

RACE WEEKLY INTELLIGENCE BRIEF

RETINA


AI · CYBERSECURITY · CLOUD · LEADERSHIP

THIS WEEK'S COVER:

THE \$600 BILLION QUESTION

In January 2025, a cheap Chinese model erased \$589 billion from Nvidia in a single day, the largest one-day loss in market history and the whole industry flinched. Big Tech will still pour roughly \$600 billion into AI this year, while 95% of corporate AI pilots show no profit at all. "Is it a bubble?" is the question every board is asking, and the wrong one. The right one is quieter: if the buildout wobbles, how much of what you now depend on wobbles with it?

This issue answers it, with a little help from the railways, the fibre-optic glut, and a sock puppet.



ALSO INSIDE

The World in 5 · Signal · Hyperfocal · Echoes · Perspectives · The Bubble-Proofing Playbook · Metric of the Week · Fun Fact · CANDID · Book Shelf · Research Radar · RACE Pulse

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Echoes
- Perspectives
- The Bubble-Proofing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

FROM THE EDITOR'S DESK

For eighteen issues we've written about what AI can do. This week: whether the people paying for it are about to find out the hard way.

On 27 January 2025, a little-known Chinese lab called DeepSeek released a model that matched the American frontier at a fraction of the training cost. By Monday's close, Nvidia had lost \$589 billion in market value — the biggest single-day wipeout any company has ever suffered, larger than the entire worth of all but a handful of US firms. The thing that spooked the market wasn't that AI didn't work. It was the opposite: that AI might work for far less money than the half-trillion-dollar buildout assumed. A whole thesis blinked. Then, mostly, it carried on.

That carry-on is the story of 2026. The five biggest hyperscalers are on track to spend somewhere around \$600 billion on AI infrastructure this year (CreditSights puts the top five near \$602 billion; J.P. Morgan's tally of five names is nearer \$697 billion), and Goldman Sachs reckons global AI investment tops \$1 trillion. Against that, the returns are conspicuously thin: MIT's Project NANDA found that 95% of enterprise generative-AI pilots produce no measurable profit-and-loss impact at all. Spending like a certainty; earning like an experiment. That gap is what the word "bubble" is pointing at.

But here's the trap in the question. "Are we in a bubble?" invites a yes/no about the stock market — which almost no CXO can act on and almost none should try to. There are really two bubbles stacked on top of each other: an infrastructure bubble (hyperscaler capex, circular financing, valuations richer than the 2000 telecom peak) and a much healthier application layer where real customers pay real money for things that work. Confusing the two is the costliest mistake in the whole debate. Your job isn't to call the top. It's to know which bubble your own spending, and your own vendors, are sitting in — and what happens to you if the first one deflates while the second one keeps compounding. History, as this issue shows, is unusually generous with clues.

THE WORLD IN 5

GLOBAL HEADLINES A CXO SHOULD BE ABLE TO REFERENCE

- From the Editor's Desk
- **The World in 5**
- Signal
- Hyperfocal
- Echoes
- Perspectives
- The Bubble-Proofing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

Five signals from the past fortnight, worth a note at your next leadership table.

1. The Spend: A Trillion-Dollar Bet, Mostly On Faith

Global AI investment will cross \$1 trillion in 2026 (Goldman Sachs), with the top five US hyperscalers alone near \$600–700 billion. Capital intensity has hit levels last seen in utilities and telecoms, Oracle and Microsoft now spend 45–57% of revenue on capex. The money is real and committed. Whether the demand that justifies it is equally real is the entire argument.

2. The Gap: Revenue Isn't Keeping Up

Sequoia's David Cahn framed it as "AI's \$600 billion question": the annual revenue the industry must generate to justify the build and in 2026 the gap is widening, not closing. Cloud AI revenue is growing fast (Azure AI reportedly +123% year on year), but from a base far below the spend. Capex is running roughly 46% ahead of sales growth; in the 2001 telecom bust, that divergence was 32%.

3. The Pilots: 95% Show Nothing

MIT's Project NANDA studied enterprise generative-AI deployments and found 95% produced no measurable P&L impact, on tens of billions of corporate spending. The failure is rarely the model; it's deploying a tool in search of a problem. This is the single most important number for a buyer, because unlike the market, it is entirely within your control.

4. The Plumbing: Money Going In Circles

A web of deals now links chipmaker, cloud and lab, Nvidia investing in the customers who buy its chips, take-or-pay compute commitments, GPU-collateralised debt. It manufactures a demand signal that looks independent but isn't, echoing the vendor financing that inflated Lucent and Nortel before 2001. Analysts also flag that depreciating GPUs over five to six years, when their useful life may be two to three, flatters reported earnings across the group.

5. India: Exposed As A Buyer, Not A Builder

India isn't funding this buildout, it's consuming it, through cloud contracts, model subscriptions and SaaS priced in dollars. That's a subtler exposure: not stock-market risk, but dependency risk. If the infrastructure layer reprices, Indian enterprises inherit the new prices without ever having placed the bet. The hedge is the discipline India already knows well, frugal, ROI-first adoption.

Source:

<https://know.creditsights.com/insights/technology-hyperscaler-capex-2026-estimates>
<https://www.jpmorgan.com/insights/banking/capital-markets/financing-ai-infrastructure-data-centers>
<https://am.jpmorgan.com/se/en/asset-management/per/insights/market-insights/investment-outlook/technology-and-ai/>
<https://sequoiacap.com/article/ais-600b-question>
https://www.allianz-trade.com/en_global/news-insights/economic-insights/AI-capex-cycle-war-proof-now.html
<https://www.goldmansachs.com/insights/articles/what-to-expect-from-ai-in-2026-personal-agents-mega-alliances>

 SIGNAL

THE BUBBLE — BY THE NUMBERS

- From the Editor's Desk
- The World in 5
- **Signal**
- Hyperfocal
- Echoes
- Perspectives
- The Bubble-Proofing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

~\$600B

AI capex by the top five US hyperscalers in 2026 — roughly triple 2024 (CreditSights / J.P. Morgan)

~\$1 tn

total global AI investment forecast for 2026 (Goldman Sachs)

95%

of enterprise generative-AI pilots show no measurable P&L impact (MIT Project NANDA)

\$589B

erased from Nvidia in one day when DeepSeek launched — the largest single-day loss in market history

~46%

gap between capex growth and revenue growth in 2026, vs 32% at the 2001 telecom peak (Allianz)

7.9%

the Magnificent 7's free cash flow as a share of sales — its lowest since early 2024, as capex surges

Source:

<https://know.creditsights.com/insights/technology-hyperscaler-capex-2026-estimates/>


<https://www.goldmansachs.com/insights/articles/what-to-expect-from-ai-in-2026-personal-agents-mega-alliances>

<https://mitsloan.mit.edu/shared/ods/documents?PublicationDocumentID=10819>

<https://www.bloomberg.com/news/articles/2025-01-27/asml-sinks-as-china-ai-startup-triggers-panic-in-tech-stocks>

https://www.allianz-trade.com/en_global/news-insights/economic-insights/AI-capex-cycle-war-proof-now.html

<https://am.jpmorgan.com/us/en/asset-management/adv/insights/market-themes/artificial-intelligence/>





HYPERFOCAL

WHAT BROKE THIS WEEK

- From the Editor's Desk
- The World in 5
- Signal
- **Hyperfocal**
- Echoes
- Perspectives
- The Bubble-Proofing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

There Are Two Bubbles. You Only Need to Worry About One.

Picture the AI economy as two floors of the same building. The upper floor is infrastructure: the hyperscalers, the chipmakers, the labs, the data centres, the \$600-billion capex and the valuations richer than the 2000 telecom peak. The lower floor is application: ordinary companies using AI to draft, code, forecast, support and sell, some of it paying for itself handsomely. When people ask “is AI a bubble,” they blur the two floors together. They are not the same building risk. The upper floor could take a serious fall while the lower floor barely notices, exactly as, after 2000, hundreds of dot-coms vanished while e-commerce quietly kept growing every single year.

For a CXO, this reframes everything. You are almost certainly not an investor in the upper floor; you are a tenant on the lower one. Which means the market's bubble is not your primary risk. Your primary risk is two things the headlines rarely mention. First: are your own AI projects in the 95% that show no profit? MIT's NANDA number is not a story about AI failing, it's a story about organisations buying a

model and then hunting for a use, spending real money to produce a demo that never becomes a P&L line. That failure is entirely yours to prevent, and it has nothing to do with Nvidia's share price.

Second, and more subtly: dependency. The tools you now run your business on a copilot in every seat, a model behind your support desk, a vendor whose whole roadmap assumes cheap frontier compute- are priced today on the assumption that the buildout continues and capital stays cheap. If the upper floor corrects, some of that repricing lands on you: subscriptions rise, a favourite startup folds, a “free” tier disappears. Circular financing makes this worse, because when one node stumbles the shock travels through all the others at once. The enterprises that sail through won't be the ones who called the top. They'll be the ones who knew, in advance, which of their dependencies were load-bearing — and had somewhere else to stand.

RETINA TAKE

“Are we in a bubble?” is a spectator's question. “Which of our AI spend pays for itself, and what breaks if our vendors reprice?” is an operator's. Only one of them is answerable, and it's the one that keeps you standing whichever way the market goes.

ECHOES

What Three Bubbles Left Behind

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- **Echoes**
- Perspectives
- The Bubble-Proofing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

A new segment for this issue. We usually run Panorama's three forces here; this week the analysis is better told as history, because we have done a version of this at least three times before – and the pattern is remarkably kind to anyone who reads it.

1840s · Britain

RAILWAY MANIA

In the 1840s, Britain went mad for railways. Parliament approved thousands of miles of track; over 7% of GDP poured into it; the network nearly quintupled to 7,300 miles between 1840 and 1852. The mathematician Andrew Odlyzko, who studies these episodes, calls them “collective hallucinations”, moments when investors, press and public stop seeing risk together. The lines, when finished, earned about a quarter of the revenue their backers had promised. Fortunes evaporated.

The lesson: The trains still ran. The mania killed the investors, not the railways – which were bought cheap by second owners and carried Britain for a century. Being right about the technology and right about the share price were two entirely different bets.

1998–2002 · Telecom

THE FIBRE-OPTIC GLUT

Convinced the internet would need infinite bandwidth, telecom firms buried 80–90 million miles of fibre and spent something like \$2 trillion doing it. By 2001, an estimated 95% of that fibre was “dark” never lit, never used. Bandwidth prices fell up to 90%. WorldCom collapsed into the largest bankruptcy in US history to that point; Global Crossing and others followed; roughly \$2 trillion in market value was destroyed.

The lesson: That dark fibre became the backbone of the broadband, streaming and cloud era bought for cents on the dollar by the companies that came next. The overbuild was catastrophic for its financiers and a gift to everyone who bought its output afterwards.

2000 · Dot-Com

THE SOCK PUPPET AND THE INDEX

Tech stocks rose about 800% from 1995 to March 2000, then fell more than 80%. Pets.com, mascot, Super Bowl ad and all – went from IPO in February 2000 to liquidation by November, roughly nine months. Hundreds of companies with a domain name and a story simply disappeared.

The lesson: E-commerce, the actual idea, never stopped growing not in 2000, not in 2001, not once. The froth was in the valuations and the business models, not the technology. The application layer survived the infrastructure crash without blinking.

The Pattern

Three times, the same shape: the technology was real, the capital was early, the assets outlived the investors, and the next wave's winners bought the last wave's overbuild cheap. AI's data centres and GPUs are this cycle's rails and fibre. Which raises the only question that matters for a buyer – in this story, are you the one financing the overbuild, or the one who benefits from it later?



PERSPECTIVES

DEEP TECH — ANALYSIS BY THE RACE TEAM

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Echoes
- **Perspectives**
- The Bubble-Proofing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs



Dr. Shinu Abhi

Professor and Director
Corporate Training, RACE,
REVA University

<https://www.linkedin.com/in/shinu-abhi-phd-711a339/>

Almost every week, a senior leader asks us to “build them an AI.” I’ve learned to listen for one thing in the first five minutes, because it predicts almost everything that follows. The ones whose projects fail begin with the technology: they’ve seen a demo, a competitor’s press release, a board member forwarded a video, and they want that. The ones whose projects work begin with a number, a cost they want down, a cycle time they want halved, a conversion rate they want to move and only then ask whether AI is the right tool for it. Same budget, opposite outcomes. MIT’s 95% figure did not surprise

me at all. I watch it get made.

The bubble talk, from where I sit, is mostly aimed at the wrong floor of the building. The enterprises losing money on AI are not losing it because the market is frothy or because Nvidia is expensive. They are losing it because they funded a capability and forgot to fund a problem, no baseline measured before, no metric agreed, no owner who has to answer for the P&L after. That is not a technology risk you hedge by watching the Nasdaq. It is a management discipline, and it is completely independent of whether the analysts are right about a correction.

There is also, quietly, an opportunity in all this for Indian organisations, and it is the one history keeps pointing at. We are buyers in this cycle, not the ones financing the overbuild and DeepSeek showed the whole world that capability is getting dramatically cheaper, not more expensive. The disciplined

move is not to place a huge frontier bet; it is to build the muscle to get real value from the cheapest model that clears the bar, and to keep dry powder for the day the overbuilt capacity is on sale. We teach frugal AI not because it is modest, but because in a cycle like this one, frugality is the strategy that survives whichever way the market breaks.

RACE Perspective

The bubble will sort itself out on a timeline no CXO controls. What you control is whether your own AI spend is anchored to a number you’re already tracking. Fund problems, not models; default to the cheapest tool that works; and keep some powder dry for the sale that history says is coming.

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Echoes
- Perspectives
- **The Bubble-Proofing Playbook**
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs



THE BUBBLE-PROOFING PLAYBOOK

SPENDING ON AI WITHOUT
BETTING ON THE BUBBLE

Five moves, ordered from fastest to most strategic, that any leadership team can act on.

MOVE 1

JUDGE YOUR SPEND BY YOUR P&L, NOT THE MARKET'S MOOD

The bubble debate is about the upper floor — hyperscaler valuations. Your AI projects live on the lower floor and should be judged there. A tool that pays for itself doesn't stop paying because the Nasdaq falls, and a project that shows nothing isn't rescued because the market rises. Separate the two questions completely.

The Move: Stop asking 'is AI a bubble.' Ask 'does this specific deployment beat its own baseline' — and let that, not market sentiment, gate the budget.

MOVE 2

ESCAPE THE 95%: FUND A PROBLEM, NOT A MODEL

Every project in MIT's failed 95% has the same origin story — a capability bought before a problem was named. Invert it. Start with a costed business metric, measure the baseline before you build, and put a date on when the project must move that number or be killed. This one discipline separates the 5% from the rest more reliably than any tool choice.

The Move: No AI budget without a named metric, a measured baseline, and a kill date. A pilot that can't show P&L by its date is a sunk cost, not a strategy.

MOVE 3

MAP YOUR LOAD-BEARING DEPENDENCIES

If a key vendor doubled its price, lost its funding, or retired a model you rely on, what would break? Because of circular financing, a stumble in one part of the AI supply chain can hit several of your tools at once. Know which of your AI dependencies are load-bearing, and make sure each has a substitute or an exit before you need one.

The Move: Build an AI dependency map with an exit path for every critical one. Prefer portable and open options where the switching cost of a reprice would otherwise be brutal.

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Echoes
- Perspectives
- **The Bubble-Proofing Playbook**
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs



THE BUBBLE-PROOFING PLAYBOOK

SPENDING ON AI WITHOUT
BETTING ON THE BUBBLE

MOVE 4

KEEP DRY POWDER FOR THE SALE THAT HISTORY PREDICTS

Railways, fibre, dot-com capacity — each overbuild became a fire-sale bargain for the wave that followed. AI compute and tooling are unlikely to be different. The organisations that gained most from past manias weren't the ones who bet biggest at the peak; they were the ones with cash and a shortlist ready when prices fell.

The Move: Don't commit every rupee at the top of the cycle. Keep a shortlist and a reserve, so a correction is a buying opportunity for you rather than a crisis.

MOVE 5

MAKE FRUGAL THE DEFAULT, NOT THE FALLBACK

DeepSeek's lesson wasn't geopolitical — it was economic: capability is getting cheaper fast, and paying frontier prices for a routine task is now a choice, often a poor one. Default to the smallest, cheapest model that clears your quality bar, and treat every frontier call as something you must justify, not assume. Frugality isn't caution here; it's the posture that wins across the whole cycle.

The Move: Set the cheapest model that passes as your default, and require a reason to spend up. In a cycle this uncertain, low-cost capability is the most durable bet on the board.

Source:

<https://mitsloan.mit.edu/shared/ods/documents?PublicationDocumentID=10819>

<https://sequoiacap.com/article/ais-600b-question>

<https://sequoiacap.com/article/ai-optimism-vs-ai-arms-race>

https://www.allianz-trade.com/en_global/news-insights/economic-insights/AI-capex-cycle-war-proof-now.html

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Echoes
- Perspectives
- The Bubble-Proofing Playbook
- **Metric Of The Week**
- **Fun Fact**
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

METRIC OF THE WEEK

ONE NUMBER, PROPERLY UNDERSTOOD

\$600 billion

the annual revenue AI must generate to justify the infrastructure being built – and the gap is widening

What it is: Sequoia partner David Cahn's "AI's \$600 Billion Question", an estimate of the yearly end-customer revenue the industry needs to earn to pay back the compute it is buying. He first posed it at \$200 billion in 2023, raised it to \$600 billion in 2024, and in 2026 the shortfall between that number and actual AI revenue is growing, not shrinking.

Why it matters now: It converts a vague

worry into arithmetic. Infrastructure is being built for a level of AI demand that has not yet shown up in revenue. That doesn't mean it never will, e-commerce eventually dwarfed the dot-com build but it does mean today's spending is a bet on tomorrow's adoption, and the size of the bet is now measured in hundreds of billions a year.

How leaders use it: Treat it as a weather report, not a verdict. It won't tell you when a

correction comes, but it tells you the pressure is building which is the cue to do the unglamorous things in this issue's playbook now, while capacity is plentiful and your vendors are still competing for you.

Source:
<https://sequoiacap.com/article/ais-600b-question>

FUN FACT

SOMETHING TO REPEAT AT DINNER

Corning, the company that makes the glass fibre the internet runs on was one of the most essential firms of the dot-com era. Its stock went from about \$12 in 1995 to over \$100 in October 2000, then to roughly \$2 by October 2002. It lost about 98% of its peak value. Corning was not a fraud, not a fad, and not wrong about the future: fibre really did become the nervous system of the modern world. It was simply, for a while, priced as though that future had already arrived. Being essential to a revolution and being a good investment at the top of a mania are, it turns out, completely different things.

Source: <https://sequoiacap.com/article/ais-600b-question> | <https://stockpricehistory.com/quotes/GLW/history/2002-10-01>



CANDID



Biswajit Mishra

Cognitive Business, SoundHound AI

<https://www.linkedin.com/in/biswajit-mishra-21a14820/>

Q1. Big Tech will spend around \$600 billion on AI infrastructure this year, while 95% of enterprise pilots show no profit. How should a board tell the difference between an AI investment that is merely early and one that simply isn't working?

One of the key lessons from past technological disruptions from electricity to IT, is the J-curve effect: technological capability often arrives years before measurable enterprise productivity. AI innovation projects may therefore lose money initially and still be strategically successful if they generate valuable, actionable insights. These include whether the organisation is too early in the adoption cycle or ready for rapid scaling; how to reduce the unit cost of experimentation so that each iteration becomes faster and cheaper; and what non-linear outcomes, operating models and organisational redesign may be required.

Conversely, a project may demonstrate a small positive ROI yet remain strategically irrelevant if its economics cannot scale. AI investment should therefore combine a Skunk Works operating philosophy small, capable and autonomous teams with the board-level mindset of a corporate venture capitalist:

CXO VOICE – IN CONVERSATION THIS WEEK

Small exploration cost → actionable insight → strategic option → selective scaling

Rather than: Large investment → simulated business case → optimistic forecast → hoped-for ROI

Q2. For most enterprises, the real bubble risk isn't their share portfolio; it's the vendors and tools they now depend on. How should a CXO be thinking about that dependency before, rather than after, any correction?

Vendor lock-in and excessive dependency erode enterprise agility, resilience and competitive advantage—especially during uncertainty. Enterprises should retain ownership and control of their data, business logic, workflows, institutional knowledge, prompts and business context. Models, GPUs, inference capacity and specialised AI services should follow a “rent and substitute” approach.

The shared priority of procurement and business teams must be to preserve optionality, the ability to switch providers without disrupting core operations. This is equally critical for disaster recovery, enabling enterprises to move between model and inference providers during service outages.

Q3. Every past infrastructure mania- railways, fibre-optic cable- left behind assets the next wave bought cheaply. Is there a version of that second-owner opportunity in AI, and how would you position an organisation to capture it?

In every innovation cycle, pioneers fund discovery and reveal what works,

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Echoes
- Perspectives
- The Bubble-Proofing Playbook
- Metric Of The Week
- Fun Fact
- **Candid**
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs



CANDID

CXO VOICE – IN CONVERSATION THIS WEEK

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Echoes
- Perspectives
- The Bubble-Proofing Playbook
- Metric Of The Week
- Fun Fact
- **Candid**
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

while capable later entrants commercialise those lessons. As with railways and fibre-optic networks, later entrants in AI can benefit from infrastructure built by others, provided they invest in the capabilities required to convert it into value.

These capabilities span people, processes and organisational design across the entire chip-to-customer value chain. Organisations must prepare to exploit emerging AI assets as owners, operators and customers by:

- Developing context-aware AI talent across hardware, infrastructure and applications
- Accelerating use-case discovery, experimentation and deployment
- Strengthening data management and governance at scale
- Building API-first, provider-neutral architectures
- Nurturing a strong AI-native partner ecosystem

Above all, organisations must preserve the financial capacity to act when opportunities emerge, while running small, autonomous, organisation-wide Skunk Works experiments to capture insights, build capability and prepare for selective scaling.

Q4. Cheaper models like DeepSeek keep undercutting frontier costs. Does that strengthen the case for waiting until prices fall further, or for moving now with frugal tools, and how do you tell the two apart?

Waiting to adopt AI risks falling behind the AI Darwinian curve as capabilities advance and inference costs decline. Organisations must put

“skin in the game” through frugal pilots using increasingly capable open-source models.

The guiding principle should be: Learn now; commit later—with scalability and optionality built in.

While inference is becoming cheaper, organisational learning is becoming more demanding. It requires new technical capabilities, workflow redesign and sustained change management.

A practical response is an AI barbell strategy:

- Run numerous small, inexpensive experiments to accelerate learning.
- Make substantial investments only in demonstrated winners with strong scaling potential.
- Build provider optionality into the commercial AI stack from the outset.

Note: CANDID features direct contributions from CXO leaders.
To contribute, contact: retina@reva.edu.in

BOOK SHELF

RETINA'S READ OF THE WEEK

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Echoes
- Perspectives
- The Bubble-Proofing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- **Book Shelf**
- Research Radar
- RACE Pulse
- RACE Programs

RETINA RECOMMENDS · THIS WEEK'S THEME

Devil Take the Hindmost: A History of Financial Speculation


Edward Chancellor · Farrar, Straus and Giroux, 1999

Chancellor's history reads like a gallery of the anecdotes in this issue's Echoes segment, told at full length, the railway mania, the South Sea Bubble, the 1920s, each with the same cast of hallucinating investors and the same morning-after. Published in 1999, on the very eve of the dot-com peak, it became an unintentional prophecy; read in 2026, its patterns are almost uncomfortably familiar. The book's quiet argument is that speculation is not an aberration in capitalism but a recurring feature of it, and that the details rhyme far more than anyone in the middle of a mania wants to admit.

What it offers a leadership team is not a market call, Chancellor is careful never to pretend he can time a top. It is pattern

recognition: the tells that a boom has detached from its fundamentals, the specific ways financing structures hide risk, and the reliable human machinery of "this time is different." A CXO who has read it will listen to an AI vendor's growth story with a more useful kind of ear.

Pair it with this week's cover and the takeaway sharpens. Chancellor's manias almost always sat on top of a genuine innovation, the technology was usually real, and often transformative. The money was simply too early and too sure. That is precisely the distinction between AI's two floors. Read it before your next AI investment committee, not for pessimism, but for the calm that comes from having seen the movie before.



- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Echoes
- Perspectives
- The Bubble-Proofing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- **Research Radar**
- RACE Pulse
- RACE Programs



RESEARCH RADAR

RACE PUBLICATIONS THIS WEEK

EDITOR'S NOTE

Two new publications from RACE — one on making cybersecurity operations faster and cheaper to run, one on turning supply-chain uncertainty into a number a planner can act on. Standard format: a concise technology summary followed by clear enterprise relevance.

Cybersecurity · Incident Response · Applied LLMs

AI-Powered Incident Report and Response System

Nirupam Ghosh; Rashmi Agarwal; Anupam Ghosh; Nishanth Krishna

Modern IT estates drown analysts in unstructured log data, and manual triage is slow and error-prone. Rather than bolt on one more isolated tool, this work combines three that usually run apart — Random Forest for severity prediction, BERT for classification, and an LLM for contextual remediation advice — into a single pipeline. Tested on logs from operating systems, applications, firewalls and synthetic sources, it reached 89% accuracy in category classification and 82% in severity prediction, with real-time ingestion via REST APIs and independently scalable components. The reported payoff is lower analyst workload and a shorter Mean Time to Response.

Enterprise Relevance: This is the ROI cut of security that fits the week's cover, value measured in analyst hours saved and MTTR reduced, not tools bought. For a CISO, it shows how combining cheap, proven models beats waiting for a single frontier system, and how to put a number on a security investment a board will actually approve.

Innovative Computing and Communications (ICICC 2026) · Lecture Notes in Networks and Systems, vol 2021, pp 593–604 · Springer
https://link.springer.com/chapter/10.1007/978-3-032-28310-8_45

Supply Chain · Predictive Analytics · Healthcare Manufacturing

Order Visibility: Supply Chain Order Milestone Prediction

Shivpriya V; Krishnakumar Tiwari; Rashmi Agarwal

Off-the-shelf ERP and planning suites are hard to customise and slow to react, costly in healthcare manufacturing, where supply volatility hits production, inventory and customer trust. This work builds a machine-learning model to predict the delivery date of a new made-to-order medical device using deliberately few, explainable features, so it stays scalable and doesn't need constant reconciliation against heavyweight enterprise planning systems.

Enterprise Relevance: The discipline here is exactly what this issue's playbook argues for a lean model tied to a concrete business metric (on-time delivery) that pays for itself, rather than a sprawling platform bought on faith. For operations and supply-chain leaders, explainable, low-feature prediction is often the higher-ROI path than a full optimisation-suite overhaul.

2024 IEEE 3rd World Conference on Applied Intelligence and Computing (AIC), Gwalior · IEEE
<https://ieeexplore.ieee.org/document/10730951>

Explore the full RACE research archive at <https://raceresearch.reva.edu.in>. For collaboration, licensing or research-partnership enquiries: race@reva.edu.in

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Echoes
- Perspectives
- The Bubble-Proofing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- **RACE Pulse**
- RACE Programs

RACE PULSE

EVENTS · COMMUNITY · WHAT'S COMING

RACEx360 SEASON 2

25 SEPTEMBER 2026 · TAJ WEST END, BENGALURU

One day, one stage, four frontiers.

REVA Academy's Emerging Technology Conference returns, in partnership with The Economic Times — 750+ builders, buyers and decision-makers across four tracks that matter: GenAI & Agentic AI · Cloud & Platform Engineering · Cybersecurity in the AI Era · Deep Tech & Sustainability.

Not a day of slides — a day of proof: an Innovation Experience Zone with live demos across AI, robotics, cybersecurity, cloud and quantum; live build sessions; an AI Hackathon and Cybersecurity Challenge; and keynotes that close with a 90-day action plan you can take back to your team.

The evening turns gold with the Top 10 Women Tech Leaders India Awards 2026, now in its 4th edition — India's only annual awards programme dedicated to women technology leaders.

Early Bird 4,000 + GST until 15 August · Standard 5,000 + GST. Passes cover all four tracks, the Expo Hall and the Experience Zone.

GET YOUR PASS:

<https://race.reva.edu.in/racex360-season-2>

BECOME A SPONSOR:

<https://race.reva.edu.in/racex360-sponsorship>



RACEx360
Emerging Technology Conference 2026



In partnership with  **THE ECONOMIC TIMES**

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Echoes
- Perspectives
- The Bubble-Proofing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- **RACE Pulse**
- RACE Programs

RACE PULSE

EVENTS · COMMUNITY · WHAT'S COMING

NAAC
GRADE A+
ACCREDITED UNIVERSITY

REVA
RACE

REVA Academy for
Corporate Excellence

THE ENTERPRISE
GAVE AI THE KEYS
You get 90 minutes to prove it shouldn't have.

EC-Council
HACKai
Challenge
Prompt injection. RAG poisoning. Jailbreaking.
Agentic privilege escalation.

QUALIFIER
10th to 18th September 2026

FINALE
25th September 2026 at
The TAJ West End, Bengaluru

CASH PRIZE
₹1.75 Lakhs
+
CEH Vouchers

REGISTER TODAY

Registration closes on 10th September 2026

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Echoes
- Perspectives
- The Bubble-Proofing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- **RACE Pulse**
- RACE Programs

RACE PULSE

EVENTS · COMMUNITY · WHAT'S COMING

SEPTEMBER 2026 BATCH · APPLY BY 10 AUGUST 2026

Applications open for all six RACE Master's & PG Programmes — Business Analytics (#1 India), Artificial Intelligence (#3 India), Cybersecurity (#5 India), and Cloud Architecture & Security.

Weekend format. AICTE/UGC approved. Industry mentors from Amazon, Microsoft, AWS, EC-Council, Terralogic, IBM and more.

To know more, call:
+91 89040 58866

enquiry@race.reva.edu.in ·
Apply now: race.reva.edu.in

NEW BATCH **Certified Agentic AI Engineer** 26 SEPTEMBER 2026

Practical expertise in LLMs, AI agents, multi-agent systems, agentic workflows, governance and capstone-based implementation. Real-world use cases, deployable solutions.

To know more, call:
+91 89040 58866

Apply now:
<https://race.reva.edu.in/certified-agenticai-engineer>

RACE SANDBOX SATURDAY 19 SEPTEMBER 2026

Run your career simulation, live on campus. Sit in a live Master's class, walk the AI, Cybersecurity and Cloud labs, watch participants and alumni demonstrate real projects, and take a 1:1 on where your career goes next. Experience today, decide your tomorrow.

9:30 AM – 2:00 PM
RACE, REVA University,
Bengaluru

Reserve your slot:
+91 89040 58866 ·

enquiry@race.reva.edu.in ·
<https://race.reva.edu.in/>

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Echoes
- Perspectives
- The Bubble-Proofing Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- **RACE Programs**

Master's & PG Programmes

New Batch: September 2026 · 12–24 months ·
Weekend · For Working Professionals

#1 IN INDIA · ANALYTICS INDIA MAGAZINE

ANALYTICS

PG Diploma / M.Sc. in
BUSINESS ANALYTICS
AWS Academy · Microsoft Azure

12–24 Months

#3 IN INDIA · TECHGIG

ARTIFICIAL INTELLIGENCE

PG Diploma / M.Sc. in
ARTIFICIAL INTELLIGENCE
AWS Academy · Microsoft Azure

12–24 Months

#5 IN INDIA · ANALYTICS INDIA MAGAZINE

CYBERSECURITY

PG Diploma / M.Sc. in
CYBERSECURITY
EC-Council · Microsoft Azure · Terralogic.

12–24 Months

#3 IN INDIA · TECHGIG

ARTIFICIAL INTELLIGENCE

M.Tech. in
ARTIFICIAL INTELLIGENCE
AWS Academy · Microsoft Azure

24 Months

#5 IN INDIA · ANALYTICS INDIA MAGAZINE

CYBERSECURITY

M.Tech. in
CYBERSECURITY
EC-Council · Microsoft Azure · Terralogic

24 Months

CLOUD

PG Diploma / M.Sc. in
**CLOUD ARCHITECTURE &
SECURITY**
AWS Academy · Microsoft Azure

12–24 Months

Short-Term Certification Programmes

100 Hours · Project-Driven · Customisable for Organisations

AGENTIC AI | 4–6 weeks

AGENTIC AI SYSTEMS DESIGN

Multi-agent systems using LangChain, LangGraph, CrewAI. Capstone:
deployed autonomous agent.

AI ENGINEER | 8–12 WEEKS

AI ENGINEERING CERTIFICATION

Data pipelines, model training, LLMs, deployment. Real production
projects every batch.

CLOUD | 8–12 WEEKS

CLOUD ARCHITECTURE & DEVSECOPS

Multi-cloud, IaC, secure CI/CD on real AWS/Azure/GCP environments. Not
demos.

CYBERSECURITY | 8–12 WEEKS

CYBERSECURITY & SOC OPERATIONS

Threat detection, incident response, SIEM/SOAR. Hands-on in RACE's live
Security Operations Centre.

DATA & AI | 12 WEEKS

DATA SCIENCE & BUSINESS ANALYTICS

Statistics to ML to BI. Python, SQL, Power BI, Tableau on real business
datasets.

LEADERSHIP | 8–12 WEEKS

TECHNO-FUNCTIONAL LEADERSHIP

Executive presence, decision-making, leading tech-driven teams. Built for
mid-to-senior managers.

Enroll with us · race@reva.edu · race.reva.edu.in



REVA Academy for
Corporate Excellence

RETINA

AI · CYBERSECURITY · CLOUD · LEADERSHIP

THIS WEEK, ASK YOUR LEADERSHIP TEAM

“Of every rupee we’ll spend on AI this year, how much is tied to a business number we’re already tracking, and if the buildout corrects, which of our tools and vendors would we lose?”

REVA Academy for Corporate Excellence,
REVA University
Rukmini Educational Charitable Trust
Kattigenahalli, Yelahanka, Bangalore, Karnataka, India, 560 064
Phone: +91 89040 58866

RETINA PUBLISHES EVERY SUNDAY · FORWARD TO A CXO YOU RESPECT



REVA
UNIVERSITY
Bengaluru, India



Visit our website
race.reva.edu.in