



Reject the Doomer/ Utopian drama in AI

by **Pietman Roos**

F It turns out that sprinkling artificial intelligence (AI) magical dust on a business does not automatically turn it into gold. That is the dumbed-down version of the recent MIT study into the implementation of 300 generative AI projects¹, of which 95% experienced zero return on investment. This relatively unknown paper, first released in July 2025, was picked up by Fortune² and published again in August 2025, and in combination with warnings from OpenAI's head, reportedly led to a 3.5% drop in Nvidia and 10% drop in Palantir equity prices³.

The different responses to the research are more revealing than the findings themselves. On the "Utopian" side, it's a blip on the supercycle of AI development that will end in superintelligence. The dot-com crash is frequently cited as a comparison to the market reaction to the MIT report⁴, which at the time looked bleak, but ultimately has led to the world changing internet far beyond early millennial expectations.

On the "Doomer" side, which views too-powerful AI as a systemic risk in the long run, the report is seen as evidence that AI has always been overhyped and that the chaos from the introduction of AI into every facet of business has now come to pass. Doomers probably feel a mixture of relief and sheepishness in that generative AI is not becoming as powerful as quickly as expected. Ultimately, suggesting that both sides share a common view that AI will reach superintelligence, the difference being whether it will be malign or not.

Some of the most salient findings from the MIT report are that custom-built generative AI fails due to integration complexity and poor fit with existing workflows, and that seven out of nine industries (including finance) show no evidence of structural change despite investment into generative AI.

One finding compares it to an employee that must be reminded not to make the same error; often custom generative AI models do not learn the intricacies of the business, so its usefulness is capped to simple tasks or queries. And like an employee that needs constant reminders, the temptation is to either do the work yourself or ask someone else, in this case the subscription-based large language models (LLMs). The report goes so far as to call this a "shadow AI economy" where employees use their personal accounts to automate some work, often without IT department knowledge, and often where a custom model is already available, with the estimate that 90% of the survey reported this phenomenon.

The problem with a Doomer/Utopian approach to risk management (and regulation) is that it simultaneously over-and-underestimates the risk involved with generative AI. The probability of a superintelligence that will either exterminate us or solve all the world's problems is frankly overdone because it is imprecise to extrapolate mass societal change beyond the immediate and observable technological capacity. Conversely, the actual problems with generative AI are

1. MIT NANDA, [The GenAI Divide STATE OF AI IN BUSINESS 2025](#), July 2025

2. Fortune, [MIT report: 95% of generative AI pilots at companies are failing](#), 18 August 2025

3. Fortune, [U.S. tech stocks slide after Altman warns of 'bubble' in AI and MIT study doubts the hype](#), 20 August 2025

4. Business Insider, [3 reasons everyone is talking about an AI bubble](#), 25 August 2025

overlooked because they are boring when compared to a science-fiction saga⁵.

The same applies to capital markets; overdue focus on sentient rogue trading algorithms that will cause mass market failure or grand scale Utopian visions of Billion-Dollar One-Person companies⁶ misses a more fundamental question: are the basics getting done right? Or rather, is your staff actually using the custom AI-tool you spent millions on? A recent IOSCO report shows that the majority of the use cases for AI across financial sectors are AML/CFT, internal productivity support, and market analysis and trading insights⁷. But with the MIT report findings in mind, it could be that despite the investments in models to assist with screening or random tasks, staff are opting out of custom AI-tools and using public LLMs.

This has implications on risk and what regulators would want to understand of their licensees for any technology or process. Actual usage is something that regulators can query and test, which is conceptually less demanding than having to delve the innards of an AI model. It is safer and arguably more accurate for a regulator to simply ask “who uses this,

how much and for what tasks, how do you record this use, what can go wrong, and how do you mitigate that risk?” than to prescribe statistical metrics to, for example, ward against model bias. Of course, over time these metrics will enter the conversation. But this assumes custom generative AI is used at all. For public LLMs, regulators would have to form a view of each provider.

Anyone growing up in the 90s would be familiar with the then flourishing cyberpunk genre influenced by the nascent internet, and its fixation on cyberspace and how we would all uplink through neural implants at the turn of the millennium. Strangely, that has not come to pass, but cybersecurity and constant IT updates is a fact of life. The same could be said of the drama currently attached to generative AI, and the need to perhaps think more boringly, more bureaucratically.

Pietman Roos is Senior Manager, Market Regulation, The International Stock Exchange (TISE)

5. [The Atlantic, The AI Doomers Are Getting Doomier, 21 Aug 2025, Wong M](#)

6. [Forbes, The Future Is Solo: AI Is Creating Billion-Dollar One-Person Companies, 17 February 2025](#)

7. [IOSCO, Artificial Intelligence in Capital Markets: Use Cases, Risks, and Challenges, March 2025](#)