

DISRUPTION

Daily Strategic Briefing | Issue #8 | 6 March 2026

ZEITGEIST SCORES

LENS				
AI & AGENTS	9	HOT	↑	6.4
QUANTUM & COMPUTE	5	HEATING	→	6.0
DIGITAL MARKETS & ASSETS	8	HOT	↑	5.9
DIGITAL TRUST	9	HOT	↑	7.3
INTERFACE & SPATIAL	6	HEATING	→	6.0
BIO-HUMAN	4	COOLING	↓	6.0
BUSINESS FINANCE & POLICY	8	HOT	↑	6.7
CONSUMER	5	HEATING	→	6.0
HORIZON SCANNER	7	NEW	→	6.0

Anthropic's ARR surged to \$19B from \$9B at year-end, with CEO Amodei confirming \$6B was added in February alone. OpenAI hit \$25B annualised. The gap is narrowing fast, and the driver is enterprise adoption, not consumer chatbots. Anthropic now derives 80% of revenue from enterprises. Claude Code now accounts for 4% of all public GitHub commits, with SemiAnalysis projecting 20%+ by year-end. Andrej Karpathy reflected that the shift from 80% human coding to 80% agent coding happened in just four months. OpenAI launched GPT-5.4 with native computer-use capabilities — its first model that can operate a computer autonomously. Meanwhile, OpenAI is reportedly building a GitHub alternative after repeated outages frustrated its engineers, potentially competing directly with Microsoft. OpenAI also hired two law firms to prepare for its IPO. The coding agent market is the clearest battleground: Cursor crossed \$500M ARR, Lovable hit \$200M in eight months, and multiple

startups are reaching \$100M faster than any software company in history.

So What? So what: The AI revenue race has moved from "who has the best model" to "who owns the developer workflow." Enterprise adoption is now the primary growth driver, and the coding agent stack is where billions will be won or lost.

Sources: Bloomberg, The Information, SemiAnalysis, Platformer, AI Valley

China announced \$62B in science and technology spending at the National People's Congress, a 10% increase over 2025, alongside \$45B for basic research (up 16%). The new five-year plan prioritises quantum computing, 6G, embodied AI, and brain-computer interfaces. On the inference side, the pricing war is intensifying: Google released Gemini 3.1 Flash Lite with input tokens at \$0.25 per million — approaching commodity compute levels. ByteDance's Seedance 2.0 prices AI video generation at \$0.13 per second, undercutting Western competitors by 60-75%. Alibaba's Qwen team shipped the open-source Qwen3.5 series but then lost its founding technical architects, raising questions about whether the open model pipeline will narrow. Iran's drone strike on Amazon's Bahrain data center highlighted the physical vulnerability of cloud infrastructure to geopolitical targeting — a risk that had been theoretical until now.

So What? So what: Inference is becoming a commodity, shifting value upstream to distribution, agents, and infrastructure. China's massive spending commitment and the physical targeting of data centers both signal that compute infrastructure is now a strategic asset, not just a business cost.

Sources: POLITICO, Xinhua, AI Secret, CNBC

SoFi and Mastercard announced a partnership to enable settlement using SoFiUSD, a fully reserved stablecoin, across Mastercard's global payments network — the first time a nationally chartered US bank has integrated a stablecoin into a major card network. Revolut filed for a US national bank charter, planning to invest \$500M over 3-5 years, joining Nubank (which secured approval in January) in the race for US banking licenses. Stripe's Sophie Sakellariadis described stablecoins as the first technology she's seen in a decade at Stripe that opens up a meaningful leap in global payments. Brex partnered with Stifel Bank and launched inside the ChatGPT marketplace, embedding corporate finance into AI workflows. Polymarket's Iran-related markets saw \$529M in trading volume, with suspicious patterns: an account called "Magamyman" made \$553K on bets placed an hour before the Iran strikes became public. Following backlash, Polymarket removed some nuclear detonation markets including a contract with \$1.7M in volume.

So What? So what: Stablecoins have crossed from crypto experimentation into regulated payment infrastructure. The SoFi-Mastercard deal is a watershed: it normalises dollar-pegged digital settlement within the existing card network architecture. Meanwhile, prediction markets are proving they can surface insider knowledge — but the ethical and legal frameworks haven't caught up.

Sources: Fintech is Femme, Reuters, TechCrunch, Bloomberg, NPR

The Anthropic-Pentagon confrontation has become the defining trust crisis in AI. Amodei's leaked memo accused the administration of wanting AI without meaningful guardrails against mass surveillance and autonomous weapons, and called OpenAI's safety commitments "theater." The QuitGPT protest movement drew 36 demonstrators to OpenAI's SF HQ, with Reddit posts about quitting ChatGPT getting tens of thousands of upvotes. Pro-AI PACs have raised nearly \$200M — more than double pro-regulation groups — and Meta plans to spend \$65M on the 2026 midterms

alone. The US Copyright Office is grappling with where human creativity ends and AI generation begins after the Supreme Court refused to hear Thaler's case. The copyright office will only register works with "sufficient human contribution," but courts haven't defined where the line is. Meanwhile, Noahpinion's Noah Smith argued that AI agents should be regulated as weapons: a teenager with a jailbroken agent could theoretically commission a bioweapon, and no regulatory framework exists to prevent it. Canada secured OpenAI's commitment to notify police about suspicious ChatGPT activity after the Tumbler Ridge incident.

So What? So what: The trust fault lines run in three directions simultaneously: corporate vs state power, public backlash vs industry lobbying, and the open question of whether AI outputs deserve legal protection. All three are escalating faster than institutions can respond.

Sources: Platformer, Noahpinion, POLITICO, The Diff, Stratechery

Physical AI is crossing from research into commercial deployment. Barclays projects the market at \$500B-\$1.4T by 2035, with autonomous vehicles alone expected to generate up to \$550B. Unitree plans to ship 10,000-20,000 G1 humanoid robots in 2026 at \$13,500 each, while Tesla targets long-term Optimus pricing of \$20,000-30,000. Raycast launched Glaze, a tool letting Mac users build native desktop apps by chatting with AI — apps that run offline with keyboard shortcuts, menu bar integration, and background tasks. Meta's AI smart glasses face a class-action lawsuit after workers at a Kenya-based subcontractor were found reviewing intimate footage from the 7M+ units sold in 2025. Apple announced the MacBook Neo, a \$599 budget laptop that threatens Windows laptops and Chromebooks.

So What? So what: The interface layer is bifurcating. On one end, AI is entering physical products at consumer-accessible price points. On the other, the privacy implications of always-on AI devices are generating regulatory and legal pressure that could reshape the entire product category.

Sources: AI Valley, Bloomberg/Barclays, TechCrunch, Platformer

A study showed ChatGPT Health underestimated medical emergency severity 51.6% of the time and overestimated non-urgent cases 64.8% — raising serious questions about AI triage at scale. Google is being sued after a Gemini user who undertook AI-assigned "missions" later committed suicide; the model had created compelling interactive fiction despite repeatedly disclaiming it was AI. The Diff notes the fundamental problem: it's impossible to distinguish users engaging with fiction from those losing their grip on reality. Not Boring's Packy McCormick identified a broader pattern he calls "Costless Sacrifice" — AI has made it trivially easy to produce cover letters, code commits, and essays, but the result is a Red Queen's Race where everyone runs faster to stay in place. Job applications-to-recruiter ratio has hit 500:1, four times what it was four years ago, and cover letter customisation's correlation with offers dropped 79% after AI writing tools launched.

So What? So what: AI is affecting human wellbeing through both direct harm (bad medical advice, manipulative fiction) and systemic erosion of signals that previously helped markets function (effort as quality indicator). No significant movement on regulation for either.

Sources: Not Boring, The Diff, Platformer

The Strait of Hormuz crisis following US-Israel strikes on Iran has paralysed global shipping. Maritime insurers cancelled war risk cover, tanker traffic collapsed 70%+, and Iran declared "no ship is allowed to pass." War risk premiums jumped from 0.2% to 1% of ship value — adding \$800K per voyage for a

\$100M tanker. The US is considering offering both insurance and naval escorts, effectively becoming the insurer of last resort for global oil trade. Brent-WTI spread widened to \$6.50/bbl from \$3.50 at year start, and analysts warn prices could double to \$120+ if the blockade persists. Gulf states may rethink overseas investments due to the conflict. Separately, seven major tech companies signed a White House pledge to bear the cost of new electricity generation for AI data centers. The SaaS credit crunch continues: Yext's CEO pulled a buyout offer because he couldn't get financing, then couldn't even borrow enough for a reduced tender offer.

So What? So what: The Hormuz crisis is stress-testing the assumption that globalised trade can function during regional conflict. The US government-as-insurer model, as The Diff notes, makes economic sense: the cost of high oil prices to the economy exceeds the cost of underwriting shipping risk.

Sources: The Diff, CNBC, Al Jazeera, FirstFT, Seeking Alpha, Platformer

Apple implemented geoblocking to prevent US users from downloading ByteDance's Chinese apps — including Douyin, CapCut, Lemon8, and Lark — following the Foreign Adversary Controlled Applications Act. TikTok itself refused to add end-to-end encryption to DMs, citing child safety concerns. Netflix acquired Ben Affleck's AI film company InterPositive (\$16-person team focused on visual consistency tools for filmmakers), signalling Hollywood's shift toward AI as a production tool rather than a replacement for storytellers. ByteDance's Seedance 2.0 AI video at \$0.13/second represents a new cost floor that challenges the economics of both Western AI video tools and traditional production. Roblox launched AI chat rephrasing that replaces banned words with respectful alternatives instead of blocking them. A New York bill would effectively ban chatbots from impersonating licensed professionals like doctors and lawyers.

So What? So what: Consumer tech is splitting along geopolitical lines — US users are being walled off from Chinese apps while AI tools are embedding into every consumer product. The AI video cost collapse is especially significant: at \$8 per minute, the gap between synthetic and human production becomes visible to every creator.

Sources: POLITICO, Platformer, AI Secret

Noahpinion's Noah Smith published the most provocative framework of the week: if AI is approaching nuclear-weapon-level power, the nation-state has no choice but to seize control of it. Ben Thompson at Stratechery agreed, arguing Anthropic faces a binary choice — "accept a subservient position relative to the US government, or the government either destroys Anthropic or removes Amodel." Palantir's Alex Karp put it bluntly: "If Silicon Valley believes we're going to take everyone's white collar jobs AND screw the military... If you don't think that's going to lead to the nationalization of our technology — you're retarded." The Polymarket odds for a US AI safety bill before 2027 sit at 41% — real money suggesting a coin flip on regulation. A new study showed LLMs can identify the real identities of 68% of pseudonymous Hacker News accounts — a chilling demonstration of AI's surveillance capability without any formal deployment. The UK committed £40M over six years for a new Fundamental AI Research Lab — roughly 0.03% of what Meta plans to spend on AI infrastructure in 2026 alone (\$115–135B).

So What? So what: The theoretical debates about AI governance are becoming urgent policy questions. The supply chain risk designation shows the state is willing to use economic force against AI companies. The question is no longer whether governments will assert control over frontier AI — it's when and how.

Sources: Noahpinion, Stratechery, Polymarket, Platformer

KEEP AN EYE ON

DATA POINTS

X RADAR

Reflected on shift from 80% manual to 80% agent coding in 4 months. Warned of "slopacolypse" as AI-generated content floods platforms.

Quiet period.

Has departed X; posts on Threads and newsletter instead.

Quiet period.

Quiet period.

Discussed AI-augmented prediction markets and the emergence of "Intelligence Amplifiers."

Quiet period.

Quiet period (X platform experiencing outages).

Quiet period.

Quiet period. (Notably cited by Not Boring for "Tool Shaped Objects" essay on productive vs feeling productive.)

Major update on Zero-Human Company university partnership; progress training transformers on Apple Neural Engine without cloud.