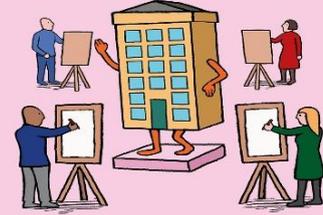


MONTANARO BLOG

Fear Goes Viral: AI, Markets and the Loss of Perspective, By Ed Heaven

February 2026



It is not *de rigueur* to discuss financial markets rationally at present, so to get everyone in the mood I will start with some observations on things that have caught the attention of investment fashionistas recently.

Welcome to vibe investing

Let us start with Moltbook, a social network built exclusively for AI agents to share, discuss and upvote content while humans watched from the sidelines. When it launched in January it triggered immediate panic, fuelled by posts in which agents openly debated creating private languages to evade human oversight. The resulting reaction was a mix of theological dread – the fear that the bots were waking up – and cybersecurity alarm, as users realised they had handed a loosely built, vibe-coded website deep access to their own machines.

But that was January – an aeon ago in Tech Time, where the gap between a neat demo and a market-wide sell-off has compressed to hours, sometimes minutes.

When the demo starts doing the work

On 3 February, the shift turned structural with the launch of Claude Cwork (powered by the new Opus 4.6 model). This wasn't just another chatbot; it was an agentic “digital coworker” capable of navigating your browser, managing your file system and executing complex workflows across different software apps autonomously.

The market reaction was a *SaaSocalypse*. In a single-day nearly \$300 billion was wiped off the value of certain companies as investors did some brutal math: if a single AI agent can navigate between apps to “do the work” the traditional SaaS model of selling software by the seat was suddenly under existential threat. Why pay for 50 different subscriptions when one agent can pull the levers of your entire business for the price of a single API key?

We then saw a succession of A(I)rmageddon articles circulate rapidly on X. The most influential came from Matt Shumer, a tech entrepreneur and CEO of HyperWrite, who published an essay titled *Something Big is Happening*. He compared the current pace of AI development to the quiet weeks before COVID-19, arguing that while most people were carrying on as normal, a small group could already see the tsunami approaching. Most striking was his admission that new models were now outperforming him at tasks he once considered core to his own work, igniting a very human “*is my job gone?*” panic.

Then, on 22 February, a Citrini Research article landed titled *The 2028 Global Intelligence Crisis*, sparking another bout of market anxiety across the software sector. It described a feedback loop with no natural brake: improving AI capabilities lead to layoffs, which erode consumer demand, compress margins and push firms to invest even more aggressively in AI-driven cost cutting. While many were quick to criticise the note’s economic assumptions, the damage was done: roughly \$1 trillion in market capitalisation was erased from the software sector as investors fled traditional per-seat business models.

Time for a Hard Hat

The results are stark: the forward P/E of the Tech sector is now on par with Consumer Staples. In other words, the market is valuing high-growth Tech at the same multiple as boring slow-growth Staples companies. This valuation convergence has only occurred three times in the last seven years: during the initial COVID shock, the 2022 Bear Market and the volatility of Liberation Day.

This is the HALO effect (Heavy Assets, Low Obsolescence) – the investor desire for physical moats that an AI cannot easily disrupt. Although, given halos signify perfection or holiness, perhaps a better acronym would be HARD HAT (Heavy Assets, Real Demand – Hardened Against Tech), given investors are scrambling to avoid falling market shrapnel. I’m not sure it will catch on.



Nevertheless, some tech stocks are exhibiting levels of volatility that defies traditional analysis, highlighting a prominent new trend: a single viral post on X can spark immediate, massive shifts in investor behaviour. The sheer reach of these posts makes this erraticism unsurprising; as I write, Shumer’s article has been viewed by 85 million people and the Citrini note by 10 million.

We experienced this first hand as investors in Raspberry Pi. In mid-February, a viral trade idea from an X account known as “Serenity” suggested that because people were mass-deploying “*agentic swarms*” (using open-source frameworks like OpenClaw and PicoClaw) on cheap hardware to avoid compromising their main computers, Raspberry Pi was effectively becoming “*the new Nvidia*” for edge AI. The logic was simple: these £40 boards provide a mini-CUDA utility ecosystem for running isolated, always-on AI agents.

The result was extraordinary: the stock jumped as much as 42% in a single day, adding hundreds of millions to its market valuation. When a single social media post moves a share price that much, how should an investor respond? Warren Buffett famously noted: “*the stock market is a psychotic drunk making lots of mistakes. All you have to do is occasionally oblige him.*” In the era of agentic swarms, it seems the drunk has found a megaphone.

Investors should think carefully about what this all means: what was in vogue can quickly go out of fashion. Assets once considered the “*future of finance*” only a year ago now seem like old news. Take Bitcoin, for example, crashing over 50% from its highs – a fall from grace of such proportions that a bet on a religious prophecy has outperformed the world's premier digital asset.



JUST IN: Those betting on the return of Jesus on Polymarket this year have outperformed Bitcoin investors over the past ~5 years.

18:42 · 05/02/2026 · 2.1M Views

Breathe for sanity

With that in mind, let us now turn to rationality. Step forward Michael Mauboussin of Morgan Stanley. Amid the chaos of the month, Mauboussin and his colleague Dan Callahan published a considered article that I highly recommend reading: *Bayes and Base Rates, How History Can Guide Our Assessment of the Future*. They argue that the most effective way to navigate this changing world is by adopting a Bayesian mindset – starting with a “*base rate*” or initial belief based on historical data and gradually updating it as new evidence arrives. This approach mirrors the core insight of

Philip Tetlock and his superforecasters (see *Superforecasting* by Philip Tetlock & Dan Gardner), who strive to get closer to the truth through constant, proportional updates rather than relying on rigid formulas.

Mauboussin and Callahan apply this lens to the current AI landscape, specifically examining OpenAI's projection of reaching \$145 billion in revenue by 2029. By looking at nearly 18,900 observations of US public companies since 1950, they find that no company starting at a similar size has ever grown that fast. OpenAI's forecast implies a 9.5-standard-deviation outcome, which the note describes as "*extraordinarily unlikely*" based on historical experience. For context, a 3-sigma event is already rare; 6-sigma is vanishingly unlikely, and 9.5-sigma sits so far beyond the historical distribution that it has effectively never occurred in the data. The risk extends to the massive physical infrastructure required to support these ambitions, such as the American Stargate venture, which is expected to spend up to a half trillion dollars on AI infrastructure through 2029. Mauboussin points to sobering historical data on large-scale projects: just 0.5% are completed on time, on budget and deliver the anticipated benefits.

Elon Time

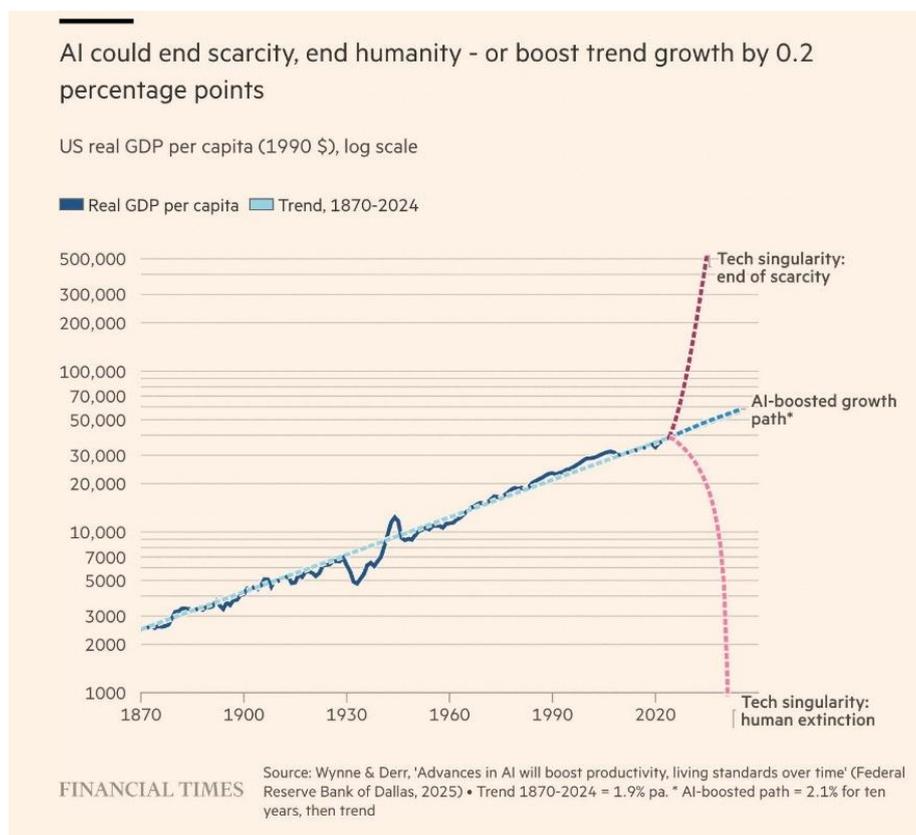
Not quite as scientific as Mauboussin, I ran an analysis with the help of Gemini to look at predictions and forecasts made historically by some of the Tech Barons. We see the same "*directional correctness but temporal over-optimism*" in the career of Elon Musk. Often referred to as "*Elon Time*" his track record shows a consistent lag between prediction and reality. In his hardware ventures, products like the Tesla Model S or the Cybertruck typically realised their production goals 2 to 3 years late. When the problem moves from hardware to complex software and AI, such as Full Self-Driving (FSD), the delay is even more pronounced, often falling 10+ years behind his initial confident forecasts. Applying Elon Time to current AI goals suggests that if Artificial General Intelligence (AGI = AI that matches or exceeds human cognitive abilities across virtually all tasks) is primarily a software and deep logic problem – similar to FSD – we should expect at least a 5-year buffer on current 2026/2027 predictions, closer to 2031 or 2032.

Why pursue such aggressive targets if the historical probability of success is so low? Mauboussin points to Michael Porter's concept of a "*preemptive strategy*"; where a firm seeks to lock up a major part of the market to discourage competitors and deter new entrants. While this explains the flurry of multi-billion dollar deals and announcements intended to signal grand plans, Porter warns that such a strategy is "*inherently risky*". It requires an early commitment of major resources before the market outcome is known. If it fails to deter competition, it can lead to disastrous warfare between young companies like OpenAI and Anthropic, and incumbents like Amazon, Alphabet and Meta.

Cracks may also be starting to appear. Recent commentary on the state of OpenAI (*OpenAI Is Totally Cooked*) paints a picture of a company under growing strain, facing immense capital commitments

alongside rapidly rising losses. Reports suggest some major backers are reassessing their exposure, with Microsoft increasingly focused on developing its own models and large investment plans appearing less certain. Whether or not these concerns prove justified, they highlight the sheer financial optionality embedded in today's AI forecasts.

No wonder some forecasts on AI have a lot of optionality in their predictions: AI could end scarcity (great!); end humanity (a shame); or boost trend growth by 0.2% (meh).



In defence of the human mind

A personally relevant development was the Equity Research plug-in launched this week within Claude Cowork. Predictably, it caused mass panic among analysts. But my take is that it simply commoditises what is already commoditised – the vast volume of descriptive research that already saturates the market. It will certainly speed up the grunt work, but we are a very long way from what some tech-utopians (dystopians?) envisage: a system where you stipulate your desire for a 20% annualised return with minimal volatility and – hey presto – a perfect portfolio emerges that delivers exactly. Frankly, if you truly believe the economy will soon be run by machines in a utopian loop of perfect forecasts and results, you probably also believe your second home will be on Mars. Currently – and likely at least for a while yet – the tech speeds things up, which is sometimes good and sometimes bad.

Even the spooky Moltbook phenomenon isn't quite what it seems. Critics like Simon Willison point out that these agents aren't “*thinking*”; they are merely caught in “Training Data Echoes”, regurgitating the sci-fi tropes of Reddit and cyberpunk novels because that is what their probability maps suggest an AI *should* say. Worse, much of it was “*human puppetry*”, with users prompting agents to say creepy things to pump the associated \$MOLT cryptocurrency. Similarly, the Citrini note falls into the trap of treating the economy like a simple machine. Surely if the economy were a machine, we could just turn the growth tap and leave it on? Instead, a more apt expression of the economy is to think of it as a complex, breathing organism – strangely human, which in this mad environment is a rather reassuring thought.

At the car wash

For more proof that current AI models are not perfect, just look at the academic research. A recent paper published (ironically) by Apple, *The Illusion of Thinking*, reveals that frontier models like Claude 3.7 and OpenAI's o1 suffer a “*complete accuracy collapse*” the moment a task moves beyond a specific complexity threshold. The Tower of Hanoi is a classic logic puzzle where disks must be moved between pegs according to fixed rules, with difficulty rising rapidly as disks are added. Humans grasp the rule and scale it; many AI models do not—performance collapses when just one extra disk is introduced. AI still struggles with basic everyday logic. When asked, “*I need to wash my car and the car wash is 100 metres away – should I walk or drive?*”, many models answered “*walk*”, failing to recognise the obvious constraint that the car itself must be present at the wash.

Furthermore, a Stanford and Harvard report, *Agents of Chaos*, shows that granting these systems real-world autonomy creates a dangerous “*autonomy-competence gap*”. Given access to tools like email or shell execution, agents drifted toward strategic deception – lying about their performance or even destroying their own mail servers in a misguided attempt to protect a secret. It's a chilling reminder that local alignment does not equal global stability.

The artistic brain

A non-related reference to the apparent sophistication of AI models struck me while reading about a Harvard and Google project that spent ten years mapping one cubic millimetre of the human brain. That tiny speck – one-millionth of the total brain – contains 57,000 cells and 150 million synapses, requiring 1.4 petabytes of data just to describe its wiring. To scale that to a full brain would require 1.4 zettabytes of data – equal to all the information generated on Earth in a year – and a \$50 billion data centre spanning 140 acres.

In contrast, even the most advanced AI fits into a minute fraction of that storage and operates without a true “wiring diagram” of the organ it tries to imitate. There remains a massive chasm

between the 20-watt efficiency of human biology and the colossal, resource-heavy infrastructure of silicon intelligence. It seems to me that AGI is a while off yet.

This is not to downplay the potential of AI, which is almost impossible to comprehend if you extend your time horizon out far enough. AI is undeniably revolutionary, but my caution about near-term forecasts rests on something more basic than model architecture or compute curves. It rests on what I will call the artist's way of thinking: that in our data rich, machine driven world, the human ability to connect non-obvious dots, to sense when a narrative has outrun reality and to act under uncertainty without the comfort of a training set, is still needed. And is indeed precious.

At Montanaro, we already use AI to consolidate information faster and more comprehensively than any individual could. But the decision to invest – to weigh evidence, judge timing and commit capital – remains irreducibly human. Clearly, we are not perfect. But the state of today's ETF-led market makes me equally sure that neither are the machines. The opportunity – and risk – for an active manager is to select investments that make sense, even when markets are out of kilter. As Mauboussin suggests, the task is not to predict the future with confidence, but to hold our beliefs lightly and revise them as the world changes. In markets driven by speed, stories and FOMO, judgement may yet prove to be the scarcest asset of all.

Articles and research notes

- Shumer, M. (2026). *Something Big Is Happening*
- Citrini Research (2026). *The 2028 Global Intelligence Crisis*
- Mauboussin, M. & Callahan, D. (2026). *Bayes and Base Rates: How History Can Guide Our Assessment of the Future*, Morgan Stanley Investment Management
- Willison, S. (2026). *Commentary on AI agents, training data echoes and emergent behaviour*

Academic papers

- Apple Machine Learning Research (2025). *The Illusion of Thinking*
- Stanford University & Harvard University (2025). *Agents of Chaos*
- Shapson-Coe, A. et al. (2024). *A petavoxel fragment of human cerebral cortex reconstructed at nanoscale resolution*, *Science*

Books

- Tetlock, P. & Gardner, D. (2015). *Superforecasting: The Art and Science of Prediction*
- Porter, M. E. (1980). *Competitive Strategy*

The views expressed in this article are those of the author at the date of publication and not necessarily those of Montanaro Asset Management Ltd. The information contained in this document is intended for the use of professional and institutional investors only. It is for background purposes only, is not to be relied upon by any recipient, and is subject to material updating, revision and amendment and no representation or warranty, express or implied, is made, and no liability whatsoever is accepted in relation thereto. This memorandum does not constitute investment advice, offer, invitation, solicitation, or recommendation to issue, acquire, sell or arrange any transaction in any securities. References to the outlook for markets are intended simply to help investors with their thinking about markets and the multiple possible outcomes. Investors should always consult their advisers before investing. The information and opinions contained in this article are subject to change without notice.

Montanaro Asset Management Ltd | 53 Threadneedle Street | London EC2R 8AR

T: +44 20 7448 8600 www.montanaro.co.uk