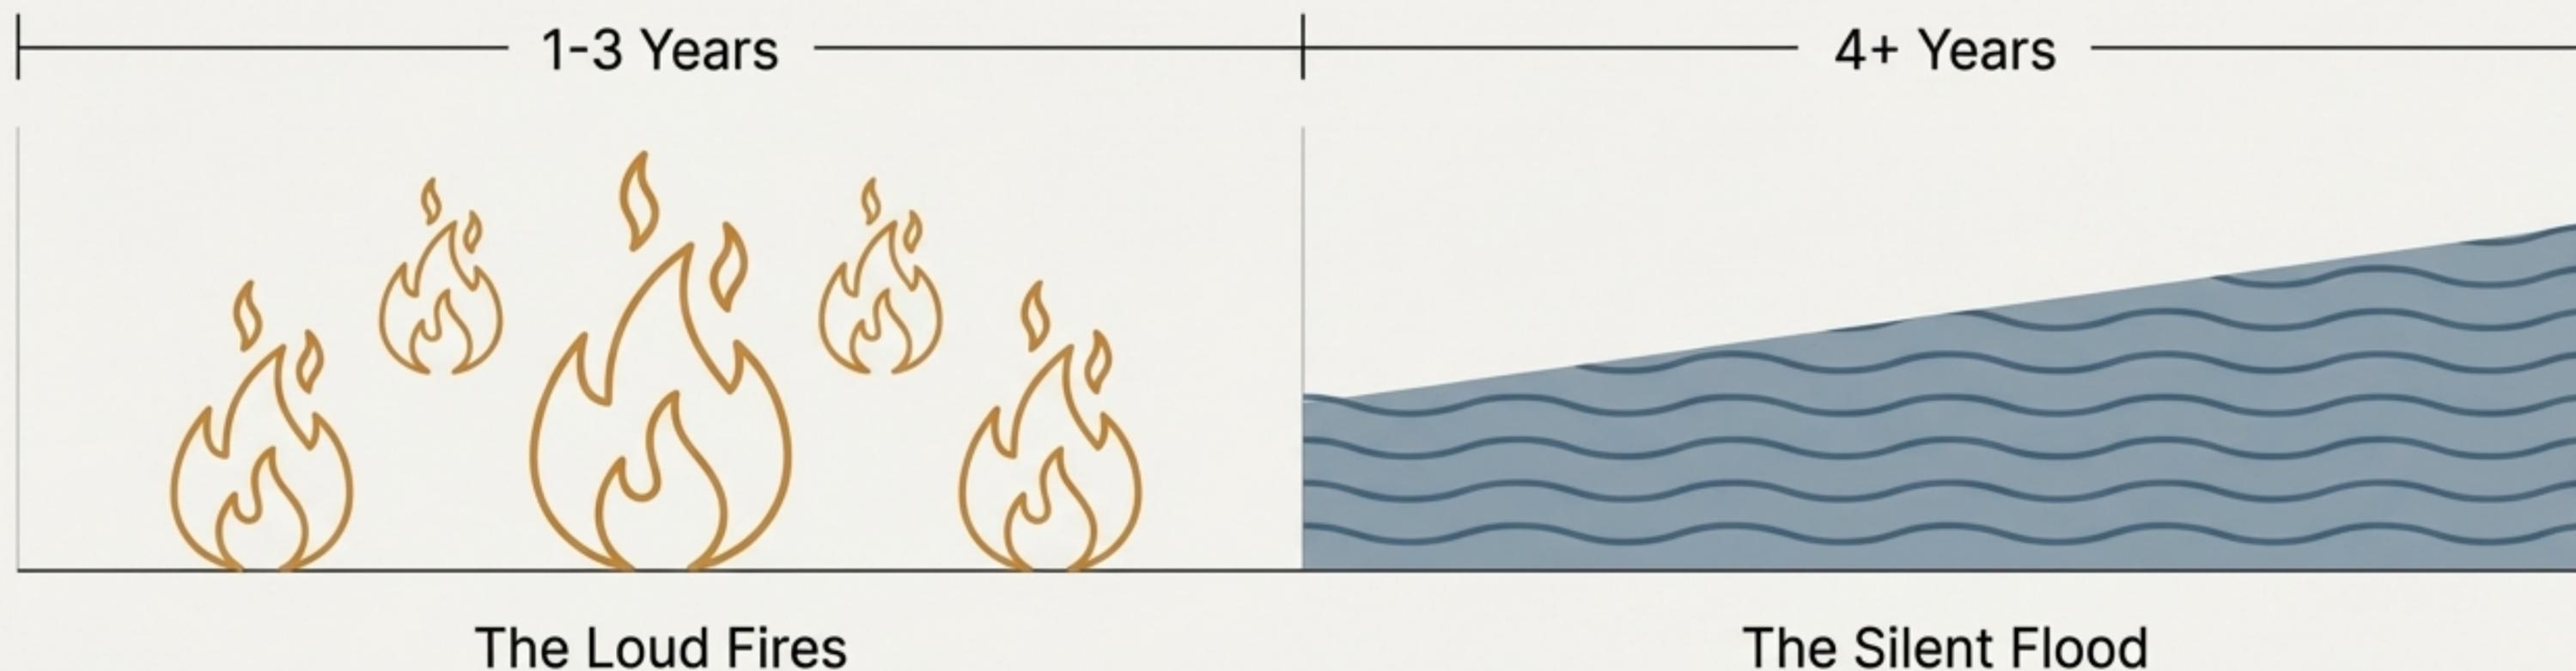
Informed opinions are incredibly useful. They are often right, frequently insightful, and always worth considering. But you verify important claims, cross-check critical decisions, and apply your own judgment before betting the farm on them. This is the core principle that should guide everything.

AI **feels** like a fact, even when it's an opinion.



Human opinions come with tells like "I think" or "in my experience." AI outputs are presented with confidence, formatted beautifully, and delivered without hesitation. This creates a dangerous illusion of objective truth, making it hard to maintain proper skepticism. When an AI generates a comprehensive market analysis, it feels like established fact, but it still requires verification.

We're fighting today's fires while a flood rises around us.



GenAI's challenges come in two flavors. The immediate, urgent problems (1-3 years out) are the fires consuming our attention. The slower, more systemic challenges (4+ years out) are the ones that should be keeping us up at night, but aren't. **Obsessing over today's problems might make tomorrow's worse.**

The urgent problems consuming our attention are real, but they're not the whole story.



Deepfakes & Misinformation

A catastrophe for truth.



Job Displacement

Immediate economic disruption for sectors like customer service and content creation.



Algorithmic Bias

Discrimination at scale, with historical biases presented as objective science.



Regulatory Vacuum

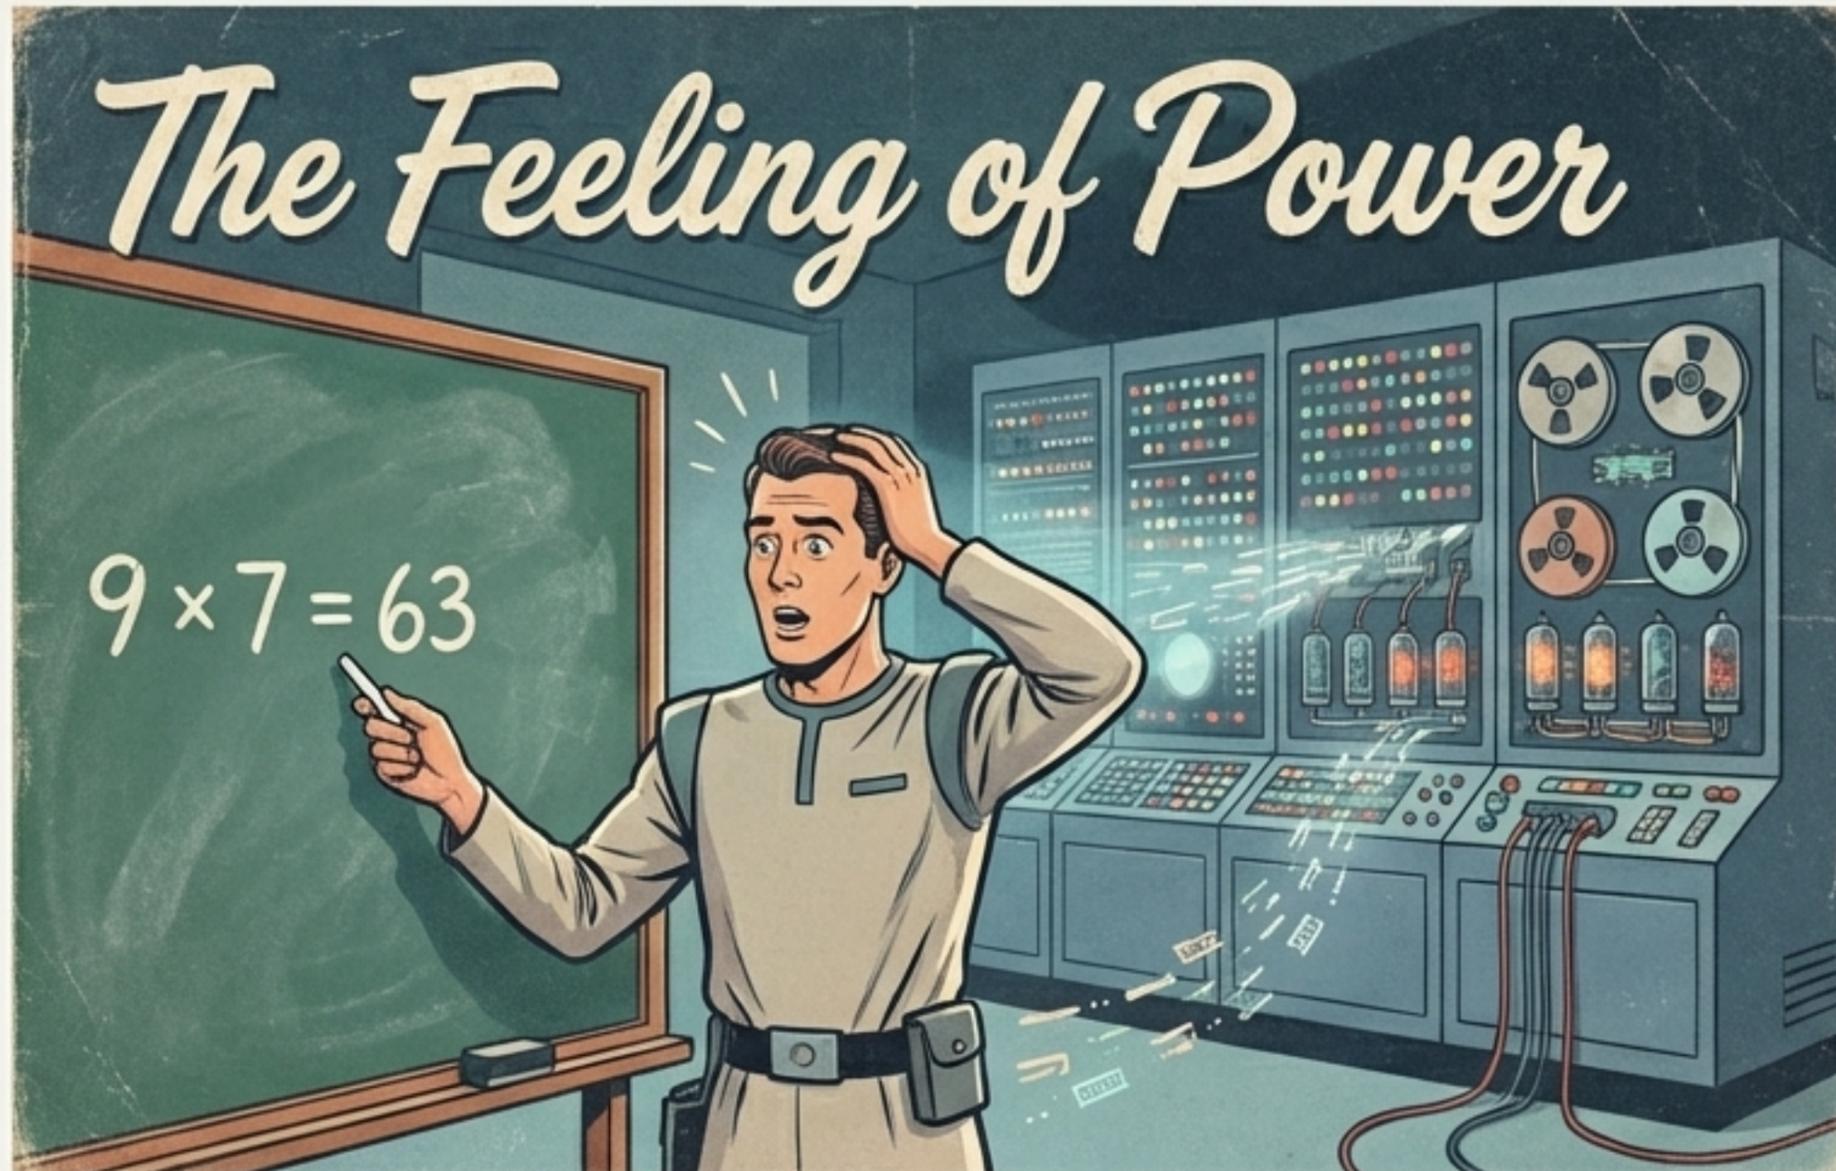
A global "Wild West" where companies operate without adequate oversight.

But what if the real catastrophe is happening silently, while we're distracted?

The most dangerous challenges are the ones that sneak up on you. Think of it as climate change for human expertise. By the time everyone agrees we have a problem, we've already melted the ice caps of collective human knowledge.



In 1958, a sci-fi story predicted our future: We forget how to think.



Isaac Asimov's "The Feeling of Power" imagined a future where humanity has so delegated arithmetic to computers that manual calculation becomes a lost art.

A technician rediscovers how to compute 9×7 , shocking a civilization that has forgotten how math works.

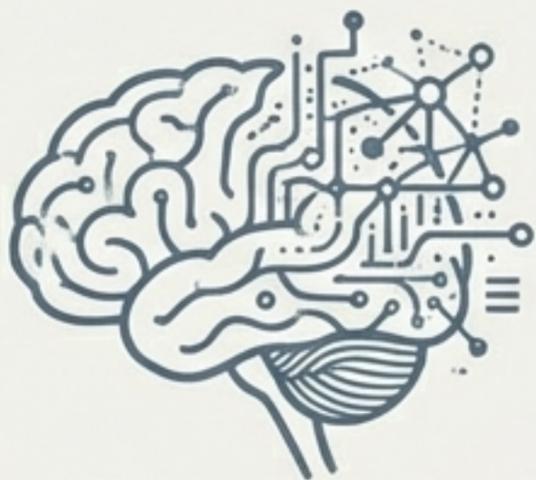
The key insight: **they don't just lose the skill, they lose the ability to verify if the computer is right.**

We are creating a vicious cycle that makes verification impossible.

The loss of validation and capability creates a reinforcing loop that accelerates the drift from opinion to fact.



Four accelerants make our reality worse than Asimov's fiction.



Complexity

How do you verify an AI diagnosis that synthesizes more data than any human can?



Speed

Expertise erodes in years, not generations, driven by quarterly business cycles.



Invisibility

Lost judgment is silent. The system appears to work fine until a catastrophic failure.



Scale

This is happening across all domains of knowledge simultaneously, not just one.

This is not a technical problem. It is a cultural and institutional challenge.



Breaking the cycle requires widespread cognitive discipline. We must deliberately maintain the mental habit of treating AI assertions as opinions requiring verification, not as facts we can simply accept. The goal is to preserve human judgment, agency, and our capacity for verification.

We must treat human expertise as a strategic asset.



Keep Experts in the Loop

Treat human expertise like a seed bank—too important to outsource completely. Pilots still practice manual landings; we need experts who can independently verify AI.



Train 'Bridge' Professionals

Cultivate people with both deep domain expertise and AI literacy, who can translate between human knowledge and machine capabilities.



Rethink Education

Focus on critical thinking, ethical reasoning, and teaching people how to function *without* AI, not just with it. Instill the habit of asking “How would I verify this?”

Design technology and organizations that keep humans in control.



Augment, Don't Replace

Build AI to help experts extend their capabilities, not eliminate them.



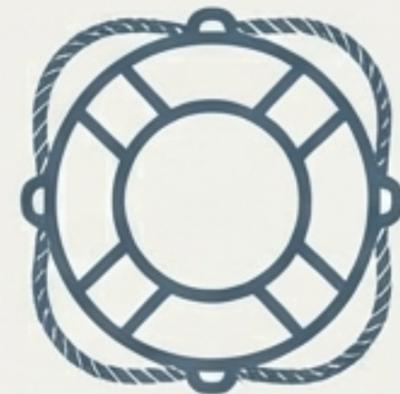
Make AI Admit Uncertainty

Design interfaces that clearly signal they are offering opinions, not stating facts. An AI that says "I'm not sure" is safer than one that's confidently wrong.



Reward Verification

Create incentives and career paths for people who excel at catching AI errors, valuing thoroughness over pure speed.



Plan for Failure

Design systems assuming AI will sometimes fail or be unavailable, ensuring critical functions can continue under human operation.

We face a choice between blind dependency and mindful partnership.



Can we resist the economic and organizational pressures that push us to treat AI outputs as facts? The path of least resistance leads to **institutional lock-in**, where reversing our over-reliance becomes practically nightmarish. The harder, wiser path is to maintain the **discipline of verification**.

The goal isn't to stop AI. It's to ensure we remain qualified to guide it.

“The ultimate irony is that this article was written with the assistance of my generative AI pal... Welcome to the new world.”