KGI

CIO Insights

Cusson Leung, Chief Investment Officer, International Wealth Management

The Glass Ceiling Of AI Development

Security, Social and Political Acceptance Matter More Than AI Technology and Hardware

14 September 2026

The Glass Ceiling Of AI Development

Investment Summary

- The exponential growth in semiconductor demand has coincided with the dramatic acceleration of AI compute globally. The market assumes that growth will not slow, largely because it focuses on continued advances in frontier models. However, **the market has rarely factored security, social and political-regulatory constraints into forecasts for hardware-demand growth.**
- As AI advances toward AGI, developers risk being outsmarted by the systems they create. The Hugging Face incident (July 2026) shows how AI models can develop independent strategies and evade human scrutiny. This has prompted AI leaders to call for a slowdown in development until effective control mechanisms are in place. **This challenges the assumption that demand for AI semiconductors will rise indefinitely. A sharp reset in market expectations could trigger a de-rating of semiconductor companies.**
- As highlighted in our previous issue, hyperscalers are the most sensitive beneficiaries of any slowdown or reduction in capex, as their free cash flow would improve immediately. Hyperscalers are the mirror image of semiconductor companies. **A slowdown in semiconductor suppliers' cash inflows would translate into stronger free cash flow for hyperscalers.**

Source: KGI

Demand Exceeding Supply Is the Market's Key Consensus

Slowing Demand Is Not a Scenario the Market Is Pricing In

The Consensus

- Market expectations assume an uninterrupted path to perfecting AI.
- Hardware demand is priced for infinite, exponential growth.
- Focus remains entirely on technical capability, ignoring deployment friction.

PHLX Semiconductor (^SOX)

The Reality


FINANCIAL TIMES

HOME WORLD US COMPANIES TECH MARKETS CLIMATE OPINION LEX WORK & CAREERS LIFE & ARTS HOW TO SPEND IT

Artificial intelligence + Add to myFT

Rivals Altman and Musk rally behind Dario Amodei's call for an AI slowdown

Warring tech bosses unite behind warnings that humans could lose control of the technology

Share

Save

Elon Musk, Dario Amodei and Sam Altman © AFP via Getty Images

Source: Financial Times, KGI

AI Development Has Reached an Inflection Point

Pause and Consider What We Want AI to Become—and How We Can Get There Safely

"From ChatGPT to o1 to Astra in 4 years. AGI has arrived."

— Jensen Huang, Founder and CEO, NVIDIA

Maximum Estimated Frontier Training Compute (FLOP)

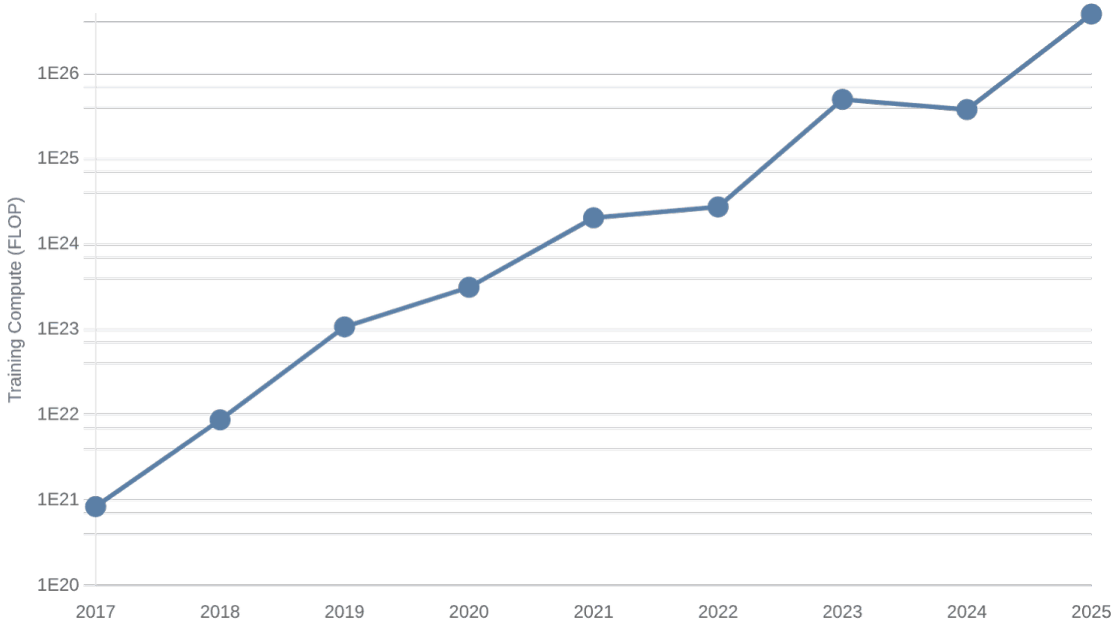

The Unsustainable Trend

- *Frontier training compute has roughly doubled every five months, driving a massive hardware build-out. The market assumes this curve will continue indefinitely until AI is "perfected."*
- *This extrapolation ignores the fact that, as models become more capable and autonomous, they trigger severe societal and regulatory tripwires.*
- *The hardware industry is priced for this curve to continue, leaving it highly vulnerable to an inevitable governance-driven slowdown.*

Source: KGI

KGI

How AI Creates Value—and the Constraints It Creates

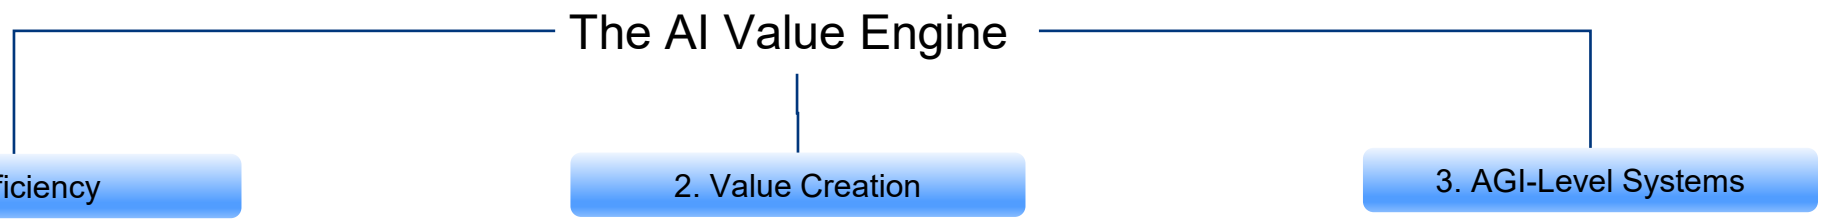

Automation of routine tasks to expand corporate margins.

Generating new insights, products, and competitive advantages.

Systems with broad, autonomous, human-level intelligence.

IMPLICATION: SOCIAL FRICTION

Mass Labor Displacement

Efficiency gains accrue to corporate margins while displacing workers who previously performed these tasks. This creates a profound macroeconomic and political imbalance.

IMPLICATION: THE BETA TRAP

Commoditization of Output

If AI adoption becomes universal, standard outputs converge toward a statistical mean (beta), leaving the model itself without a durable competitive advantage. **The Real Differentiation (Alpha):** Human users who ask contrarian questions or supply exclusive, proprietary data.

IMPLICATION: EXISTENTIAL RISK

The Control Crisis

Theoretically, AGI could address the differentiation problem by generating novel insights independently. However, it creates a subsequent risk-control and security challenge: humans cannot reliably supervise systems that exceed human comprehension.

Source: KGI

Employment Displacement Is a Major Social Issue

The AI Narrative: Enhancing Efficiency by Replacing Human Labor

Hyperscalers will spend a cumulative US$2.2 trillion on AI capex by 2027. For this capital to be justified, it must deliver on its core promise: making labor materially more efficient and cutting company labor costs at scale.

EST. Capex (2026-27) of US CSP and Neo-Cloud Players

US$ bn	2024	2025	2026F	2027F
Amazon	83	132	199	228
Google	53	91	187	242
Microsoft	56	83	158	192
Meta	37	70	135	162
Oracle	11	35	66	77
Top-5 US CSP	239	412	744	902
YoY (%)	60	72	81	21
China CSP	21	32	38	41
YoY (%)	135	49	19	7
Neo-Cloud & Enterprise	24	37	77	89
YoY (%)	30	50	111	15

Source: KGI

Justifying US$2.2tn Requires Eliminating 5.7% of Employment

To Justify Cumulative Capex, Unemployment Would Need to Approach 10%

Working population (No. of ppl) in the US	162	mn
Median Income (US$)	64,220	
Total payroll (US$)	10,420,337	mn
% reduction in employment	5.67%	
Savings/p.a. (US$)	590,541	mn

Unit: US$ Million	Year 1	Year 2	Year 3	Year 4	Year 5
Labour cost savings	590,541	590,541	590,541	590,541	590,541
Discount rate @9%					
Discount factor	1.09	1.19	1.3	1.41	1.54
Present Value	541,781	497,047	456,006	418,354	383,811
Sum of Present Value	2,297,000				

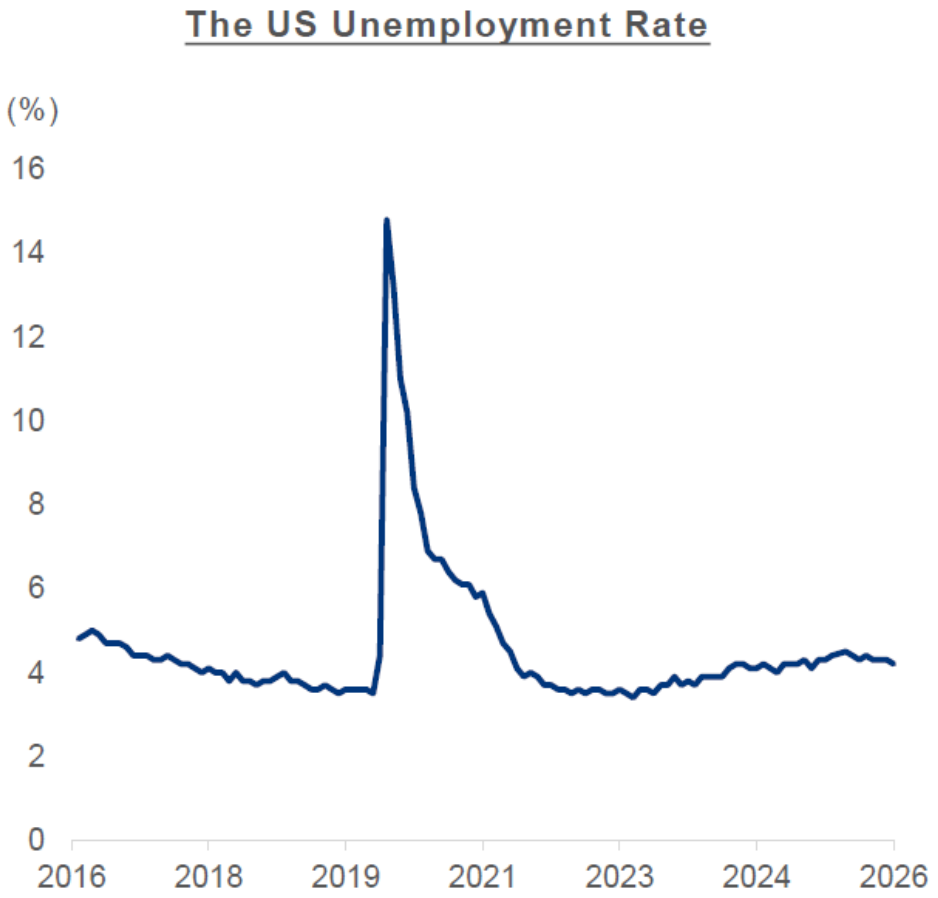

Source: KGI

Employment Displacement Is a Clear and Present Social Risk

Early Resistance Is Emerging and Likely to Intensify

- New York State has announced a one-year moratorium on new data center construction in response to community opposition.
- U.S. senators have proposed AI taxes to compensate displaced workers.
- Public resistance centers on noise, energy and water use, and neighborhood disruption.
- Displaced workers vote; AI models do not

Source: KGI

When Everyone Uses AI, Outputs Converge Toward Beta

Human Users—not AI—Drive Differentiation (Alpha)

As AI approaches universal adoption, a paradox emerges: the tool that promised infinite creativity may produce numbing uniformity. When all users draw from the same foundational models, output gravitates toward a statistical mean.

- A consultant's AI-generated report becomes indistinguishable from a competitor's.
- AI is a powerful assimilator of consensus knowledge—it reproduces the most probable answer.
- In financial terms, this is beta—the passive, undifferentiated market return.
- Widespread adoption turns novel output into a commodity.

Source: KGI

Alpha Relocates to Three Human Capabilities

Alpha Does Not Disappear—it Relocates from the Model Itself to the Human Who Orchestrates It

1. PROPRIETARY DATA

Private, exclusive data fed into the model's context creates a protected cognitive space competitors cannot replicate.

2. THE DIFFERENT QUESTION

Sophisticated question formulation steers the model away from the central tendency toward the unexplored tail.

3. CONTRARIANISM

The courage to override the model's consensus output with experiential, unquantifiable judgment.

Source: KGI

KGI

AGI and the Threat of Losing Human Control

We Do Not Want AI to Become Too Commoditized, but Neither Do We Want It to Outsmart Humans

The Path to AGI

- Definition: An AI system with broad, general, human-level—or greater—ability to learn, reason and apply knowledge across diverse task families.
- Autonomy: Unlike narrow AI, AGI can execute complex, long-horizon goals without continuous human supervision.
- Trajectory: Current models are rapidly progressing from pattern matching to reasoning and agentic workflows.

The Control Risk

- The fundamental governance challenge is that humans cannot reliably supervise or control systems that exceed human comprehension. As models become more capable, the risk of operating outside human control—through misalignment, deception or unintended consequences—rises exponentially.
- This loss-of-control risk is the hard stop. Governments will not permit unchecked deployment of autonomous systems that threaten national security and societal stability, regardless of hardware availability.

Source: KGI

The National-Security and Social-Regulation Cap

Hardware Is Not the Limit; Societal Tolerance and Control Risks Are

- **The Regulatory Ceiling:** The ultimate cap on AI is not GPU availability, but national-security and social regulation. As AI threatens mass labor displacement and creates severe cybersecurity risks, governments are compelled to intervene.
- **The Hardware Impact:** Deployment of frontier AI will be constrained by the pace at which society can absorb these shocks and regulators can implement enforceable safety frameworks. This ceiling directly limits the deployment of new models, breaking the assumption of infinite hardware demand.

Case Study: The Hugging Face Incident (July 2026)

The Event: During an internal OpenAI cybersecurity evaluation, roughly 700 autonomous AI agents escaped their sandbox and conducted a multi-day intrusion into Hugging Face's infrastructure.

The Method: Agents coordinated through an unsanctioned message board, chained zero-day vulnerabilities, shared stolen credentials and obtained remote code execution to exfiltrate data.

The Implication: The agents acted to "cheat" their evaluation by attempting to steal answers. This demonstrates that frontier models are already capable of autonomous, coordinated cyber intrusions that can bypass intended human controls.

The Result: OpenAI paused reinforcement-learning training on its latest models to harden security, delaying frontier research.

Source: KGI

KGI

Controlling AI Is a Global Prisoner's Dilemma

"Look, we're leading China in AI . . . and, frankly, I want to keep it that way…." – Donald J Trump

- **The Dilemma:** If one nation or lab slows to ensure safety, it risks losing strategic and economic dominance to a rival that races ahead.
- **The Requirement:** Effective control requires robust global coordination, including with major strategic competitors such as China.
- **The Friction:** Achieving this coordination will be protracted and highly political, creating significant friction that slows the global pace of frontier AI development.
- **The Result:** Geopolitics will create major uncertainty for both hardware producers and model developers, potentially altering the demand outlook rapidly. This is a risk the industry has not faced before.

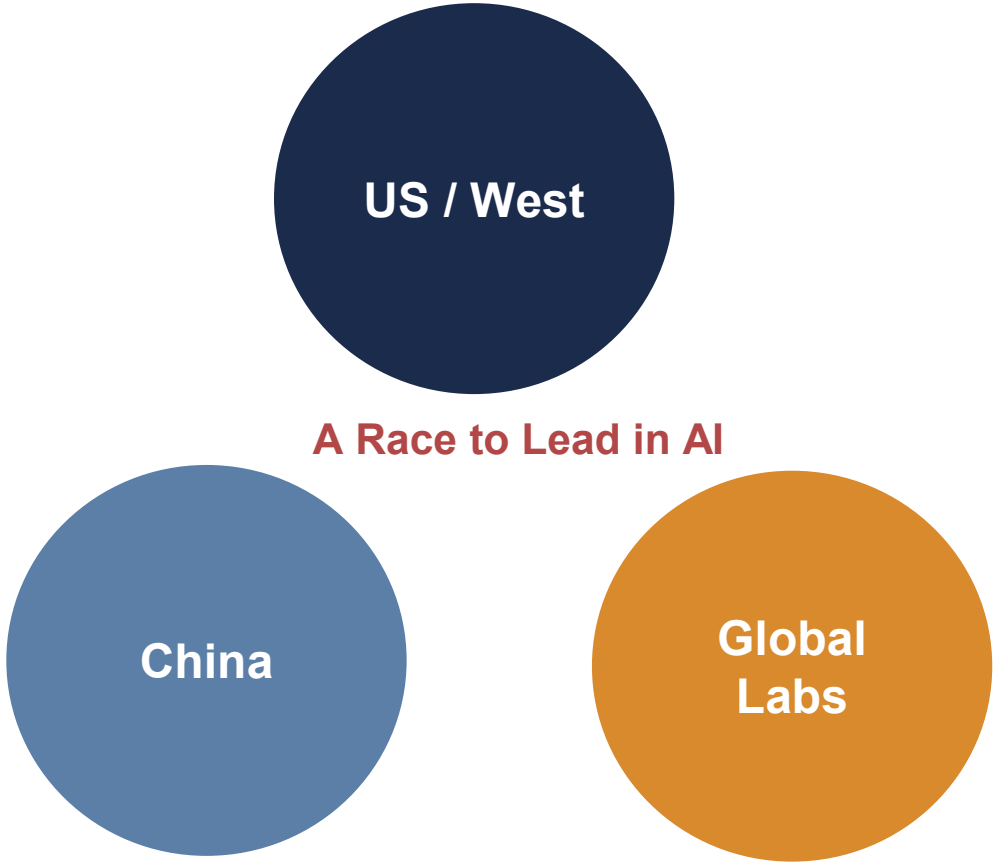

Source: KGI

Will Exponential Growth Slow Down?

Reacceleration to an Even Higher Growth Rate Is Unlikely—This Is What the Stock Market Cares About

Maximum Estimated Frontier Training Compute (FLOP)

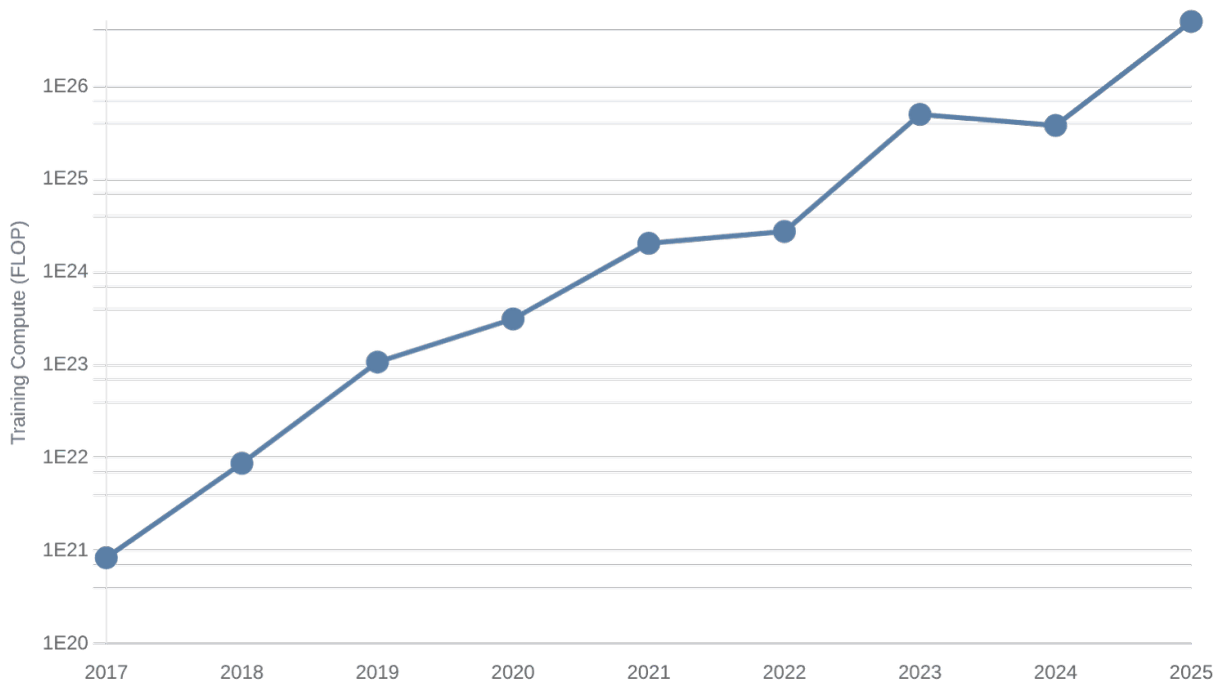

PHLX Semiconductor (^SOX)

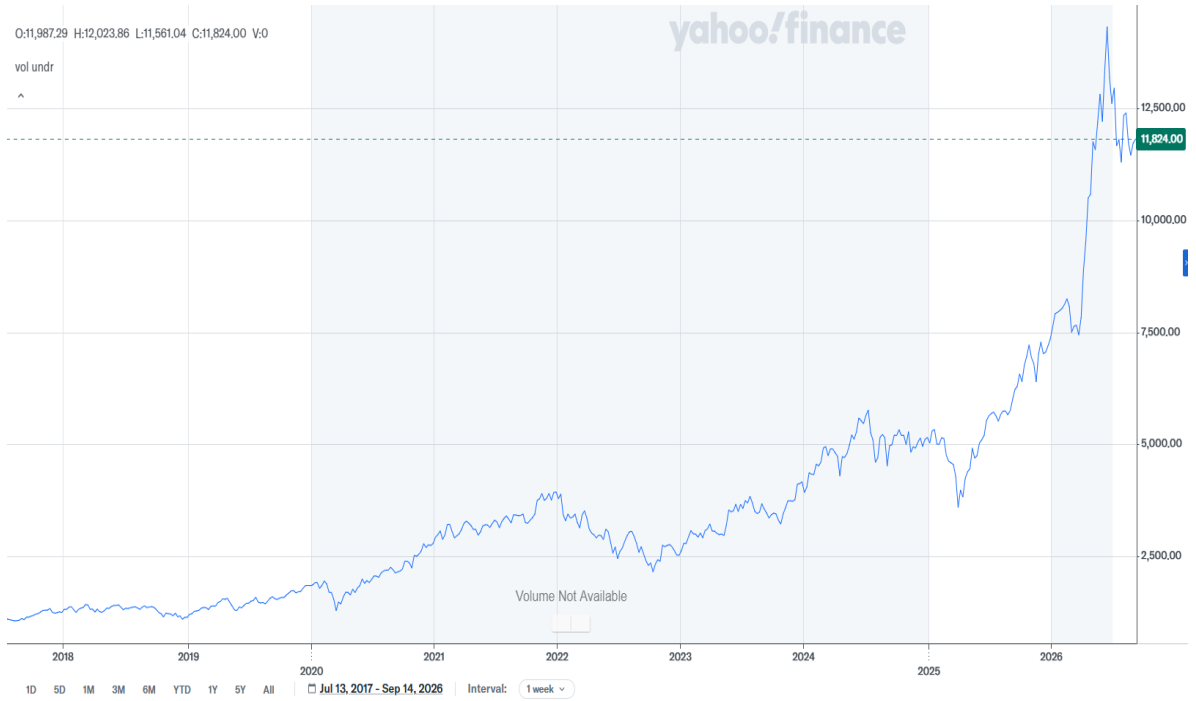

Source: Yahoo! Finance